

1-31-2023

## Copyright Throughout a Creative AI Pipeline

Sancho McCann

*Peter A. Allard School of Law, University of British Columbia*

Follow this and additional works at: <https://digitalcommons.schulichlaw.dal.ca/cjlt>



Part of the [Computer Law Commons](#), [Intellectual Property Law Commons](#), [Internet Law Commons](#), [Privacy Law Commons](#), and the [Science and Technology Law Commons](#)

---

### Recommended Citation

Sancho McCann, "Copyright Throughout a Creative AI Pipeline" (2023) 19:1 CJLT 109.

This Article is brought to you for free and open access by the Journals at Schulich Law Scholars. It has been accepted for inclusion in Canadian Journal of Law and Technology by an authorized editor of Schulich Law Scholars. For more information, please contact [hannah.steeves@dal.ca](mailto:hannah.steeves@dal.ca).

# Copyright Throughout a Creative AI Pipeline

*Sancho McCann*\*

## INTRODUCTION

Consider the following fact pattern.<sup>1</sup>

Alex paints some original works on canvas and posts photos of them online. Becca downloads those images and uses them to train an AI (training configures the AI's model parameters to useful values). Becca posts the resulting trained parameter values on her website under a license that reserves to Becca the right to use the parameters commercially. Cory uses those parameter values in a program that is designed to produce artwork. Cory clicks *create* and the program produces a work. This work is new to Cory, but it looks a lot like one of Alex's original canvas images. Cory sells the work. Advise Cory about their potential copyright liability to Alex (for the substantially similar work that the program produced and that Cory subsequently sold) and to Becca (for taking Becca's parameters and using them commercially, contrary to the license).

Cory clicks *create* again. The program produces another work, this time quite different from any of Alex's original paintings. Cory shares new work on Instagram. Danny copies this image from Cory's Instagram feed and sells a bunch of postcards that feature that image. Advise Danny about their copyright liability to Cory.

These scenarios are not as contrived as they might initially seem. People frequently use copyrighted works when training an AI (more precisely: when training an AI's parameters).<sup>2</sup> The resulting trained parameters *are* being shared under licences that assume the parameters are the subject of copyright. People do use these parameters in programs that can produce novel content.<sup>3</sup> The resulting

---

\* PhD (Computer Science), University of British Columbia, 2014. JD Candidate, Peter A. Allard School of Law, 2021. Sancho would like to thank Professor Jon Festinger, Q.C. for many helpful discussions while supervising this work and Professor Graham Reynolds for valuable feedback on an earlier draft.

<sup>1</sup> This mode of introduction is inspired by William J Braithwaite, "Derivative Works in Canadian Copyright Law" (1982) 20:2 Osgoode Hall LJ 191.

<sup>2</sup> "The training process can involve reproductions of the training data . . . It is unclear whether the use of copyrighted works for training an AI system is considered copyright infringement if the . . . copyright owner's permission is not obtained." House of Commons, Standing Committee on Industry, Science and Technology, *Evidence*, 42-1, No 114 (11 May 2018) at 1625 (Maya Medeiros), cited in House of Commons, *Statutory Review of the Copyright Act: Report of the Standing Committee on Industry, Science, and Technology* (June 2019) (Chair: Dan Ruimy) at 85 [INDU, *Statutory Review*].

<sup>3</sup> *This footnote was generated by a creative AI, seeded only with these first fifteen words. I*

work can be quite surprising to the end-user and there are generally no checks in place to ensure that the new works do not take too directly from the original training data. However, many of the new works will be quite different from any content already in the world. And the end-users of the creative program often claim copyright ownership over the resulting novel work.<sup>4</sup>

These are real issues—a fact reflected in the explosion of articles and international attention being devoted to the topic.<sup>5</sup>

I will first present the training and use of a creative program based on a neural network, a popular model that forms the basis of state-of-the-art creative AIs. Then, I will examine each of the issues just raised:

1. Does the person managing the automatic training of a neural network's parameters obtain a copyright in the resulting trained parameters?
2. Does a person using a program that produces artistic output obtain a copyright in that output?
3. The automatic training of a neural network requires large amounts of example data (a *training set*). Can images from around the internet be copied for the purpose of training a neural network?
4. What if a person uses an AI to produce a work that looks substantially similar to one of the training examples? Is that an infringement? And who is infringing?

On the first and second issues, I conclude that under current Canadian copyright law, it will almost always be the case that nobody will hold the copyright to the algorithmically trained parameters. However, works produced by *using* these trained neural networks will often, but not always, attract copyright protection.

---

should clarify a bit, though. This was merely an experiment; it's an analysis of my favorite video game, the Super Mario Bros. trilogy. The results of that experiment have shown that AI does not play by a formula. It is more humanlike than you might think. That said, it would be nice if AI was more humanlike. So how to change our AI? Why would we want to? Let's start with a classic argument: "if machines can't do a job, why should we?"

<sup>4</sup> See e.g. Bhautik Joshi, Kristen Stewart & David Shapiro, "Bringing impressionism to life with neural style transfer in *Come Swim*" in Stephen N Spencer, ed, *DigiPro '17: Proceedings of the ACM SIGGRAPH Digital Production Symposium, Los Angeles, California, July 2017* (New York: Association for Computing Machinery, 2017), DOI: < 10.1145/3105692.3105697 > .

<sup>5</sup> See e.g. WIPO, "WIPO Begins Public Consultation Process on Artificial Intelligence and Intellectual Property Policy", PR/2019/843, (13 December 2019); Carys Craig & Ian Kerr, "The Death of the AI Author" (2021) 52:1 Ottawa L Rev 31; James Grimmelmann, "There's No Such Thing as a Computer-Authored Work—And It's a Good Thing, Too" (2016) 39:3 Colum J L & Arts 403, DOI: < 10.7916/jla.v39i3.2079 > ; Bruce E Boyden, "Emergent Works" (2016) 39:3 Colum J L & Arts 377, DOI: < 10.7916/jla.v39i3.2077 > ; Enrico Bonadio & Luke McDonagh, "Artificial Intelligence as Producer and Consumer of Copyright Works: Evaluating the Consequences of Algorithmic Creativity" in Margaret Llewelyn, ed, *Intellectual Property Quarterly* (London, UK: Sweet & Maxwell, 2020) 112; Daniel J Gervais, "The Machine as Author" (2020) 105:5 Iowa L Rev 2053; Pamela Samuelson, "AI Authorship?" (2020) 63:7 Communications ACM 20, DOI: < 10.1145/3401718 > .

This distinction is normatively justified, on the basis that the purpose of copyright in Canada is to provide balanced protection of an author’s expression of skill and judgment and because there are technological means to keep one’s trained parameters secret even while allowing others access for use.

Third, I conclude, in agreement with the *Statutory Review of the Copyright Act*,<sup>6</sup> that copying existing works for the purpose of building a training set is *prima facie* infringement and that Canada should clarify that this is a purpose allowed under Canada’s fair-dealing user’s right. To include this purpose under the fair-dealing user’s right would avoid chilling educational and research activity that is dependent on large collections of example works.

Finally, I conclude that it is open for the output of a creative AI to be an infringing work and that the burden is properly on the trainer to ensure they have not created an infringement machine. This is exactly where that burden should be placed, as the trainer has more information and is the least-cost avoider.

## 1. A GENERAL MODEL

In this Section, I present a general model for how today’s state-of-the-art artificial intelligence (AI) produces (or is used to produce) visual or written works<sup>7</sup> like images or prose. While the works produced by this technology might be art-like or may potentially qualify as art, I avoid calling them art at this stage because such a categorization may be seen to assume the conclusion that this AI activity is creative or original in a way that attracts copyright. Also, while I may talk about AIs as *producing* these works, it may be more useful to view this activity as a person *using the AI* to produce the works. This distinction will be relevant for my analysis of who obtains copyright in the eventual output.

A neural network is a kind of computational model that transforms input data into output data through a series of sequential transformations. This can be as simple as classifying a 20-pixel by 20-pixel black-and-white image of a handwritten digit as one of ten Arabic numerals. The input in that case would be the  $20 \times 20$  image. The output would be the classification as a 0, 1, 2, . . . or 9. To give a more complex example, the input could be pixels sampled from a random distribution (noise, essentially) and the output could be a portrait of a person who has never actually existed.<sup>8</sup> A neural network like this underlies almost every AI that produces art-like works today.

Typical neural-network architectures used in these processes contain between 150 million and 1.5 billion parameters.<sup>9</sup> Each parameter is a single, real-valued

<sup>6</sup> See INDU, *Statutory Review*, *supra* note 2.

<sup>7</sup> While I have used the term-of-art *works* to describe these artifacts, I use this term without assuming that any work was done (in the labour sense) or that these artifacts are *works* that would properly be the subject of copyright.

<sup>8</sup> See e.g. Philip Wang, “This Person Does Not Exist”, online: *This Person Does Not Exist* < [thispersondoesnotexist.com](http://thispersondoesnotexist.com) > .

number.<sup>10</sup> A neural network that is untrained—one whose parameters are set to random values—will typically not exhibit any useful behaviour: its output will not appear to be connected with the input in any meaningful way. For the operations making up a neural network to be useful, the network must be trained. Training adjusts the parameters of the network to values that cause the network to do something useful. And the same neural network can be trained for different tasks; the difference in behaviour is captured in the trained parameters.

Important for the legal analysis that will follow is the disconnect between the choices a human makes and the resulting trained parameter values or AI behaviour. The human will have chosen the training data, the network architecture, and the training algorithm. However, there are many configurations of the parameters for which the difference between them is “in practice irrelevant.”<sup>11</sup> The person managing the training of the neural network is generally unconcerned with the particular parameter configuration that arises, and it is actually difficult to get insight into the significance of any particular parameter value.<sup>12</sup> These neural networks are black boxes, and what it would mean for such a network to be interpretable is often poorly defined.<sup>13</sup>

Even when the human who is managing the training selects a particular goal for an AI, the manner in which the AI learns to achieve that goal can be surprising.<sup>14</sup> When tasked to play a video game, for instance, an AI will discover

<sup>9</sup> See Yann LeCun, Yoshua Bengio & Geoffrey Hinton, “Deep Learning” (2015) 521:7553 *Nature* 436 at 436, 440, DOI: <10.1038/nature14539>. See e.g. Leon A Gatys, Alexander Ecker & Matthias Bethge, “A Neural Algorithm of Artistic Style” (2015) [unpublished, archived at arXiv, Cornell University] (neural style transfer), online (pdf): *arXiv* <arxiv.org/abs/1508.06576>; Christian Szegedy et al, “Inception-v4, Inception-ResNet and the Impact of Residual Connections on Learning” (2016) [unpublished, archived at arXiv, Cornell University] (inception), online (pdf): *arXiv* <arxiv.org/abs/1602.07261>; Karen Simonyan & Andrew Zisserman, “Very Deep Convolutional Networks for Large-Scale Image Recognition” (2015) [unpublished, archived at arXiv, Cornell University] (VGG-19), online (pdf): *arXiv* <arxiv.org/abs/1409.1556>.

