

This is a section of [doi:10.7551/mitpress/10913.001.0001](https://doi.org/10.7551/mitpress/10913.001.0001)

# **Beyond the Creative Species**

## **Making Machines That Make Art and Music**

**By: Oliver Bown**

### **Citation:**

*Beyond the Creative Species: Making Machines That Make Art and Music*

**By: Oliver Bown**


**DOI: 10.7551/mitpress/10913.001.0001**

**ISBN (electronic): 9780262361750**

**Publisher: The MIT Press**

**Published: 2021**

The open access edition of this book was made possible by generous funding from University of New South Wales.



**The MIT Press**

## Preface

This book is about the practice, the technology, and the deeper implications of the automation of creative tasks in artistic domains. It was written primarily to provide a broad introduction to the field—which is in fact really a cluster of different fields that come together to inform this fascinating topic. It is aimed first and foremost at professionals working in creative industries who want to understand how their field might be influenced by technological advances in the automation of cognitive tasks, some part of which involves cutting through the noise we find today around the impact of AI, whichever the area. As such it is also aimed more generally at anyone interested in learning more about this area of exciting and rapid development. This bringing together of different fields, from AI, design, social theory, the psychology of creativity, creative practice research, and elsewhere, also gives the book a secondary function, to collate the breadth of relevant ideas and their interconnections in a way that may be useful for experts in any specific one of these subdisciplines. It may be especially useful for engineers and data scientists unfamiliar with work in social, psychological, and creative domains, perhaps offering some common language and reference points that can be used to work across the richly multidisciplinary communities that must come together to work in this field.

As such, this book is unashamedly jack-of-all-trades-master-of-none. Most importantly, it is only partly about the technology and what it can do. To narrow the topic to just this question would be hopelessly lacking in context; every engineer respects the importance of understanding a problem clearly before you can solve it. The question of how to make convincing or effective creative algorithmic systems is a core theme. But to answer this question merely from the perspective of algorithms, or even combining

algorithms and cognitive science, is insufficient. I am certainly not the first person to stress in this context that human creative artistic activities are not only social in important ways but *primarily* social. This is probably the majority view. But exactly how creative practice is best viewed through a social lens remains more than a little bit challenging. To understand in any depth what people are doing (and why) when they create, and hence understand what the technical challenges are, is at least as vexing as those technical challenges themselves. This book aims to provide informed discussion across this spectrum of topics and may seem to do so sometimes at the expense of depth or coherence. There is no subject covered in this book that couldn't have been given more time and detail. The book's reviewers were thorough in pointing out omissions of work that they themselves considered important in their respective fields of specialist interest. I addressed some of these omissions, but I did not labor each subject exhaustively where I felt that the underlying idea had been covered.

Having thought about this subject for many years, situated in different university departments and talking with people from very different walks of life, I feel certain that covering such broad ground can be done equally well from a wide number of starting points and, for each starting point, looking in many different directions. There is no obvious order of events, and my choice of structure is undoubtedly grounded in various biases and niches of interest I have accumulated from researching computational creativity in those contexts I have encountered it, from anthropology to the science of artificial life, to music psychology, design, and creative practice research. For practical reasons, this book has a broad arc, starting and ending in more accessible and broad discussion, while digging deeper into academic depth toward the middle of the book. Its progression through the psychology of creativity, social creativity, the arts, algorithms, design, evaluation, and social impacts follows the narrative that I feel is most relevant but is certainly not the only possibility.

I have also tried not to become trapped by the shackles of one academic filter bubble over another, so wherever possible—and it has not always been easy—I've tried to talk in plain language and maintain some skepticism about the limits of respective fields. I'd rather write something that jumps position and risks consistency than to write something too deeply entrenched in a certain framework. At the same time, I'm liberal in drawing in diverse theoretical ideas that may not be always be very well supported

by evidence, and I hope that this liberty is understood as a way of getting ideas on the table, not as a naive, unfiltered acceptance of any particular body of theory. Likewise, this book may appear to omit my own view and have no overt thesis. This is largely by design; I try to speak through others as much as possible, and I don't believe a subject as sprawling as this can be easily tamed by any one overarching thesis. But the choice and ordering of material is mine and outlines a thesis of sorts.

These considerations can be summed up in four principles that I followed for writing this book: what you leave out can be as important as what you include; order knowledge to help others; speak through other people; work in manageable units.

It is worth commenting on the timing of writing of this book. When I pitched the concept in 2016 to MIT Press there were no full single-authored books in this field and even as recently as then, there was some question as to whether people would be much interested in, or take seriously, the idea of AI-automated creative production as anything more than speculative. The small International Conference on Computational Creativity was, in my view, the only real place to be showing work in this area, and progress wasn't occurring at lightning pace. To my simultaneous delight and horror, as I set out to write the book in 2017, and in the subsequent two years getting distracted and delayed by grant writing, teaching, children, and creative projects, the subject exploded in a big way in the wake of the deep learning revolution, with this book still sitting in draft status. Google were making music generation tools, mainstream artists were using machine learning to compose albums, ML-generated paintings were being sold for large sums at auction, and the number of academic conferences and journals and other public forums hosting special sessions on AI and creativity had multiplied many times. My social media feeds brimmed with a stream of articles asking with tedious repetition whether "an AI" might be the next Bach or Picasso, and every day a new world-first in AI art or music was announced. I was happy for (if a little overwhelmed by) the constant feed of new material to reference, but also a little anxious that this ship was sailing far more rapidly than I could keep up with.

