MUSICAT: A COMPUTER MODEL OF MUSICAL LISTENING AND ANALOGY-MAKING

Eric Paul Nichols

Submitted to the faculty of the University Graduate School
in partial fulfillment of the requirements
for the degree
Doctor of Philosophy
in the Departments of Computer Science and Cognitive Science,
Indiana University
December 2012

Accepted by the Graduate Faculty, Indiana University, in partial fulfillment of the requirements for the degree Doctor of Philosophy.

Doctoral Committee:

Dr. Douglas R. Hofstadter
(Principal Adviser)

Dr. Eric Isaacson

Dr. Donald Byrd

Dr. Michael Gasser

December 7, 2012

For Helga and Steve

Acknowledgements

So many people have been friends and colleagues, providing emotional and academic support over my eight years in Bloomington, it seems difficult to list them all. Please accept my apologies for any omissions.

First, I would like to thank my immediate family for a lifetime of support: my parents Ann and Clay and my three little sisters Jenni, Sara, and Julie. I'm lucky to have fairly frequent communication with everyone thanks to email and cell phones; in a way I feel closer to my family in Bloomington, thanks to technology, than when I was an undergraduate at Montana State in Bozeman. A side note: my mom used to run a cat-breeding business, and we gave musical names to all the cats. The business was named "Mewsicats Cattery", an amusing source of inspiration for the name of my project (Musicat).

Second, I owe a debt of gratitude to a host of my professors at Indiana University. My first class in Bloomington was an excellent artificial intelligence course taught by David Leake that got me off to a great start. Soon after I started at IU, Chris Raphael joined the faculty, and after auditing some of his courses, he and I had some very successful collaborations on various music informatics projects. I consider him an unofficial member of my thesis committee, and I'm happy about the hours of mathematical and musical discussions we've had over the years. My official committee has also been extremely helpful. Eric Isaacson taught a seminar on Music Cognition during my first semester, which, in combination with Dr. Leake's course, was a perfect way to start my grad school career; it gave me a great introduction to the field. Mike Gasser's courses on biologically-inspired computing and natural language processing were also extremely insightful. Don Byrd was a great musical presence in Music Informatics group meetings and at ISMIR conferences, and always has given me great feedback on this project. Finally, it goes without saying that my

main advisor, Doug Hofstadter, has been a tremendous influence — not only during my eight years at IU, but indeed ever since I was 18 and started reading *Gödel, Escher, Bach.* I'm especially grateful that he made it possible for me to tag along to Paris in 2010 — *il m'a emmenné dans ses valises,* as they say — when he spent a semester there on sabbatical. I expected to learn about artificial intelligence when I came to Bloomington, but Doug actually taught me about cognitive science, communication, writing, and even French pronunciation, among many other things. It's truly been an honor and a privilege!

Third, thanks to my many grad student compatriots from Bloomington: current and past students including Mike Brady, Trista Chan, Annet Czeglédi, Yupeng Gu, Yushen Han, Chris Harshaw, Chris Honey, Andrew Kalafut, Ian Knopke, Olga Rass, and Toshi Uchino; fellow FARGonauts Dave Bender, Harry Foundalis, Ben Kovitz, Francisco Lara-Dammer, Abhijit Mahabal, Damien Sullivan, and Will York; and cyclists extrordinaires George Kachergis and Richard Veale. While in Paris, although I was only there for a semester, I was fortunate to make four great friends: Cinthia Altieri, Tatiana Catanzaro, Rim Roudies, and Junie Terrier. Special thanks to my FARGonaut friends and housemates Matthew Hurley and Alexandre Linhares, both of whom were extremely influential in both my academic and personal life — many great discussions, fun times, crazy projects, and general hilarity!

Fourth, I would like to thank my other close friends in Bloomington (and elsewhere), who formed a sort of surrogate extended family. Alia Alkasimi quickly became my best friend and we've remained like family even after she moved out of town — *merci pour tous, liwliw*. She has also been one of the best sources of moral support helping me succeed with this dissertation, not to mention life in general. Other dear friends include Annika Wallendahl, Lisa Cooper, Zach Saul, Jing Pan, Claire Houck and her dog Jonie, the musically gifted, wonderful Welsh girl Sara Angharad, the incomparable human Lindsay "Z" Arcurio, my "alien twin" Quynh Neutron, my extreme-waltz partner Katherine Wiley, and, of course, the

amazing and artistic musician Priscilla Borges — *obrigado por todos as cores!* Thanks to you all for all the support, friendship, and love.

Finally, on a more somber note, thanks are due to both Helga Keller and Steve Larson, but unfortunately they both passed away in 2011 — a double blow to the CRCC family. Every dissertation like this one from FARG mentions Helga's support and contribution to the Center and her great impact on the lives of the graduate students at FARG. Helga was instrumental in my decision to come to Indiana University, not only by putting me in touch with Doug Hofstadter via phone but also by sending me a surprisingly large collection of CRCC technical reports, which made me very excited to move to Bloomington. Throughout my Ph.D., she was a tremendous help with administrative matters, including lots of extra work on an NSF grant application, but more importantly I enjoyed dinner parties she hosted, a campus tour she gave to my parents, random conversations on quiet days at the office, and a spontaneous road trip to southern Indiana with her and Howard. Helga's presence at CRCC, and in Bloomington, is deeply missed.

Steve Larson, like Helga, directly influenced my decision to come to IU. We met online when he responded to a discussion group post in which I asked where I should apply to grad school. He recommended that I read *Fluid Concepts and Creative Analogies* and encouraged me to apply to IU in Bloomington. We met in person in Eugene, Oregon and discussed his Seek Well project, which led to this dissertation. My meetings in person and via phone with Steve over the years were extremely valuable and his ideas will continue to influence those of us in music cognition for years to come. Thanks, and so long, Steve.

In addition to all the social support mentioned above, I would also like to acknowledge several sources of financial support for this work. I was generously supported by fellowships from CRCC and from the Music Informatics program at IU, as well as by NSF Grant IIS-0738384, awarded by the NSF's CreativeIT program.

Eric Paul Nichols

Musicat: A Computer Model of Musical Listening and Analogy-Making

What happens when people listen to music? What sorts of mental structures are formed? How do we make sense of a melody as its notes fly by in rapid succession? Can we model the experience of listening to music in real time?

This dissertation describes the computer program Musicat, which was designed to explore these questions. Musicat is a model of real-time melody perception by people. The program "listens" to monophonic Western tonal melodies one note at a time (presented not as audio recordings, but rather in a symbolic form much like sheet music) and generates an internal representation of the musical structures it "hears". These structures include groups of adjacent notes, meta-groups comprised of smaller groups, expectations for upcoming structures, and, most importantly, analogies between groups (and meta-groups) of various sizes. In the model, listening is not a passive process; instead, it is an active, dynamic process of creating mental structures. Thus when Musicat listens to a melody, I consider such an act to be creative, and I call it a "listening performance".

Musicat follows in the tradition of previous computer models of the Fluid Analogies Research Group: it is based on the architectures of Copycat, Tabletop, and related programs, with modifications made to accommodate music's temporal nature. In addition to giving a technical discussion of Musicat's architecture, this dissertation includes copious examples of the program's listening performances on a variety of melodies ranging from simple children's songs to more complex popular-song forms. A discussion of the results suggests that Musicat captures several previously unmodeled aspects of musical listening, such as analogy-making and hierarchical grouping in real time, but that much more work is needed to achieve humanlike listening performances on even the simplest of melodies.

Table of Contents

List of Figures

List of Tables