<sup>10</sup> By *real-valued number*, I mean, roughly, a *number* in the broadest sense that you are likely familiar with (e.g. 0, 1, 2, -1, -2, 0.1, -0.1, 0.3, 1/3, 0.5843614220), subject to the limited precision and range available within the digital representation of a computer.

<sup>11</sup> Anna Choromanska et al, “The Loss Surfaces of Multilayer Networks” (2015) 38 *Proceedings Machine Learning Research* 192 at 192.

<sup>12</sup> See Andrej Karpathy, Justin Johnson & Li Fei-Fei, “Visualizing and Understanding Recurrent Networks” (Paper delivered in the Workshop Track of the International Conference on Learning Representations, San Juan, Puerto Ric, 3 May 2016) [unpublished] at 1, online (pdf): *Stanford Vision and Learning Lab* <vision.stanford.edu/pdf/KarpathyICLR2016.pdf>.

<sup>13</sup> See generally Cynthia Rudin, “Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead” (2019) 1 *Nature Machine Intelligence* 206.

<sup>14</sup> See Joel Lehman et al, “The Surprising Creativity of Digital Evolution: A Collection of Anecdotes from the Evolutionary Computation and Artificial Life Research Communities” (2018) [unpublished, archived at arXiv, Cornell University], online (pdf): *arXiv* <arxiv.org/abs/1803.03453>.

and learn to exploit bugs that were previously unbeknownst to human players. Generative adversarial models (in which a generative neural network is tasked to fool a classifier) somehow learn to generate new, previously unseen examples, rather than merely producing near matches to the training examples.<sup>15</sup>

Given this backdrop, I will now embark on an analysis of how copyright applies throughout this AI-infused pipeline by focusing on the four issues presented above.

## 2. COPYRIGHT IN TRAINED PARAMETERS

The set of trained parameters is a list of real numbers,<sup>16</sup> each corresponding to a multiplier that is associated with a particular connection in the architecture of a specified neural network. In this section, I argue that these numbers—the trained parameters—will generally not be the subject of copyright. The parameters could be conceived of as facts, as a compilation of facts, or as a computer program. If facts, they would not be protected. But courts have construed this category—mere facts—narrowly. Other data has escaped categorization as mere fact by analogy to imaging.<sup>17</sup> I argue that the imaging analogy is inapplicable in this scenario. However, even if the parameters are mere facts, another path to copyrightability is as a compilation. I argue that this path is foreclosed due to insufficient *originality* as required in Canadian copyright law. Insufficient originality also forecloses the possibility that these parameters would be protected as a computer program. The conclusion that the automatically trained parameters are not protected by copyright is consistent with the justifications underlying Canadian copyright law.

That the parameters are a list of numbers does not help determine whether the parameters can be protected by copyright; what matters is what they represent. A digital image is also a list of numbers,<sup>18</sup> but that does not affect whether a digital image is protected by copyright. What matters in the analysis is

<sup>15</sup> See Vaishnavh Nagarajan, Colin Raffel & Ian Goodfellow, “Theoretical Insights into Memorization in GANs” (Paper delivered at the Thirty-Second Conference on Neural Information Processing Systems, Montreal, Canada, 8 December 2018) [unpublished], online (pdf): *Carnegie Mellon University, School of Computer Science* < [www.cs.cmu.edu/~vaishnan/papers/GAN\\_memorization.pdf](http://www.cs.cmu.edu/~vaishnan/papers/GAN_memorization.pdf) > .

<sup>16</sup> See note on the meaning of *real-valued number*, *supra* note 10.

<sup>17</sup> See e.g. *Geophysical Service Inc. v. Encana Corp.*, 2016 ABQB 230, 2016 CarswellAlta 742 (Alta. Q.B.), affirmed 2017 CarswellAlta 732 (Alta. C.A.), leave to appeal refused *Geophysical Service Incorporated v. EnCana Corporation, et al.*, 2017 CarswellAlta 2545, 2017 CarswellAlta 2546 (S.C.C.) [*Geophysical*].

<sup>18</sup> Each pixel in a digital image is represented by a single number (in the case of a greyscale image) or triplet of numbers (in the case of a colour image). For example, Grey is represented by (84, 84, 84) and Cadet Blue is represented by (95, 158, 160). In a greyscale image, the single number represents the brightness of the pixel. In a colour image, the triplet of numbers can represent, for instance, the mix of red, green, and blue (RGB) in the pixel. There are many variations on this, but the point is that a digital image is a list of numbers.

the connection (or lack thereof) between the representation and the purported author's skill and judgment.<sup>19</sup> The author of a digital image can express, in terms of the intended content, what they intend the numbers to be. For example, while an author using an image-editing program may not know that they wanted a particular pixel to have RGB values<sup>20</sup> equal to (123, 216, 230), the author can nonetheless express that they wanted that part of the image to be light blue or that they intended the image to capture a particular viewpoint of a scene in the world. The fact *light blue* ends up encoded digitally as numeric values is just a consequence of the digital translation. However, the numbers representing a neural network's parameter values are qualitatively different. No purported author<sup>21</sup> in the training pipeline can express their intention for the values of the neural network parameters.

In analyzing whether these trained parameters are protected by copyright, I present three analyses: parameters as facts, parameters as a compilation of facts, and parameters as a computer program.

#### (a) Parameters as Facts

This first option can be dispensed with quickly: Canadian copyright law does not protect mere facts.<sup>22</sup> But this does not ultimately take us far, given how narrowly courts understand something to be a mere fact. In *Geophysical Service Incorporated*,<sup>23</sup> Justice Eidsvik held that seismic imaging data was not a mere fact. The seismic imaging data was an "expression of GSI's views of what the image of the subsurface of the surveyed areas represents,"<sup>24</sup> but "[t]he facts themselves, the rocks, are still at the bottom of the sea available for anyone else to survey."<sup>25</sup>

Through this lens, the trained parameters could be seen as *the facts themselves*. The fact that the neural network's parameters converged to some particular values when trained on a particular training set is what gets fixed in the trained model. These parameters are not an expression of the views of a human

---

<sup>19</sup> See *CCH Canadian Ltd. v. Law Society of Upper Canada*, 2004 SCC 13, 2004 CarswellNat 446, 2004 CarswellNat 447 (S.C.C.) [*CCH SCC*].

<sup>20</sup> See note on digital images, *supra* note 18.

<sup>21</sup> Potential authors include the person who wrote the software to train the neural network, the person who designed the utility function and training set, and the person who launched and monitored the training process.

<sup>22</sup> See *Maltz v. Witterick*, 2016 CF 524, 2016 FC 524, 2016 CarswellNat 1498, 2016 CarswellNat 11536 (F.C.) [*Maltz*] ("[f]acts are facts; and no one owns copyright in them no matter what their relative size or significance" at para 31). *Maltz* concerned whether certain real-world facts expressed in a documentary could be protected by copyright. See also *CCH SCC*, *supra* note 19 ("in Canada, as in the United States, copyright protection does not extend to facts" at para 22).

<sup>23</sup> *Geophysical*, *supra* note 17.

<sup>24</sup> *Ibid* at para 97.

<sup>25</sup> *Ibid* at para 97.

author. They are not capturing an approximation of a reality external to themselves. If a court were to accept the view that the trained parameters are facts, those parameters would be ineligible for copyright protection, like rocks at the bottom of the sea.

But Justice Eidsvik acknowledges two alternative conceptions that would result in the data (seismic data, in that case) being a proper subject of copyright: 1) the data might be an original expression of an author (like a photograph), despite it resulting from facts that exist in the real world; or 2) the data might be an original *compilation* of facts.<sup>26</sup> Originality is the crux of the analysis no matter which of these alternatives is adopted.

As I have started to argue already, the first path of escape from categorization as mere fact—the analogy to a photograph—is not available in the neural-network scenario. In *Geophysical*, Justice Eidsvik found it significant that the ultimate seismic data was an attempt to depict subsurface structures.<sup>27</sup> The result was likened to a “map, plan, or chart”—a result of “recordings, from the geology of the subsurface.”<sup>28</sup> In the case of trained parameters, there is no external truth that any human author (or even the training algorithm) is attempting to approximate or to express a view about. The training algorithm merely discovers parameter values that produce some desired functionality.

A second path to copyrightability would be to treat the parameters as a compilation of facts. This requires a more fulsome analysis of the *Copyright Act* and originality, which I present in the following section.

### (b) Parameters as a Compilation—Originality

Individual facts are ineligible for copyright protection, but when facts are aggregated in a compilation, that compilation can attract copyright protection.<sup>29</sup> In this section I describe how Canada’s originality standard for copyright protection properly privileges human interaction and expression—human elements that are not reflected in the parameters of neural network.<sup>30</sup>

Courts have long interpreted the *Copyright Act* to extend copyright protection to compilations of facts—not as facts *per se*, but in their selection and arrangement.<sup>31</sup> This interpretation is now codified in the *Copyright Act*.<sup>32</sup>

<sup>26</sup> See *ibid* at para 33—36, 97—98.

<sup>27</sup> *Ibid* at para 97.

<sup>28</sup> *Ibid* at para 76.

<sup>29</sup> See *Copyright Act*, R.S.C. 1985, c. C-42, s. 2. See e.g. *British Columbia Jockey Club v. Standen (Winbar Publications)*, 1985 CarswellBC 274 (B.C. C.A.) [*BC Jockey Club*] (a compilation of facts relating to horse races was protected by copyright); *CCH SCC*, *supra* note 19 (the compilation of a “headnote and the accompanying edited judicial reasons are ‘original’ works covered by copyright” at para 33); *Geophysical*, *supra* note 17.

<sup>30</sup> See generally Craig & Kerr, *supra* note 5. Craig & Kerr present a complementary and enlightening critique of AI authorship, but avoid “mud wrestling with copyright’s originality threshold” (*ibid* at 44). Here, I get into the mud.



The *Copyright Act* today provides for copyright in “every original literary, dramatic, musical and artistic work.”<sup>33</sup> This includes compilations.<sup>34</sup> *Compilation* further includes “a work resulting from the selection or arrangement of data.”<sup>35</sup>

*Originality* is the crux of this definition. Prior to 2004, Canadian courts oscillated in their conception of originality in copyright law, moving between a “romanticization of the author-figure”<sup>36</sup> (which assumes *creativity* in a strong sense, “ingenuity,”<sup>37</sup> or a reflection of the “personalities or individuality”<sup>38</sup> of authors) and a mere reward-for-labour or *sweat-of-the-brow* justification for copyright protection.<sup>39</sup> The sweat-of-the-brow standard (independent labour that results in non-copied expression) is a low threshold for copyright protection; it would protect mere production. Both the Canadian and US Supreme Courts have now rejected the sweat-of-the-brow approach to originality.<sup>40</sup> The US has adopted the “modicum of creativity” standard for originality.<sup>41</sup> The originality standard in Canada is that the expression must be a “product of an author’s exercise of skill and judgment” that is not “so trivial that it could be characterized as a purely mechanical exercise.”<sup>42</sup> The Court explicitly intended

<sup>31</sup> See e.g. *BC Jockey Club*, *supra* note 29.