At the time of writing, the burst of interest seems to have plateaued, causing me to reflect once again on the hype-cycle that a field such as this regularly passes through, as I discuss in chapter 1. There is no doubt that with such a fast-moving subject the material in this book will need refreshing

frequently, but it is also hopefully possible to recognize the continuity of many themes, tracing back to earlier work in the field. On the one hand, we are already inhabitants of the fantasy world of someone else's future, where machines routinely perform magic; our worldview has already changed significantly with respect to our relationship to AI and the computational built environment as we adapt to it. On the other hand, machines remain brittle and glaringly incapable except in the most structured of contexts. The primary takeaway for me, becoming another principle for this book, is to avoid committing to predictions! Another useful piece of advice, from a colleague with several books under her belt, is that everyone who writes a book on a subject will write a different book that appeals to a different audience, so not to worry about getting in there first, and to focus on clearly and patiently communicating the themes I've developed in my work over several years.

I wish to acknowledge my potential biases. Some of the academic fields related to this topic are historically male dominated, and consequently the literature and history of the field are too. This book reflects that history, but I acknowledge my potential complicity in reinforcing a biased view of the field. The contemporary field has a greater gender diversity but still has far to go. In writing this book, I have been inspired by the work of many women, including Anna Jordanous, Anna Kantosalu, Kate Compton, Gillan Smith, Amy Hoover, and Alice Eldridge, and by the significant work of creativity researchers such as Margaret Boden and Teresa Amabile and social researchers Georgina Born, Genevieve Bell, Kate Crawford, and Lucy Suchman. The book is also grounded in a predominantly European and North American community of research and practice, and may not evenly represent the world's contributions to this field.

Writing this book was made possible first and foremost by the work of the computational creativity research community, whose work you will find discussed throughout. This includes several people whom I've had the pleasure of collaborating with, and who have taught me much in the process: Geraint Wiggins, Alice Eldridge, Jon McCormack, Tim Blackwell, Arne Eigenfeldt, Philippe Pasquier, Andrew Brown, Kaz Grace, Alan Dorin, Rob Saunders, and Petra Gemeinboeck. An extra special thanks goes to Dan Ventura for more detailed feedback and ideas, and to my partner Alana George not only for putting up with the emotional roller coaster of such a big project but for some very practical formatting and document management

support in its final stages. My faculty management generously supported the sabbatical break to work intensively on the first draft on the manuscript and strategically dangled carrots in wait of the final publication. During that sabbatical break I had the opportunity to visit some of the most interesting people in the field in their home environments: Douglas Eck, Jesse Engel, Bill Hsu, Andy Elmsley, Pierre Barreau, Sebastian Risi, Michael Casey, Bob Sturm, Paul Brown, Georgina Born, Alan Blackwell, Julian Togelius, and Amy Hoover. In no particular order, I'd also like to extend my thanks to the following people for, variously, their engaging conversation, ideas, and collaborations: Sam Britton, Dan Wilson, Estelle Hoen, Margaret Boden, Benjamin Carey, Aengus Martin, Craig Vear, Lian Loke, Christian Guckelsberger, Christopher Redgate, François Houle, Petros Vrellis, Palle Dahlstedt, Lindsay Kelley, Roger Mills, Linda Candy, Caleb Kelly, Deb Turnbull, Dan McKinlay, Matt Yee-King, Prue Gibson, Anna Munster, Keir Winesmith, Toby Walsh, Caroline Pegram, Justin Shave, Charlton Hill, the Sydney Google Creative Lab team, Liam Bray, Steffan Ianigro, Sam Gillespie, Angelo Fraietta, and my family. I also thank Chloe McFadden for some help with the final proofs.

Needless to say, without MIT Press the book wouldn't have received such extensive reach, and I thank Doug Sery and Noah Springer for their ongoing support and patient responses, as well as the book's copyeditor, Lunaea Weatherstone, and Melody Negron at Westchester Publishing Services.

This book is dedicated to my fantastic dad, Bruce—who always encouraged me to pursue my interests and who declared that this is the first academic thing I've written that he can make sense of—and to the memory of Harold Cohen, an innovator and a most thoughtful individual.



© 2021 Massachusetts Institute of Technology

This work is subject to a Creative Commons CC-BY-NC-ND license.

Subject to such license, all rights are reserved.



The open access edition of this book was made possible by generous funding from University of New South Wales.

This book was set in Stone Serif and Stone Sans by Westchester Publishing Services.

Library of Congress Cataloging-in-Publication Data

Names: Bown, Ollie, author.

Title: Beyond the creative species : making machines that make art and music /  
Oliver Bown.

Description: Cambridge, Massachusetts : The MIT Press, [2021] |

Includes bibliographical references and index.

Identifiers: LCCN 2020017931 | ISBN 9780262045018 (hardcover)

Subjects: LCSH: Technology and the arts. | Arts—Technological innovations.

Classification: LCC NX180.T4 B68 2021 | DDC 700.1/05—dc23

LC record available at <https://lccn.loc.gov/2020017931>

10 9 8 7 6 5 4 3 2 1