Field Work in Computational Creativity

Margareta Ackerman[@*] **and Rafael Pérez y Pérez**[+*]

[@]Computer Engineering Department, Santa Clara University. `mackerman@scu.edu`.
[+]Depto. de Tecnologías de la Información, Universidad Autónoma Metropolitana. `rperez@correo.cua.uam.mx`
[*]Both authors contributed equally to this work. Alphabetical ordering is used to signal equal contribution.

Abstract

Numerous scientific disciplines, e.g. social sciences, benefit from field work. What does field work look like in Computational Creativity and what are its potential benefits to research? We refer to the effort of actively making a system or its artifacts widely accessible outside the academia world, and as such getting feedback, as 'field work'. In this paper, we reflect on our experiences taking our systems, Alysia and MEXICA, out into the wild terrain by making them broadly available. In the case of Alysia, the system itself was made accessible; MEXICA's artifacts (stories) were shared through a traditionally published book for a broad readership. We consider the utility of field work for these vastly different systems on the CC continuum (Pérez y Pérez 2018), and discuss potential benefits to other research in the area. Finally, we discuss the necessity of developing methodology to enable rigorous registration of knowledge arising from field work in Computational Creativity.

Introduction

An inherently interdisciplinary field, Computational Creativity (CC) draws inspiration from classically creative domains, such as art, literature, and music. Domain experts and (non-expert) broad audiences alike can interface with CC research by providing feedback in the evaluation of autonomous and co-creative systems (Jordanous 2012). Experts can ascertain whether an artifact measures up to higher standards in their creative domain, while feedback from a large broad audience can further help determine the value of an artifact or identify whether a co-creative system is successful in supporting human creativity. Attaining broad audience and expert feedback requires stepping outside of the computational creativity community, often outside the computing discipline all together. We refer to the effort of actively making a system or its artifacts widely accessible, and as such getting this type of feedback, as 'field work'.

Unfortunately, as in other disciplines, field work in computational creativity is complex and time-consuming. Further, it may actually conflict with goals put forth by tenure and promotion committees, which often promote a narrow understanding of scientific contributions. Proof-of-concept systems with minimal user interfaces (UI) more than suffice for achieving primary scientific goals, demotivating researchers from devoting time to making their co-creative systems ready for broad user feedback. Creators of autonomous CC systems aren't directly incentivized to find ways to share the artifacts of their systems with broader audiences.

While many CC systems stay within the academic realm, quite a few researchers have already shared their work with broader audiences (Twitter bots (Veale 2015), DARCI (Norton, Heath, and Ventura 2015), The Painting Fool (Colton 2012), Impro-visor (Keller and Morrison 2007), etc.)

In the absence of formal methodology for CC field work, it is not surprising that work in this direction previously focused on reporting the experience rather than reflecting on the essence of sharing CC systems and artifacts with broad audiences. For instance, one of the largest efforts in the broad exposure of CC has been the musical *Beyond the Fence*, where a multitude of CC systems were used to aid in the creation of a staged musical production. In the 2016 paper on this monumental event (Colton et al. 2016), the authors explain that "This paper acts primarily as a record of the project which led to the Beyond the Fence musical and Computer Says Show documentaries." Similarly, (Colton and Ventura 2014) share that the focus of a CC festival that they organized was to "expose audiences to the main ideas of Computational Creativity within a culturally relevant setting, rather than to study audience experiences. Hence, we did not undertake experiments to gauge reactions to the ideas, systems and outputs presented."

The aim of the current paper is to reflect on the experience of making CC systems and artifacts broadly available, and begin to pave the way towards a methodology for field work in computational creativity. To this end, we share insights resulting from our experience actively sharing our CC systems and artifacts with the outside world. We hope that our analysis may help other researchers decide whether broad exposure is appropriate for their research, encourage others who engage in field work to share their unique insights, and

ultimately lead to a methodical approach to field work in computational creativity.

In this paper, we employ two different systems as frameworks: Alysia (Ackerman and Loker 2017) and MEXICA (Pérez y Pérez and Sharples 2001). Alysia is a co-creative songwriting system. The system originally enabled the creation of vocal melodies for user-provided lyrics. After deciding to increase access to the system, several cycles of extensive user feedback led to radical improvements, including the integration of a co-creative process for lyrics and in-app voices. Alysia was launched on the App Store in January 2019, allowing anyone, regardless of their musical expertise or training, to easily create original songs.