<sup>32</sup> See *Copyright Act*, *supra* note 29, s. 2. The statutory definition of *compilation*, and the inclusion of *compilation* in the definition of “every original literary, dramatic, musical, and artistic work” was added in *North American Free Trade Agreement Implementation Act*, S.C. 1993, c. 44, s. 53.

<sup>33</sup> *Copyright Act*, *supra* note 29, s. 5(1).

<sup>34</sup> *Ibid.*, s. 2 (“[e]very original literary, dramatic musical and artistic work includes every original production in the literary, scientific or artistic domain, whatever may be the mode or form of its expression, such as compilations . . . [and] translations”).

<sup>35</sup> *Ibid.*, s. 2.

<sup>36</sup> Carys J Craig, “The Evolution of Originality in Canadian Copyright Law: Authorship, Reward and the Public Interest” (2005) 2:2 Ottawa L Tech J 425 at 432, n 27 [Craig, “Evolution of Originality”]. See e.g. *Tele-Direct (Publications) Inc. v. American Business Information Inc.*, 1997 CarswellNat 2752, 1997 CarswellNat 2111 (Fed. C.A.), leave to appeal refused 1998 CarswellNat 3212 (S.C.C.) [*Tele-Direct*].

<sup>37</sup> Craig, “Evolution of Originality”, *supra* note 36 at 432, citing *CCH Canadian Ltd. v. Law Society of Upper Canada*, 1999 CarswellNat 2123, 1999 CarswellNat 3080 (Fed. T.D.) at para. 111, additional reasons 2000 CarswellNat 168, 2000 CarswellNat 5102 (Fed. T.D.), reversed in part 2002 CarswellNat 1000, 2002 CarswellNat 2841 (Fed. C.A.), reversed 2004 CarswellNat 446, 2004 CarswellNat 447 (S.C.C.), quoting *Tele-Direct*, *supra* note 36 at para 29.

<sup>38</sup> Craig, “Evolution of Originality”, *supra* note 36.

<sup>39</sup> See e.g. *CCH Canadian Ltd. v. Law Society of Upper Canada*, 2002 CAF 187, 2002 FCA 187, 2002 CarswellNat 1000, 2002 CarswellNat 2841 (Fed. C.A.), reversed 2004 CSC 13, 2004 SCC 13, 2004 CarswellNat 446, 2004 CarswellNat 447 (S.C.C.).

<sup>40</sup> See *CCH SCC*, *supra* note 19 (rejecting that “labour, in and of itself, could ground a finding of originality” at para 21); *Feist Publications Inc. v. Rural Telephone Service Co.*, 111 S.Ct. 1282 (1991) [*Feist*].

<sup>41</sup> *Ibid.*

this to be a lower threshold than creativity: “creativity is not required to make a work ‘original’.”<sup>43</sup>

To help understand what this skill-and-judgment standard privileges, we can find additional clarification in the earlier decision in *Théberge*.<sup>44</sup> In *Théberge*, the Court articulated its fullest conception of the interests in balance in Canadian copyright law. They express a view of copyright that is based in economic rights. Copyright strikes a balance between economic rights of the copyright owner and proprietary rights of a person in possession of the work—between encouraging the “dissemination of works of the arts and intellect” and giving “due weight” to the “limited nature” of the incentives.<sup>45</sup> The majority in *Théberge* limits the protection of “personality” interests to the domain of moral rights.<sup>46</sup> Given that the Court has placed personality interests in the domain of moral rights, the skill-and-judgment originality threshold is thus unlikely intended to protect those personality interests. The skill-and-judgment threshold nonetheless incentivizes human intellectual involvement in the production of a work. It results in “compensation for the exercise of ‘skill and judgment’ in the production of a work of art and intellect, and not an entitlement flowing from the investment of *labour* or *personality*.”<sup>47</sup>

For John Dewey, this interaction is essential preparation for the creation of art. Artistic expression is the result of a long period of preparation and practice: preparation to be receptive to emotions and experience, and preparation to be able to work with a medium (clay, paint, written word, spoken word).<sup>48</sup> All of this preparation sets the stage for ostensible spontaneity and aesthetic expression: “New ideas come leisurely yet promptly to consciousness only when work has previously been done in forming the right doors by which they may gain entrance.”<sup>49</sup> The nature of expression calls for us to incentivize people to practice this expressive skill and judgment so they can better share their stories and at the same time be more receptive to those of others. This interaction gives both the artistic medium *and* the author “form and order they did not at first possess.”<sup>50</sup>

In requiring the expression of skill and judgment as a prerequisite for copyright, Canada’s originality threshold privileges this “pragmatist aesthetic.”<sup>51</sup>

<sup>42</sup> *CCH SCC*, *supra* note 19 at para 25.

<sup>43</sup> *Ibid.*

<sup>44</sup> *Galerie d’art du Petit Champlain inc. c. Théberge*, 2002 CSC 34, 2002 SCC 34, 2002 CarswellQue 306, 2002 CarswellQue 307 (S.C.C.) [*Théberge*].

<sup>45</sup> *Ibid* at paras 30—31.

<sup>46</sup> *Ibid* at para 59.

<sup>47</sup> Carys J Craig, “Resisting ‘Sweat’ and Refusing *Feist*: Rethinking Originality after *CCH*” (2007) 40:1 UBC L Rev 69 at 119—120 (emphasis added) [Craig, “Resisting ‘Sweat’ and Refusing *Feist*”].

<sup>48</sup> See John Dewey, *Art as Experience* (New York: Penguin, 2005) at 102 (originally published 1934).

<sup>49</sup> *Ibid* at 76.

<sup>50</sup> *Ibid* at 67—68.

It is not sweat, labour, or creation that matters, but instead *aesthetic* work: expression of one's skill and judgment. The spectrum of potential originality thresholds has at the low end: mere production that is not a copy; and at the other end: "creativity" evincing a "true and deserving 'author'."<sup>52</sup> Canada has selected a middle ground that incentivizes intellectual labour while not requiring the "imprint of personality"<sup>53</sup> on a work.

This motivation underlying copyright has also been suggested by Jane Ginsburg: "in copyright law, an author is (or should be) a human creator who . . . succeeds in exercising minimal personal autonomy in her fashioning of the work. Because, and to the extent that, she moulds the work to her vision (be it even a myopic one), she is entitled [to copyright]."<sup>54</sup> It is the *exercise* of personal autonomy as part of *creating* the work that is significant, not any assumption that unique *individuality* has been fixed in the resulting work itself (although, that may very well result).

Despite being human-centered, this Deweyan, pragmatic conception of originality does not assume the existence of a romantic author-figure. It merely recognizes the value in having people work through the art-making process—to develop creative skill and judgment by *using* creative skill and judgment.<sup>55</sup> The "modicum of creativity" standard suggests that in the US, copyright at least envisions something more akin to novelty.<sup>56</sup> The Canadian standard ("skill and judgment . . . [not so] trivial that it could be characterized as a purely mechanical exercise"), by its own words does not require *creation*, in any sense other than production, and does not assume a romantic author.<sup>57</sup> The Canadian standard qualifies what *kind* of labour is required and the nature of the connection between that labour and the fixed expression.

*CCH Canadian Ltd. v. Law Society of Upper Canada*, which established Canada's current originality standard, also happened to be about compilations.

<sup>51</sup> This is claimed by Barton Beebe to be a more justifiable value for copyright incentives: "a notion of aesthetic progress that privileges active aesthetic *work* over objectified aesthetic *works*, everyday aesthetic practice over timeless aesthetic achievements, and seeks not so much more artworks, however fine, but more artists, fine or otherwise." Barton Beebe, "Bleistein, The Problem of Aesthetic Progress, and the Making of American Copyright Law" (2017) 117:2 Colum L Rev 319 at 347.

<sup>52</sup> Craig, "Resisting Sweat and Refusing *Feist*", *supra* note 47 at 106.

<sup>53</sup> *Ibid* at 103.

<sup>54</sup> Jane C Ginsburg, "The Concept of Authorship in Comparative Copyright Law" (2003) 52:4 DePaul L Rev 1063 at 1064.

<sup>55</sup> See generally Gervais, *supra* note 5 at 2061 ("human progress should serve as a normative guidepost"); Craig & Kerr, *supra* note 5 at 45 ("demise of romantic authorship should also spell the death of the AI author").

<sup>56</sup> "Creativity" has various connotations (production, subjective novelty, absolute novelty, a civil-law romantic-author connotation) and it is unclear which of these the SCC was rejecting when they rejected the language from *Feist*. See Craig, "Resisting 'Sweat' and Refusing *Feist*", *supra* note 47 at 99.

<sup>57</sup> *CCH SCC*, *supra* note 19 at para 25.

The Court recognized that compilations may be sufficiently original to attract copyright protection.<sup>58</sup> This does not extend copyright to the individual components or facts that are part of the compilation; it is the “over-all arrangement of them,” which is the subject of copyright.<sup>59</sup> So, selection and arrangement of data can attract copyright protection as a compilation, as long as they express a sufficient degree of the purported author’s skill and judgment.

Next, I present a group of cases that apply this test in the context of digital databases. Digital databases are compilations of data and a close analogy to the trained parameters of a neural network. These cases give guidance regarding the amount of skill and judgment required in order to obtain copyright in such a compilation. They also reveal that works whose content and form are driven by utility and function do not meet this originality threshold.

*Toronto Real Estate Board v. Commissioner of Competition*<sup>60</sup> examined whether the TREB had copyright in database of property listings. The TREB maintained a database of property listings in the Toronto area. TREB members provided the information and a third party corrected it to adhere to quality standards. The database was “in line with industry norms across Canada.”<sup>61</sup> The Federal Court of Appeal noted that drawing the line between sufficient skill and judgment and a “mere mechanical exercise” is difficult in the case of compilations.<sup>62</sup> But, several criteria are helpful. They look to the merger doctrine: “when an idea can be expressed in only a limited number of ways, then its expression is not protected as the threshold of originality is not met.”<sup>63</sup> And they look to industry standards: as “a legitimate, residual consideration.”<sup>64</sup> However, reflecting *CCH*, this factor is not determinative: “there is no steadfast rule that ‘there is no entitlement to copyright protection. . . where the selection or arrangement is directed by accepted and common industry practices.’”<sup>65</sup>

In *Distrimed Inc. v. Dispill Inc.*<sup>66</sup> the Federal Court emphasized that the threshold for originality is the same for a compilation as for any other work.<sup>67</sup>

---

<sup>58</sup> See *ibid* at para 33.

<sup>59</sup> *Ibid*.

<sup>60</sup> *Toronto Real Estate Board v. Commissioner of Competition*, 2017 CAF 236, 2017 FCA 236, 2017 CarswellNat 6861, 2017 CarswellNat 8751 (F.C.A.), leave to appeal refused 2018 CarswellNat 4555, 2018 CarswellNat 4556 (S.C.C.) [*Toronto Real Estate*].

<sup>61</sup> *Ibid* at para 33.

<sup>62</sup> *Ibid* at para 185.

<sup>63</sup> *Ibid* at para 187, citing *Red Label Vacations Inc. v. 411 Travel Buys Ltd.*, 2015 CF 18, 2015 FC 18, 2015 CarswellNat 427, 2015 CarswellNat 4554 (F.C.) at para. 98, affirmed 2015 CarswellNat 7643, 2015 CarswellNat 12245 (F.C.A.), citing *Delrina Corp. v. Triolet Systems Inc.*, 2002 CarswellOnt 633 (Ont. C.A.) at paras. 48-52, additional reasons 2002 CarswellOnt 3220 (Ont. C.A.), leave to appeal refused 2002 CarswellOnt 4080, 2002 CarswellOnt 4081 (S.C.C.) [*Delrina ONCA*].

<sup>64</sup> *Toronto Real Estate*, *supra* note 60 at para 188.

<sup>65</sup> *Ibid* at para 189, citing *Geophysical*, *supra* note 17 at paras 100—101 (relating to seismic data).

Justice de Montigny recognized that “when the content and layout of a form is largely dictated by utility and/or legislative requirements, it is not to be considered original.”<sup>68</sup> When training a neural network, there is literally a fitness function or a “utility function” that guides the updates of the parameters. The entire purpose of the automatic parameter configuration is to find a parameter setting that allows the neural network to perform better according to the fitness function.

This last aspect of originality (that only an original expression in a work that is severable from that which is dictated by utility) has both statutory and common-law sources.

In Canada, section 64(1) of the *Copyright Act* deems it not an infringement to reproduce “any features of [an] article that are dictated solely by a utilitarian function of the article.”<sup>69</sup> This is a statutory declaration of non-infringement for certain kinds of copying. However, it only applies to *articles*: “any thing that is made by hand, tool or machine.”<sup>70</sup> This section seems to conceive of these articles as physical things, capable of having “shape, configuration, pattern or ornament.”<sup>71</sup> Based on the text and the scant case law interpreting this section, it seems unlikely that weights of a neural network (either as a compilation or as a computer program as discussed below) would be treated as an article for the purpose of this section.<sup>72</sup> Section 64.2(2) even explicitly sets out that “incorporation of a computer program into an integrated circuit product or the incorporation of a work into such a computer program may constitute an infringement of the copyright or moral rights in a work.”<sup>73</sup> Thus, if copying a portion of a computer program that is dictated solely by a utilitarian function is exempted from infringement, such an exemption is likely not due to these statutory sections.

However, regardless of whether sections 64ff of the *Copyright Act* exempt function-driven features of a computer program or compilation from infringement, the common-law interpretation of originality has already exempted function-driven expression from copyright protection. *Delrina Corp.*

---

<sup>66</sup> *Distrimedic Inc. v. Dispill Inc.*, 2013 CF 1043, 2013 FC 1043, 2013 CarswellNat 3663, 2013 CarswellNat 5722 (F.C.) [*Distrimedic*].

<sup>67</sup> See *ibid* at para 321.

<sup>68</sup> *Ibid* at para 324, cited with approval in *Toronto Real Estate*, *supra* note 60 at para 191.

<sup>69</sup> *Copyright Act*, *supra* note 29, s. 64.1(1)(b).

<sup>70</sup> *Ibid*, s. 64(1).

<sup>71</sup> *Ibid*.

<sup>72</sup> I identified seventeen Canadian cases on CanLII that refer to s. 64 and all the works at issue are physical works. See e.g. *Lainco inc. c. Commission scolaire des Bois-Francis*, 2017 CF 825, 2017 FC 825, 2017 CarswellNat 5087, 2017 CarswellNat 9467 (F.C.), additional reasons 2018 CarswellNat 761, 2018 CarswellNat 941 (F.C.) (an arched roof of a soccer stadium); *Import-Export René Derhy (Canada) inc. c. Magasins Greenberg ltée*, 2004 CarswellQue 566 (C.A. Que.) (clothing).

<sup>73</sup> *Copyright Act*, *supra* note 29, s. 64.2(2).

*v. Triolet Systems Inc.*<sup>74</sup> recognized as a general principle of copyright law that “the expression of the idea is not copyrightable if the expression does no more than embody the elements of the idea that are functional in the utilitarian sense.”<sup>75</sup> To protect function would be to protect an idea—a way of doing something—and nobody has copyright “in an arrangement or system or scheme or method for doing a particular thing.”<sup>76</sup>

We have traditionally not used copyright to protect function-driven works because there are complementary regimes for that purpose such as patent law and industrial-design legislation.<sup>77</sup> Trademark law has likewise adopted a *doctrine of functionality* and “does not protect the utilitarian features of a distinguishing guise.”<sup>78</sup> And integrated-circuit topography, driven almost exclusively by functionality and efficiency, has its own protection regime.<sup>79</sup> Blaine Thacker, former Parliamentary Secretary to Minister of Consumer And Corporate Affairs, noted in 1990: “A clear international consensus has emerged to treat intellectual property in chip topography as a *sui generis*] regime outside the framework of . . . copyright.”<sup>80</sup> When law prohibits the copying of utilitarian function, this grants the original creator a monopoly over ideas or activity, rather than expression. This is a very different privilege and is not appropriately managed through copyright or trademark law.

---

<sup>74</sup> *Delrina Corp. v. Triolet Systems Inc.*, 1993 CarswellOnt 174 (Ont. Gen. Div.), affirmed 2002 CarswellOnt 633 (Ont. C.A.), additional reasons 2002 CarswellOnt 3220 (Ont. C.A.), leave to appeal refused 2002 CarswellOnt 4080, 2002 CarswellOnt 4081 (S.C.C.) [*Delrina Trial Decision*].

<sup>75</sup> *Ibid* at para 143.

<sup>76</sup> *Moreau v. St. Vincent*, 1950 CarswellNat 4 (Can. Ex. Ct.) at para. 10. See also *Cuisenaire v. South West Imports Ltd.*, 1968 CarswellNat 56 (S.C.C.) (rods were held to be merely “devices which afford a practical means of employing the method” and thus not a proper subject of copyright. See *ibid* at para 8). But see *Bulman Group Ltd. v. “One Write” Accounting Systems Ltd.*, 1982 CarswellNat 4, 1982 CarswellNat 412 (Fed. T.D.) (while the accounting forms at issue were “functional” in a sense, they contained original literary expression in headings and instructions sufficient to be the subject of copyright).

<sup>77</sup> See Wendy J Gordon, “Fair Use in *Oracle*: Proximate Cause at the Copyright/Patent Divide” (2020) 100 BUL Rev 389 (describing how copyright law has *deferred* to patent law with respect to protection of function). See also “Bill C-60, An Act to amend the Copyright Act and to amend other Acts in consequence thereof”, 2nd reading, *House of Commons Debates*, 33-2, (26 June 1987) at 7692 (Lynn McDonald): “It would be quite inappropriate and certainly not good for the economy to extend [copyright protection] for examples of work that should only receive the lesser protection of industrial design.” Bill C-60 introduced ss 64ff of the *Copyright Act*, *supra* note 29 respecting industrial designs and useful articles.

<sup>78</sup> *Kirkbi AG v. Ritvik Holdings Inc. / Gestions Ritvik Inc.*, 2005 SCC 65, 2005 CarswellNat 3631, 2005 CarswellNat 3632 (S.C.C.) at para. 43.

<sup>79</sup> See *Integrated Circuit Topography Act*, S.C. 1990, c. 37.

<sup>80</sup> *House of Commons Debates*, 34-2, Vol 7 (9 March 1990) at 9041 (Blaine A Thacker, Parliamentary Secretary to Minister of Consumer And Corporate Affairs).

### (c) Parameters as a Computer Program

The parameters may also be seen as a computer program. Copyright for computer programs is dependent on their inclusion as literary works. Prior to 1988, *literary work* was interpreted by courts to include written computer programs and their translation into machine-readable representations.<sup>81</sup> Parliament explicitly added computer programs as copyrightable subject matter in 1998.<sup>82</sup> *Literary work* now “includes . . . computer programs”<sup>83</sup>, and *computer program* means “a set of instructions or statements, expressed, fixed, embodied or stored in any manner, that is to be used directly or indirect in a computer in order to bring about a specific result.”<sup>84</sup>

Canadian courts have not been asked to stretch this definition much beyond written computer code in its ordinary sense and the translation of that written code into various representations fixed in various media. Along this spectrum, courts have recognized high-level source code, mid-level assembly code, low-level object code, and the embedding of that object code into computer memory all as computer programs. The copyright protection of a computer program can also extend to the output of a program when it “reflects exactly and is a visual reproduction of the instructions that the creator of the program embodied on the tape or disk” (e.g. user interfaces).<sup>85</sup>

Parameter values are not generally understood as code in the ordinary sense; they are more like inputs to a system. I have not identified a case where the court was asked to decide whether a particular form of *input* to a computer system would fall within the *Copyright Act*'s definition of *computer program*. All the reported cases I have identified begin with written computer code in the ordinary sense. They take as a starting point that this computer code is, in principle, copyrightable subject matter and then ask whether the code is in fact sufficiently original or whether a derivative product of the code in question retains copyright.

But there is a sense in which the parameters are like code. A neural network architecture executes an unspecified series of transformations on the input data. It is a configurable computational graph. For example, a ten-layer neural network performs ten transformations in succession to the data. The *parameters* determine what each of those transformations does. In this sense, the parameters

---

<sup>81</sup> See *Apple Computer Inc. v. Mackintosh Computers Ltd.*, 1986 CarswellNat 606, 1986 CarswellNat 705 (Fed. T.D.), additional reasons 1987 CarswellNat 675 (Fed. T.D.), varied 1987 CarswellNat 720, 1987 CarswellNat 887 (Fed. C.A.), affirmed 1990 CarswellNat 736, 1990 CarswellNat 1027 (S.C.C.).

<sup>82</sup> See *An Act to amend the Copyright Act and to amend other Acts in consequence thereof*, S.C. 1988, c. 15, ss. 1(2), 1(3) (adding today's definition of *literary work* and *computer program* to the *Copyright Act*, *supra* note 29).

<sup>83</sup> *Copyright Act*, *supra* note 29, s. 2.