MEXICA is an agent that produces narratives about the old inhabitants of what today is México City. The MEXICA project aims to contribute to the understanding of the creative process. For many years, MEXICA has lived "isolated" inside a laboratory. In December 2017, for the first time, the agent's stories reached a much broader audience, most of whom did not have computer science or cognitive science backgrounds. This was completely new territory for both MEXICA and its author, Rafael Pérez y Pérez.

This paper attempts to reflect on the broad exposure on Alysia, MEXICA and their designers. Based on the CC continuum(Pérez y Pérez 2018), Alysia is focused on supporting the creative process of human beings while MEXICA attempts to contribute to the understanding of the creative process. We hope to illustrate that, although the fundamental intentions of each of these systems are different, both benefit from exposure to (potentially) massive audiences.

Furthermore, the study of such agents, within the framework we are proposing, allows contrasting their main characteristics: because Alysia is a co-creative system, the audience interacts with both the system itself and its product; while in the case of MEXICA, the audience interacts with an artefact that has gone through a human production process: a book. We hope that this joint exploration will give the reader a broader perspective than considering the systems separately.

Alysia Field Work

Alysia is a co-creative system, made with the aim of helping anyone create original songs. Much like EMI (Cope and Mayer 1996), which was created to help David Cope get out of writer's block, Alysia was originally made to support my (Margareta's) desire to write original songs. The first version of Alysia(Ackerman and Loker 2017), which took three months to create, allowed me to write original songs for the first time (notably, after several years of failed attempts at doing so using traditional methods). I was deeply inspired by the first-hard experience of making co-creative CC systems that successfully addressed a challenge that I have been facing for years.

At the time, the system consisted of a co-creative process for making original melodies for user-provided lyrics. The integration of musical generation with natural language processing was a significant research challenge.

The project quickly became central to my research program, and, unexpectedly got the attention of the media even before its original publication, when it was put on Arxiv. After New Scientist[1], NBC News[2], and others released articles featuring Alysia, users began to contact us asking to interact with the system. At the time, Alysia was a young system that lacked a user interface and was difficult to install, and as such could not be shared with anyone beyond one-on-one demos.

Driven by the co-creative goals of the research, I wondered how far the interactive aims of Alysia could go. As Alysia was central in helping me achieve my musical goals, I was inspired to explore the bounds of its co-creative potential by sharing it with others. Why not share it with everyone by putting it online? Since all of my work until that point was done in the academic context, it took about two years before I finally decided to take the steps to make Alysia accessible.

I expected that putting Alysia online with a minimal user-interface will give the freedom of self-expression through songwriting to the masses. This did not turn out as expected, which my team and I quickly learned through large volumes of user feedback. It is worth noting that as a publicly available system, user feedback was significantly more extensive, varied, and direct than the feedback we had gotten through more a controlled user survey we did in the academic context[3].

We heard from a wide range of users of different ages, musical expertise, and stylistic preferences. Feedback came in a variety of ways. Perhaps most useful was unsolicited feedback of users sharing their experience (positive and negative), and asking for new features. Our team also conducted extensive one-on-one user studies and in-depth surveys.

It quickly became apparent that, even though it was sufficient to let me create original songs, at the time, Alysia was not yet able to provide the same support to everyone else who wished to express themselves musically. It only solved one part: Creating original melodies for lyrics - and even that part needed work. Exposing Alysia to a broad audience quickly revealed its shortcomings, and showed me how far they were from our goals.

---

[1] https://www.newscientist.com/article/mg23231043-500-machine-learning-lets-computer-create-melodies-to-fit-any-lyrics/

[2] https://www.nbcnews.com/mach/technology/machine-made-melodies-spotlight-artistic-partnership-between-ai-humans-n698486

[3] An unpublished manuscript on the system included a user study that compared ALYSIA's rankings of vocal melodies to how humans would rank the same melodies. While helpful, the insights resulting from that study was limited in scope, particularly when compared with the wealth of diverse feedback received when taking Alysia to broader audiences.

Figure 1: Alysia demo in a middle school. Line up of students waiting to try Alysia first hand. One of many facets of evaluating Alysia via broad audiences pre-launch. Student faces blurred to protect privacy.