<sup>84</sup> *Ibid.*, s. 2.

<sup>85</sup> *Delrina* Trial Decision, *supra* note 74 at para 99.

are like machine code for a small and limited virtual computer.<sup>86</sup> This would be the argument that would bring the neural network parameters within the *Copyright Act's* definition of *computer program*, despite the parameters themselves being output from a different computer program.

If the neural network parameters constitute a computer program, the question remains whether the parameters are in fact copyrightable. This will turn on an originality analysis much as presented above regarding compilations, which I will not reproduce again. However, courts have given some more specific guidance regarding to what extent a computer program is protected by copyright.

Particularly, when analyzing whether a computer program is sufficiently original expression, a tension arises that is most stark in this domain: the idea/expression dichotomy. In *Delrina*, the Ontario Court of Appeal recognized that the “non-protection of ideas embraces the view that there is no copyright in any arrangement, system, scheme or method for doing a particular thing or process.”<sup>87</sup> This invites a dissection of a work into copyrightable portions or aspects and a “weeding-out” of those that have “merged” with the idea and thus are not copyrightable.<sup>88</sup> While the Supreme Court has rejected that such an approach is generally applicable to many types of works, it has left open the possibility that is nonetheless useful for computer programs.<sup>89</sup> If accepted, this approach would impose a further hurdle to copyrightability of the trained network parameters.

#### (d) Applying the Law

The automatically-trained parameters of a neural network are likely not protected by copyright. Interpreted either as a computer program or a compilation (of facts), the trained parameters are not an expression of any author’s skill and judgment. There is no author “discern[ing]”<sup>90</sup> what values the parameters should take on. The parameters are closer to facts: that when trained

<sup>86</sup> See generally Andrej Karpathy, “Software 2.0” (11 November 2017), online (blog): *Andrej Karpathy* <karpathy.medium.com/software-2-0-a64152b37c35>. To complete the analogy of parameters as code, we would also have to conceptualize the network’s architecture as the computer, albeit a virtual one, itself software rather than hardware. The network architecture would be akin to the Java Virtual Machine (see Tim Lindholm et al, “The Java Virtual Machine Specification” (28 February 2013), online: *Oracle* <docs.oracle.com/javase/specs/jvms/se7/html/index.html>).

<sup>87</sup> *Delrina ONCA*, *supra* note 63 at para 35.

<sup>88</sup> *Ibid* at paras 4350 (the ONCA noting that the trial judge appropriately undertook such a “weeding-out” without necessarily applying the “abstraction-filtration-comparison” approach developed in the US).

<sup>89</sup> See *Robinson c. Films Cinar inc.*, 2013 SCC 73, 2013 CarswellQue 12345, 2013 CarswellQue 12346 (S.C.C.) at paras. 34-35 [*Cinar*] (calling the “weeding-out” or “abstraction-filtration-comparison” approach a “reductive analysis”; all said in the context of substantiality of copying rather than originality).

<sup>90</sup> *CCH SCC*, *supra* note 19 at paras 16, 30.



using the selected dataset, the parameters converged to certain values. Even as a compilation or computer program, they are entirely “dictated by utility”<sup>91</sup> and encode a “scheme or method.”<sup>92</sup> And the training process adopted will generally be an “accepted and common industry practice.”<sup>93</sup>

Yes, a person selects the training data and selects or creates a fitness function, which an algorithm uses to tune the parameters. These selections may require some skill and judgment and may properly be the subject of copyright. But this does not result in an expression of that skill and judgment in the resulting trained parameters. This resulting configuration happens mechanically, driven by function and randomness. The person would neither care nor notice if the neural network settled on one or another among a large number of adequately performing configurations.

#### (e) Normative Justifications

To give copyright in the algorithmically-trained parameters to the algorithm designer or training manager would not advance either of the romantic or pragmatic visions of copyright law. It would not advance a romantic vision given the disconnect between those purported authors and the resulting parameters. It would not advance a pragmatic vision because the work that it would be incentivizing (requiring) people to do in order to get access to similarly-performing neural networks would not demand nor develop skill or judgment of an authorial or expressive nature. For a person to learn to work with this medium, it would be sufficient for them to execute the training phase of the neural network for enough iterations to observe its convergence behaviour. But to obtain the high-quality parameter settings found by the training algorithm when it is fed the largest training sets and run for sometimes weeks of training requires nothing more from the user but time and money. If somebody has already made this investment and released the parameters to the world, then there is nothing gained by having another person reproduce that work. And for somebody wanting to protect a market advantage, it would be easy to keep their trained parameters protected as trade secret by keeping them server-side.

The conclusion that algorithmically-trained parameters do not attract copyright protection is also consistent with a causal model of copyright.<sup>94</sup> The causal approach asks which person (if any) in a creative process has translated their “idea into an expression,” whether directly or via a proxy acting at their direction.

Last is a tangential policy issue: the work of training neural networks consumes a lot of energy. It has been estimated that to train some modern neural networks results in several times the carbon footprint of a car’s average

---

<sup>91</sup> *Distrimedica*, *supra* note 66 at para 324.

<sup>92</sup> *Delrina ONCA*, *supra* note 63 at para 35.

<sup>93</sup> *Toronto Real Estate*, *supra* note 60 at para 189.

<sup>94</sup> See Shyamkrishna Balganesh, “Causing Copyright” (2017) 117:1 Colum L Rev 1.

lifetime.<sup>95</sup> If we were to grant copyright over parameters, this would limit the ability of other parties to use these parameters, causing them to have to replicate this activity at a high cost to society.

### 3. COPYRIGHT IN ARTISTIC OUTPUT

The next issue is whether the resulting artifacts produced by the trained neural network are eligible for copyright protection and if so, in whom the copyright would be vested.

Unlike the programmer or trainer, who has a quite tenuous expressive connection to the values stored in the automatically trained parameters, the end-user of an AI-infused program might have a meaningful connection with the eventual output. The end-user may actually care (in terms of content) and control, to a degree, what the output looks like. They may very well take steps that cause that output to be a sufficient expression of their skill and judgment.

Here are four scenarios that display the of the range of creative control a human may have over the output.<sup>96</sup> I will refer to these scenarios throughout this section.

1. A person may merely click *create*.<sup>97</sup>
2. A person may choose the neural network architecture, fitness function, and training examples prior to training the neural network. (I call this person the *training manager* or *trainer*. When they also program the neural network framework, I call them the *programmer-trainer*.)
3. A person may exercise control over some human-adjustable input parameters that affect the behaviour of the neural network. For example, choosing seed images,<sup>98</sup> or adjusting parameters of the generation pipeline to obtain a desirable output.<sup>99</sup>
4. A person may generate a variety of outputs and then select from among those examples those that they find aesthetically pleasing.

Whether any of these people would obtain copyright over the artifacts produced using the AI depends on *originality*—whether the output is an expression of the person’s skill and judgment.<sup>100</sup>

<sup>95</sup> See Emma Strubell, Ananya Ganesh & Andrew McCallum, “Energy and Policy Considerations for Deep Learning in NLP” (Paper delivered at the Annual Meeting of the Association for Computational Linguistics, Florence, Italy, 30 July 2019) [unpublished, archived at arXiv, Cornell University], online: *arXiv* <arxiv.org/pdf/1906.02243.pdf> .

<sup>96</sup> Cf Boyden, *supra* note 5 at 392 (presenting a similar spectrum: composition tools, programmed output, user—programmer combinations, and unpredictable output).

<sup>97</sup> See e.g. Wang, *supra* note 8 (generates face images of people who have never existed using a generative adversarial network).

<sup>98</sup> See e.g. Matthias Bethge et al, “Deep Art”, online: *Deep Art* <deepart.io> (upload your own photo and select or upload a style and the AI will re-style your image).

<sup>99</sup> See e.g. Joshi, Stewart & Shapiro, *supra* note 4 (in this pipeline, they exposed to the user two parameters that affected the quality and degree of style that is transferred).

One approach would be to award copyright in the resulting output to the person who wrote the software.<sup>101</sup> Annemarie Bridy recognizes that while the author of the software would not be the *author-in-fact* of the resulting output, it may be suitable to deem them to be an *author-in-law*.<sup>102</sup> However, this suggestion did not foresee today's division of AI software into code and parameters.<sup>103</sup> While the traditional, literary code dictates some aspects of the production software (loading and displaying images, loading the neural-network parameters, presenting an interface to the user, pre- and post-processing, etc.) much of the "creative" production would be controlled by the network parameters discussed in the previous section. The output is explained more by those automatically-learned parameters than by the traditional software code that was written by a human. The suggestion that the programmer would be awarded copyright also fails to recognize that the end-user and training manager may play significant roles in the creative process.

I have not identified a reported Canadian decision that includes a claim of this nature (that the author of software would hold copyright in the substantive output), but there are examples from the US.<sup>104</sup> In *Rearden LLC v. Walt Disney Company*,<sup>105</sup> a company who authored a motion-capture software system claimed to hold copyright in the output files. The district judge relied on Ninth Circuit precedent that suggested "the copyright protection afforded a computer program may extend to the program's output if the program 'does the lion's share of the work' in creating the output and the user's role is so 'marginal' that the output reflects the program's contents."<sup>106</sup> Rearden failed to allege that their

---

<sup>100</sup> The UK seems to have answered this question through statute: "[i]n the case of a . . . artistic work which is computer-generated, the author shall be taken to be the person by whom the arrangements necessary for the creation of the work are undertaken." *Copyright, Designs and Patents Act 1988* (UK), s. 9(4). Although, this leaves open the question of whose activity constitutes the "arrangements necessary for the creation of the work."

<sup>101</sup> See Annemarie Bridy, "Coding Creativity: Copyright and the Artificially Intelligent Author" (2012) 5 *Stan Tech L Rev* 1 at 21–28.

<sup>102</sup> See *ibid.*

<sup>103</sup> See Karpathy, *supra* note 86.

<sup>104</sup> See generally Samuelson, *supra* note 5 (presenting how non-human authorship has been conceived of in the US since the 1960s). There are several examples where Canadian courts have recognized that the author of software may hold copyright in the program's *display*—the user interface, for example. See *Delrina* Trial Decision, *supra* note 74; *Conexsys Systems inc. c. Aime Star Marketing inc.*, 2003 CarswellQue 2003 (C.S. Que.), affirmed 2005 CarswellQue 969 (C.A. Que.) ("The screen displays of the two programs are not simply the products of the two programs. They are the most essential part of them." *Ibid* at para 267). But see *Harmony Consulting Ltd. v. G.A. Foss Transport Ltd.*, 2012 CAF 226, 2012 FCA 226, 2012 CarswellNat 3341, 2012 CarswellNat 5994 (F.C.A.) at paras. 87-89 (expressing skepticism about the *Delrina* position regarding screen displays).

<sup>105</sup> See *Rearden LLC v. Walt Disney Company*, 293 F.Supp.3d 963 (N.D. Cal., 2018) [*Rearden*].

program had done the “lion’s share of the work,’ and in particular ‘the lion’s share of the creativity,’”<sup>107</sup> so the claim was dismissed on summary judgment.

A similar approach was adopted in Australia in *Acohs Pty Ltd. v. Ucorp Pty Ltd.*<sup>108</sup> The purported copyright holders had written computer routines that were ultimately used to produce HTML (website layout code) when combined with other input data. The question was whether the programmers of the routines could hold copyright in the output HTML. The court held that the output HTML was a “separate work created by the operation of the . . . program on other elements entered into the . . . system. It did not emanate from authors.”<sup>109</sup>

These approaches seem correct, at least with respect to whether a program’s author would be awarded copyright in the program’s output. If elements of a program’s output are near-deterministic, intended by the program’s author, and not reflecting any additional creativity of the end-user or other input, then that output is a product of the program author’s skill and judgment.<sup>110</sup> This is also consistent with Canada’s requirement that any joint author “contribute original expression to the [work]” as opposed to merely contributing ideas or facilitating the other’s expression.<sup>111</sup> An example would be when a video game displays a character on a computer screen. As long as that character’s appearance was dictated by the program and the user played no role in customizing its appearance, then the author of the game’s software holds copyright in that character’s on-screen appearance. Even though the end-user was the proximate cause of the literal fixation of that depiction into computer memory and onto the screen, the depiction would be such a direct reflection of the program’s contents that the character may be deemed fixed by the program author as part of the program’s contents.

Others have suggested that the creative machine itself, as “autonomous entities” should hold copyright in these works.<sup>112</sup> This would be a significant

---

<sup>106</sup> *Ibid* at 969, citing *Design Data Corp. v. Unigate Enter., Inc.*, 847 F.3d 1169 (9th Cir., 2017) [*Design Data*], quoting *Torah Soft Ltd. v. Drosnin*, 136 F.Supp.2d 276 (S.D.N.Y., 2001) at 283 [*Torah*].

<sup>107</sup> *Rearden*, *supra* note 105 at 971.

<sup>108</sup> [2012] FCAFC 16 (F.C.A.F.C.) [*Acohs*].

<sup>109</sup> *Ibid* at para 57.

<sup>110</sup> *Cf* Boyden, *supra* note 5 at 379 (suggesting that whether a person could “predict the work’s content with reasonable specificity” is a potential criterion for authorship).

<sup>111</sup> *Neudorf v. Netzwerk Productions Ltd.*, 1999 CarswellBC 2774 (B.C. S.C.) at para. 163ff, additional reasons 2000 CarswellBC 1711 (B.C. S.C.) [*Neudorf*].

<sup>112</sup> Shlomit Yanisky-Ravid & Luis Antonio Velez-Hernandez, “Copyrightability of Artworks Produced by Creative Robots and Originality: The Formality-Objective Model” (2018) 19:1 Minn J L Sci & Tech 1 (arguing that “creative robots as autonomous entities are capable of holding copyrights in artworks they produce” at 2). *Cf* Andres Guadamuz, “Artificial Intelligence and Copyright”, *WIPO Magazine*, online: < [www.wipo.int/wipo\\_magazine/en/2017/05/article\\_0003.html](http://www.wipo.int/wipo_magazine/en/2017/05/article_0003.html) > (suggesting not that we grant copyright to the machine itself, but that we follow the approach of the United

departure from the human-focused motivations that justify our approach to copyright law today.<sup>113</sup>

I propose that the correct answer comes from a case-by-case analysis of the degree of skill and judgment reflected in the work by the purported copyright holder. This takes inspiration from the “lion’s share” test articulated by the Ninth Circuit.<sup>114</sup> However, that test needs refinement in order to satisfactorily address the issue presented in this section. The Ninth Circuit conceived of the question as a battle between the program author and the end-user. In the case of the AI creative pipeline, there is at least one additional actor: the person who trained the AI—the training manager. And the Ninth Circuit test is one-directional. It describes how the program author might be awarded copyright in the program’s outputs; it does not describe when the end-user would be awarded copyright in the outputs.

The program author and the training manager are separated from the AI output by a gulf of randomness and remoteness; they may even be surprised by the output.<sup>115</sup> This AI output will almost always not be an expression of their skill and judgment. In Bruce Boyden’s words, “the programmer may not be able to predict what content might emerge.”<sup>116</sup> This is akin to the scenario where a person merely contributes “suggestions and ideas” without contributing “original expression.”<sup>117</sup> When Sarah McLachlan was situated between Darryl Neudorf and the ultimate songs that were written, Neudorf could not be awarded copyright. McLachlan was “free to accept or reject” Neudorf’s ideas.<sup>118</sup>

---

Kingdom which allocates copyright to the “person by whom the arrangements necessary for the creation of the work are undertaken”).

<sup>113</sup> See e.g. US Copyright Office, *Compendium of U.S. Copyright Office Practices*, 3rd ed (2017), § 906.6, online (pdf): *U.S. Copyright Office* < [www.copyright.gov/comp3/prior-editions.html](http://www.copyright.gov/comp3/prior-editions.html) > . See also Johnathan Osha et al, “2019 Study Question: Copyright in Artificially Generated Works” (24 May 2019) at 7, online (pdf): *International Association for the Protection of Intellectual Property (AIPPI)* < [www.aippi.se/s/2019\\_SE\\_2019\\_S-study\\_Question\\_Copyright\\_in\\_artificially\\_generated\\_works\\_2019-05-24.pdf](http://www.aippi.se/s/2019_SE_2019_S-study_Question_Copyright_in_artificially_generated_works_2019-05-24.pdf) > (a majority of the countries’ groups recommended to continue to restrict copyright to works created by humans).

<sup>114</sup> See *Rearden*, *supra* note 105 at para 969, citing *Design Data Corp.*, *supra* note 106, quoting *Torah*, *supra* note 106.

<sup>115</sup> See Gervais, *supra* note 5 at 2059 (noting that the output is “at least one degree removed from the human programmer” and asks “[w]ho is the *author* of the (unpredictable) outputs created by a (deep learning) AI machine”). See also Boyden, *supra* note 5 at 379 (focusing on whether the purported author could “predict the work’s content with reasonable specificity”).

<sup>116</sup> Boyden, *supra* note 5 at 389. Recall, Boyden suggests that “[a] person would be the author of the output of a program only if he or she could predict, or rather foresee, that output, more or less.” *Ibid* at 392.

<sup>117</sup> *Neudorf*, *supra* note 111 at paras 26.

<sup>118</sup> *Ibid* at para 28.

Likewise, AI randomness separates the programmer-trainer from the resulting parameters and even further from the resulting output.

But there will be circumstances in which a single person will write a custom AI framework and carefully refine the training examples and the training algorithm according to a vision of what the resulting output should look like (Scenario 2, from above). This takes us much of the way to justifying awarding copyright in the output to the programmer-trainer. The programmer-trainer will have produced an AI-infused program that is primed to produce new artifacts. But because of the additional randomness generally involved in the generation process, the output cannot be generally described as a reflection of the program's contents as articulated by the Ninth Circuit's test.<sup>119</sup> Or, in the language from *Acohs*, the output "did not emanate from authors."<sup>120</sup> It is only when the production program is combined with the trained parameters and then initialized with random values that a particular output is produced. It would only be in the rare circumstance, where the program author and training manager (which would likely have to be the same person) can demonstrate that they have largely dictated the content of the AI output, that they would be eligible for copyright in the output.<sup>121</sup> When the programmer-trainer effectively acts as a director, to the exclusion of creative input from the end-user, they should be awarded copyright in the output artifact.

There are much clearer paths for the person *using* the AI system to be awarded copyright in the output. The artifact would need to be the product of an exercise of the user's skill and judgment. This would elevate the user's role such that the program will not have contributed the "lion's share." According to the Ninth Circuit's test, that would preclude the author of the program as a potential owner of copyright in the output.

For this end-user, there is a spectrum of potential involvement.<sup>122</sup> On one end of the spectrum, an end-user who merely clicks *create* (Scenario 1, above), having not exercised their skill or judgment, would not be awarded copyright over the output.<sup>123</sup> But I propose that two of the scenarios from above may result

<sup>119</sup> See also Gervais, *supra* note 5 at 2062 (suggests an "autonomy threshold": when the "machine. . . is no longer a tool in the user's hands or a reflection of its (human-made) program").

<sup>120</sup> *Acohs*, *supra* note 108 at para 57.

<sup>121</sup> See e.g. Ginsburg, *supra* note 54 (an author of a photograph was "the person who set up the photo, not the one who pushed the button" at 1070, n 24) citing a case I have been unable to locate: Trib gr inst Paris, 6 July 1970, (1970) RIDA 190 (affaire Paris Match). See also *Lindsay v. The Wrecked and Abandoned Vessel RMS Titanic*, 52 U.S.P.Q.2d 1609 (S.D. N.Y., 1999) (recognized that a director who exercises "such a high degree of control over a film operation" that the "final product duplicates his conceptions and visions of what the film should look like" may be the author, and that "authors may be entitled to copyright protection even if they do not 'perform with their own hands the mechanical tasks of putting the material into the form distributed to the public'").

<sup>122</sup> See Grimmelmann, *supra* note 5 at 410 (phrasing the problem as distinguishing "computer users who are genuine authors from users who merely push a button").

in the end-user holding copyright to the output: Scenario 3 (the user who is given access to production-time parameters that control the form of the output) and Scenario 4 (the user who selects from a variety of outputs).

In Scenario 3, the end-user is effectively using the AI-infused software as an elaborate tool for expression, like a person using a photoshop filter. This is how Kristen Stewart, as director of *Come Swim*, used a neural network to perform style transfer in order to create a “look in service of the story.”<sup>124</sup> Stewart and her team chose from among several alternative neural network architectures based on the look that they wanted to produce. They also experimented with various source images from which a style would be transferred. And they adjusted several meta-parameters (parameters that are not automatically trained) during production to achieve the desired artistic effect. This activity is properly a basis upon which to grant these directors copyright over the neural-network output because it is a fixed expression of the authors’ non-trivial skill and judgment.<sup>125</sup>

In Scenario 4, the end-user takes more of a hands-off approach, but is still selecting the aesthetically pleasing outputs from among those created by the AI. There is no case law from Canada relating to whether random, or accidental, authorship can attract copyright. A US case, *Alfred Bell & Co. Ltd. v. Catalda Fine Arts Inc. et al.* suggests that unintentional expression can be “adopt[ed]” by the author.<sup>126</sup> I propose that this scenario would be assessed based on whether end-user’s selection amounts to an original compilation. Even assuming that each individual production is not eligible for copyright (as in Scenario 1), a user’s selection or arrangement may create a copyrightable compilation. This would at minimum, require the end-user to select more than one of the outputs. Canadian courts have consistently discussed compilations assuming that they are about selection or arrangement of more than one item.<sup>127</sup> And as noted, it is difficult to articulate the threshold for sufficient skill and judgment in a compilation.<sup>128</sup> Courts have looked for “a complete independent work of carefully selected paintings,”<sup>129</sup> “selection and layout” of data “according to unpublished

<sup>123</sup> In this scenario, the work would likely enter the public domain, given how rare it would be for the programmer—trainer to hold the copyright.

<sup>124</sup> Joshi, Stewart & Shapiro, *supra* note 4, s. 3.

<sup>125</sup> See Bonadio & McDonagh, *supra* note 5 at 123 (“at times the user is the person who sets the parameters and provides data for the algorithm in ways that significantly influence the final work; and in some circumstances the user may even affect the way the algorithm functions”).

<sup>126</sup> *Alfred Bell & Co. Ltd. v. Catalda Fine Arts Inc. et al.*, 191 F.2d 99 (2d Cir., 1951) at 104.

<sup>127</sup> See e.g. *Rains v. Molea*, 2013 ONSC 5016, 2013 CarswellOnt 11337 (Ont. S.C.J.) [*Rains*] (“no evidence that Rains exercised skill and judgment to select a crumpled paper image and arrange it with others” at para 17). See also *CCH SCC*, *supra* note 19 at para 33 (“[i]t is not the several components that are the subject of copyright, but the over-all arrangement of them”, citing *Slumber-Magic Adjustable Bed Co. v. Sleep-King Adjustable Bed Co.*, 1984 CarswellBC 765 (B.C. S.C.) at para. 17).

<sup>128</sup> See *Toronto Real Estate*, *supra* note 60.

standards of selection for the first time in Quebec and in Canada,”<sup>130</sup> or an “intimate connection” between the selected elements.<sup>131</sup> Where the end-user does acquire copyright in the compilation, this has no bearing on whether copyright subsists in any of the constituent elements nor does it depend on or inform whether the programmer-trainer would have copyright in those individual outputs.<sup>132</sup>



This is a compilation of images generated from *This Person Does Not Exist*.<sup>133</sup> I used the program to generate about 10 × as many faces as this and selected those that allowed me to arrange them into a grid that displayed a spectrum of apparent age (left—right) and facial expression (top—bottom).

<sup>129</sup> *Rains*, *supra* note 127.

<sup>130</sup> *Édutile Inc. v. Automobile Protection Assn. (APA)*, 2000 CarswellNat 744, 2000 CarswellNat 1258 (Fed. C.A.) at paras. 14-15, leave to appeal refused 2001 CarswellNat 29, 2001 CarswellNat 30 (S.C.C.).

<sup>131</sup> *Robertson v. Thomson Corp.*, 2006 SCC 43, 2006 CarswellOnt 6182, 2006 CarswellOnt 6183 (S.C.C.) [*Robertson*].

<sup>132</sup> See *Copyright Act*, *supra* note 29 (“[t]he mere fact that a work is included in a compilation does not increase, decrease or otherwise affect the protection conferred by this Act in respect of the copyright in the work or in the moral rights in respect of the work” at s. 2.1(2)).

<sup>133</sup> Wang, *supra* note 8.



Last, I consider the possibility that the copyright might be shared in some manner between the end-user and the programmer or trainer. They may be joint authors, the copyright may be “layered” as in musical works, or they may be distinct authors of a collective work or of aspects of a compilation.

One path to shared copyright is through joint authorship. If the parties each contribute original expression to the work (subject to all the analysis above), with the “intention that their contributions be merged into. . . a unitary whole” and “regard themselves as joint authors,”<sup>134</sup> then the parties are joint authors, holding the copyright as co-owners.<sup>135</sup>

Another path to shared copyright is if the work is a *collective work*.<sup>136</sup> Instead of a one-or-the-other (or neither) approach, this alternative would recognize that the programmer-trainer contributes one aspect of the eventual expression and the end-user another aspect. But collective works are works in which there are “distinct parts by different authors.”<sup>137</sup> The mixture of original expression (if any) that is generally present in works produced using AI does not generally result in such “distinct parts.” The output is a generally a unified whole, not possible of deconstruction into constituent parts attributable to different authors.

However, this kind of fragmentation *has* been recognized in compilations. One can even consider compilations to be a special case of collective works.<sup>138</sup> For example, in *Robertson v. Thomson Corp.*, the court held that various freelance journalists held copyright to their articles, but that the newspaper held copyright to the compilation of the articles in aggregate.<sup>139</sup> And copyright in music results similarly results in a layering of copyright: that of the score and that the sound recording.

Aside from the realm of collective works, compilations, and music, courts have not fragmented or layered copyright between authors based on a separation of a work into conceptual building blocks that are identifiable in the eventual output expression. In *Robinson c. Films Cinar inc.*, the defendant argued that an infringement analysis of a children’s cartoon should follow something like the abstraction-filtration-comparison approach that has been discussed in some

---

<sup>134</sup> *Neudorf*, *supra* note 111 at paras 34, 79—80.

<sup>135</sup> See David Vaver, *Intellectual Property Law: Copyright, Patents, Trade-marks* (Toronto: Irwin Law, 2011) at 118.

<sup>136</sup> See *Copyright Act*, *supra* note 29, s. 2 (“collective work means. . . (c) any work written in distinct parts by different authors, or in which works or parts of different authors are incorporated” [emphasis added]).

<sup>137</sup> *Ibid.*

<sup>138</sup> See e.g. *Pinto v. Bronfman Jewish Education Centre*, 2013 CF 945, 2013 FC 945, 2013 CarswellNat 4833, 2013 CarswellNat 3396 (F.C.) at para. 151 (“in addition to the copyright which attaches to individual works, there may be a second, concurrent copyright to the compilation as a whole. . . this second copyright does not extinguish the copyright over individual components of the collective work” [emphasis added]).

<sup>139</sup> See *Robertson*, *supra* note 131.

software infringement cases.<sup>140</sup> The Court rejected this, viewing it a “reductive analysis” likely not suitable outside of the domain of software infringement.<sup>141</sup> While that critique was directed at an infringement analysis, it applies equally to an approach that would divide copyright allocation between parties. Compilations can easily be split into constituent contributions. Music likewise has long-recognized dividing lines between its layers of constituent copyrightable work products. A single piece of visual art or prose does not easily lend itself to such a dissection.

In summary:

**The programmer of an AI framework is likely too distant** from the eventual output to deserve an award of copyright. Their ideas, skill, and judgment have not been converted into expression in the output.

**If the programmer also acts as trainer and puts their vision into the system, they may have a claim to copyright in the outputs** if the outputs are so determined by the trained AI system that they can be said to be a direct reflection of the trained program’s contents.

**If the end-user uses the system as a tool in their own creative endeavours, this may be sufficient to award them copyright.** This will also likely displace any programmer—trainer as potential holder of copyright in the output.

Even if the end-user is limited to merely selecting from among uncopyrightable individual outputs, **they may hold copyright in their compilation of the outputs due to skill and judgment expressed in the selection and arrangement.**

**Compilations and joint-authorship are two paths that might recognize shared copyright** in the eventual expression. But it is unlikely that Canadian courts will look to dissect an eventual unitary output into constituent bases in order to fragment copyright ownership between the end-user and the programmer—trainer.

#### 4. IS BUILDING A TRAINING SET INFRINGEMENT?

Recall in the fact pattern at the outset of this article that Becca relied on Alex’s original artwork as training examples. To use Alex’s art as training examples invariably requires copying them. This is supported by the submissions of several witnesses to the Statutory Review of the Copyright Act (the Statutory Review).<sup>142</sup>

Google:

<sup>140</sup> *Cinar*, *supra* note 89.

<sup>141</sup> *Ibid* at para 35.

<sup>142</sup> INDU, *Statutory Review*, *supra* note 2.

Machine learning requires data-based training examples, and it is often necessary for the datasets to be copied, processed and re-purposed. Sometimes, this includes copyrighted material, such as an image recognition system using millions of photographs. Unless there is an exception. . . machine learning risks infringing copyright. . . . It is unclear whether this activity would fall within existing exceptions.<sup>143</sup>

Vector Institute et al.:

Teaching a computer. . . requires large data sets, tens of millions of photographs. . . and other works that can be used to train software. . . . analyzing data to teach a computer often requires making an incidental technical copy of a lawfully acquired work.<sup>144</sup>

From the BSA (The Software Alliance):

Because the machine learning process can involve the temporary creation of machine-readable copies of works, difficult legal questions can arise about the potential copyright implications of those reproductions.<sup>145</sup>

Microsoft:

Because machine learning techniques may require the incidental copying of lawfully acquired copyrighted works to make them accessible for machine learning, analyze them for patterns, facts, and insights, and use those copies for data verification, there is a risk that copyright can be asserted as a basis to block this activity unless permission of a copyright owner is granted.<sup>146</sup>

Element AI submitted that current copyright law might already allow this copying, either as a temporary reproduction or as fair dealing as research or private study.<sup>147</sup>

All of these submissions asked Parliament to clarify that copying works for the purpose of machine learning would be fair dealing. Outside of the exercise of a fair dealing user right, this copying is likely *prima facie* infringement because these training sets are often built using copyrighted images acquired from the internet.<sup>148</sup> The Statutory Review does recommend amending the *Copyright Act*

<sup>143</sup> House of Commons, Standing Committee on Industry, Science and Technology, *Evidence*, 42-1 (4 September 2019) (Brief submitted by Google Canada) at 5.

<sup>144</sup> House of Commons, Standing Committee on Industry, Science and Technology, *Evidence*, 42-1 (14 December 2018) (Brief submitted by Vector Institute et al.) at 1.

<sup>145</sup> House of Commons, Standing Committee on Industry, Science and Technology, *Evidence*, 42-1 (26 August 2019) (Brief submitted by BSA The Software Alliance) at 3.

<sup>146</sup> House of Commons, Standing Committee on Industry, Science and Technology, *Evidence*, 42-1 (4 September 2019) (Brief submitted by Microsoft Canada Inc.) at 5.

<sup>147</sup> See House of Commons, Standing Committee on Industry, Science and Technology, *Evidence*, 42-1 (3 September 2019) (Brief submitted by Element AI Inc.) at 4–5. See also *Copyright Act*, *supra* note 29, ss. 30.71, 29—29.2.

“to facilitate the use of a work. . . for the purpose of informational analysis.”<sup>149</sup>

Most of these submissions contemplated that Parliament would provide this certainty by clarifying the fair-dealing user right. Canada’s current copyright law allows one to make copies of images in order to build a training set when that activity falls within an existing fair-dealing user right.<sup>150</sup> For copying to fall within the fair-dealing user right, it must be for the purpose of “research, private study, education, parody or satire,” “criticism or review,” or “news reporting.”<sup>151</sup> This is an exhaustive list.<sup>152</sup> Copying during the course of academic research may be viewed differently than copying in a commercial setting with no research motivation.<sup>153</sup>

The private-purposes user right found at section 29.22(1)(e) may also allow one to create a training set from copyright-protected images. The private-purposes user right requires only that “the reproduction is used only for the individual’s private purposes.”<sup>154</sup> The copying need not be for research or any of the fair-dealing purposes, so this right has the potential to provide a broader scope than the general fair-dealing rights when the copying is done for “the individual’s private purposes.” No court in Canada has yet interpreted this user right, so it is unclear what scope would be given to “private purposes.”

Parliament should remove this uncertainty in line with the Statutory Review’s recommendation.<sup>155</sup> This would be consistent with a pragmatic purpose of copyright that seeks to encourage and enable more people to create more works. It is especially justified to deem training-set creation as fair dealing when the algorithm being trained is seeking only to learn from non-expressive content found in the training material.<sup>156</sup>

<sup>148</sup> See e.g. “About ImageNet”, online: *ImageNet* <image-net.org/about> (“ImageNet does not own the copyright of the images”); “Download”, online (archived): *ImageNet* <web.archive.org/web/20201112004102/image-net.org/download-faq> (“images in their original resolutions may be subject to copyright, so we do not make them publicly available on our server”).

<sup>149</sup> INDU, *Statutory Review*, *supra* note 2 (recommendation 23 at 87).

<sup>150</sup> Because fair dealing is a *user right* rather than an exception, it is owed a “large and liberal interpretation.” *CCH SCC*, *supra* note 19 at para 51. See also Michael Geist, “Fairness Found: How Canada Quietly Shifted from Fair Dealing to Fair Use” in Michael Geist, ed., *The Copyright Pentology: How the Supreme Court of Canada Shook the Foundations of Canadian Copyright Law* (Ottawa: University of Ottawa Press, 2013) 157.

<sup>151</sup> *Copyright Act*, *supra* note 29, ss. 29—29.2.

<sup>152</sup> See *CCH SCC*, *supra* note 19 at para 9. See also *Alberta (Minister of Education) v. Canadian Copyright Licensing Agency*, 2012 SCC 37, 2012 CarswellNat 2419, 2012 CarswellNat 2420 (S.C.C.).

<sup>153</sup> “[R]esearch is not limited to non-commercial or private contexts.” *CCH SCC*, *supra* note 19 at para 51. But the commercial aspect can weigh against its fairness. See *Public Performance of Musical Works, Re*, 2012 SCC 36, 2012 CarswellNat 2380, 2012 CarswellNat 2381 (S.C.C.).

<sup>154</sup> *Copyright Act*, *supra* note 29, s. 29.22(1)(e).

<sup>155</sup> See INDU, *Statutory Review*, *supra* note 2 (recommendation 23 at 87).

But even when the algorithm will adopt stylistic elements from the training material, deeming this use to be fair dealing does not leave the original works without protection. The ultimate creative program must still not infringe.

## 5. CAN THE OUTPUT INFRINGE?

The final question I address is whether the AI output can infringe. The AI training process is intended to infuse the neural network with the stylistic influences of the training examples. While it may be rare or unlikely, nothing generally precludes an AI from producing output that substantially reproduces one or more of the original training examples. The creator of generative adversarial networks (such as that used to generate the faces shown above) cannot explain why the generator does not learn to more closely mimic the training examples.<sup>157</sup>

So, for the purpose of this analysis, I am assuming the first element of copyright infringement: a *prima facie* reproduction of the work.<sup>158</sup>

The remaining elements are summarized well in *Rains v. Molea*.<sup>159</sup> The “plaintiff. . . must establish that the defendant had [a reasonable possibility of access] to the plaintiff’s work.”<sup>160</sup> Here, the court is looking for evidence of a “causal connection between the original and the allegedly copied work.”<sup>161</sup> If the plaintiff can establish similarity and the reasonable possibility of access, “the burden then shifts to the defendant to prove the second work was created independently of the first.”<sup>162</sup>

Taken literally, this test likely places too high a burden on the defendant who did not participate in the training of the AI (Cory, from the original fact pattern). If the plaintiff can establish the reasonable likelihood that the defendant had access to the original work<sup>163</sup> and the defendant is using an AI that happens to have been trained using that original work, the defendant will not be able to meet their burden of proving that the second work (the AI’s output) was created

---

<sup>156</sup> See Mark A Lemley & Bryan Casey, “Fair Learning” (2021) 99:4 Tex L Rev 743 at 750.

<sup>157</sup> See “Ian Goodfellow: Generative Adversarial Networks (GANs)” (18 April 2019), online (podcast): *Artificial Intelligence with Lex Fridman* <lexfridman.com/ian-goodfellow/> .

<sup>158</sup> *Copyright Act*, *supra* note 29, ss. 2 (“infringing”), 3(1), 27(1).

<sup>159</sup> *Rains*, *supra* note 127.

<sup>160</sup> *Ibid* at paras 47—48.

<sup>161</sup> *Atomic Energy of Canada Ltd. v. Areva NP Canada Ltd.*, 2009 CF 980, 2009 FC 980, 2009 CarswellNat 5755, 2009 CarswellNat 2991 (F.C.) at para. 35, varied on reconsideration 2009 CarswellNat 5766, 2009 CarswellNat 3522 (F.C.).

<sup>162</sup> *Rains*, *supra* note 127 at para 68.

<sup>163</sup> Many AIs are trained using publicly available (even if not free-to-use) works. A defendant may very well have access to images or data that were used to train the AI (as those images may be widely available in public) without *knowing* that those images were used to train the particular AI in use.

independently of the first. The second work will *not* have been created independently of the first because the first was part of the AI's training data. However, it does seem wrong to make this defendant, who had no part in the training, liable for infringement.

I propose that the burden on the defendant should be lower. Infringement is premised in causation. The defendant should be required only to prove that the second work was created *independently from the access that they (the defendant) had to the source work*. This would sever the causal chain through the defendant. A defendant could demonstrate this by establishing that they played no role in training the AI. This would be a refinement of the *independent creation* doctrine to account for the fact that an AI could create an infringing work independent of the access that the *user* of that AI had to the source work.

The situation is different for a defendant who was aware of the material used to train the AI. If a person selects the training examples, trains an AI, and then uses that AI to produce output that substantially reproduces one of the training examples, they will not be able to prove that the second work was created independently of the access that they had to the source work. This is likely a straightforward copyright infringement.

A more difficult case is the trainer who releases an AI to somebody else and it happens to produce output that substantially reproduces the training examples. This is Becca from the original fact pattern. As argued, the person using the AI would likely not be infringing, but would Becca be liable?

The path to liability for Becca is not through secondary infringement because “[a]bsent primary infringement, there can be no secondary infringement.”<sup>164</sup> If Becca is infringing, it is through primary infringement.

A trained AI does not generally and substantially reproduce the training data. If it does, though, then the trainer will have (perhaps accidentally) fixed a representation of the original training data in the trained AI. The program contents would contain “a direct reflection”<sup>165</sup> of the training image(s). And there is no *mens rea* requirement for non-criminal infringement.<sup>166</sup> If Becca either “communicates to the public by telecommunication” or “publishes” an AI system that she trained and it happens to routinely reproduce the original training data, she has likely infringed.<sup>167</sup>

<sup>164</sup> *CCH SCC*, *supra* note 19 at para 82. See also *Copyright Act*, *supra* note 29, s. 27(2).

<sup>165</sup> I take this phrase from the Ninth Circuit's test (discussed above) for allocating copyright in a program's outputs.

<sup>166</sup> See also 91439 *Canada Ltée v. Éditions JCL Inc.*, 1992 CarswellNat 1346, 1992 CarswellNat 1347 (Fed. T.D.), varied 1994 CarswellNat 1432, 1994 CarswellNat 1432F (Fed. C.A.) (“[n]either intention to infringe, nor knowledge that he is infringing on the part of the defendant, is a necessary ingredient. . . it is no defence that the defendant was unaware (and could not have been aware) that what he was doing infringed the copyright in the plaintiff's work” at para 17), quoting from John S McKeown & Harold G Fox, *The Canadian Law of Copyright and Industrial Designs*, 2nd ed (Toronto: Carswell, 2003) at 331—32. Criminal infringement, though, requires the defendant to *know* that they are infringing. See *Copyright Act*, *supra* note 29, s. 42.

The conclusion that Becca might infringe when she trains and then distributes a neural network may seem inconsistent with the conclusion above that the programmer-trainer would generally not be able to hold copyright in either the trained parameters or the eventual output of the program. The argument was that they are separated from both by a gulf of randomness. However, when the AI has effectively embedded a representation of the original training data in its parameters such that it is reproducing them as part of the output, this is a different scenario. This would be considered a well-known training error called *overfitting*, something a careful trainer would seek to avoid. Certain errors in selecting a network architecture or in arranging the training data can remove the element of randomization and cause a neural network to overfit and reproduce the training data. I conclude that the *Copyright Act* places a burden on the trainer to ensure they have not done this before distributing their system to others.

## CONCLUSION

AI art presents many questions that touch the heart of copyright. To answer these questions, we need to focus on what we are incentivizing. Canadian copyright incentivizes human practice converting ideas into expression—for one’s own benefit and for the benefit of others—through each particular work and with an optimism of works yet to come.

Today’s state-of-the-art “creative” AI tools are based on a technology (neural networks) that serve to separate the programmer and trainer from any of the eventual expression, even the expression stored in the automatically-learned network parameters. It would be very rare that a programmer or trainer might obtain copyright in the output from an automatically trained “creative” AI.

However, there are a multitude of ways to use such an AI to produce output, many of which would very well justify awarding copyright to the end-user, especially when they use the AI as an elaborate brush with which to bring their own ideas to life in expression.

The current methods of training these creative AI tools requires large amounts of training data: existing works often protected by copyright. It is unclear whether Canada’s fair dealing user right allows for such copying for the purpose of training a neural network, particularly when not for private purposes. When a fair dealing user right is not available, this copying would be copyright infringement: unauthorized reproduction of existing works. Canada should clarify or expand the fair dealing user right to allow for such copying.

Trainers must be careful that they have not simply embedded a representation of the training examples in the AI. If the AI effectively contains “direct reflections” of the training data such that it regularly reproduces them,

---

<sup>167</sup> See *Copyright Act*, *supra* note 29, s. 2.2(1)(i) (“publication means. . . making copies of a work available to the public”), s. 3(1) (copyright providing the “sole right to produce or reproduce the work. . .”).

distributing such an AI would be copyright infringement. The trainer has a burden to verify that they are not distributing copies of the training data.

This analysis allocates copyright in a manner consistent with a pragmatic conception of creativity and art. It keeps the focus on human expression and allows for free distribution of the material needed for more people to have more practice with creative tools while preserving protection for original expression.