Users asked us to assist with lyrics creation, many did not have the production skills to use Digital Audio Workstations to flesh out Alysia's melodies, and many more struggled with singing. We got a lot of feedback on the melodies themselves: An expert musician pointed out that the melodies consist of too many large intervals, while users wanted melodies that are less varied in both pitch and rhythm (notably, this finding directly conflicted with the user study that we had conducted in the academic context, where ALYSIA's melodies appeared too monotone based on subjects' preference for more varied tunes in their ranking). If we were to help people create songs, we had to do a lot more work. This gave our team a huge push forward, and we completed what was originally supposed to be a five year research plan (and more) in less than a year.

We radically improved the melody model, built a co-creative lyrics generator, and integrated in-app voices. We also put forth a new process for song creation that integrated human-made background tracks on top of which the lyrics, melodies, and vocals were created. This led to an end-to-end system that finally achieved the original goal: In extensive user studies, we observed people with no musical experience easily create songs for the very first time.

Thousands of songs have been created with Alysia before it launched on the app store on January 17, 2019. This put the research to an even more rigorous test, by providing large amounts of data that can be used for evaluation. Explicit user feedback is inherently limited and at times misleading. For instance, the most requested feature on an early Alysia beta was the ability to change musical keys. We rushed to add it, only to discover that hardly anyone ever used it.

Implicit feedback, enabled through user data, pro-vides an exceptionally direct form of evaluation for co-creative systems. Logs reveal how useful users find the co-creative process, letting us see how often they rely on Alysia's generations, how much they modify them, and how often the users input their own melodies and lyrics. Overall use and retention can be used to gauge the utility of the system.

It is worth pointing out that making Alysia widely accessible required much time and effort spent on UI, marketing, and other activities that are not traditionally integral to research. The effort required to make a co-creative system public is neither feasible nor appropriate for all co-creative systems.

For the Alysia project, making the system accessible to a broad audience led to extensive and rigorous evaluation and radically accelerated research. The impact of unabashedly direct, continuous user feedback cannot be overstated. I believe that Alysia had to be placed in the wild-wild west of the broad consumer market in order to reach its potential. It is possible that other co-creative systems may find similar benefits from wide exposure. Further, it may be worth considering how the CC community may be able to facilitate the exposure of our systems and their artifacts.

MEXICA field work

In December 2017, the book "MEXICA 20 years – 20 experiences"(Pérez y Pérez 2017) started to circulate. The volume includes 20 narratives, each in Spanish and English, generated by a computer agent. Its goal is to offer a different reading experience to the general audience; so, the book does not have technical or scientific intentions.

Previously, MEXICA's stories have been published in scientific journals (Pérez y Pérez and Sharples 2004; Pérez y Pérez and Sharples 2001; Pérez y Pérez 2015a), conferences (Pérez y Pérez 2014), book chapters (Pérez y Pérez 2015c; 2015b), web pages[4], specialized talks and so on. However, no such events produced the attention that the book has generated. Why? It is true that the system was improved before generating the tales for the book; for instance, it now generates stories in Spanish and English, the predefined texts are richer (e.g. now it is possible to use pronouns), the analysis of the coherence during the generation process is more robust, the evaluation process is more elaborated, and so on. However, the essence of the system, and therefore the soul of the stories it produces, are essentially much the same. So, why is the general audience giving the system more attention than ever before?

Having a physical book has been a key factor. A book is a familiar artefact that many feel comfortable with. Manuscripts have been amongst us for centuries and so people do not feel threatened by them. In this case, the originality of the cover and the quality of the printed volume has also helped. I also claim that the stories generated by the system are interesting enough to grab

---

[4]http://www.rafaelperezyperez.com

the attention of some readers. Publishing a book usually brings prestige to human authors, sometimes even fame. Thus, when a creative agent becomes the author of a book, there is a good chance that some people will notice it.

Some research domains (e.g. literature) heavily relies on books. Thus, "MEXICA 20 years – 20 stories" works as a bridge that allows me to interact with colleagues from fine arts and literature programs in ways that I did not anticipate. To illustrate my point I would like to introduce Andy Fitch, who is a writer, an editor, teaches in the University of Wyoming's MFA program, and directs the MA program in literature. He interviewed me for the blog of LA review of books. Thus, the book became a piece that both a computational creativity researcher and a writer felt comfortable discussing. In the following I will show some of the questions that Andy asked me that made me reflect about MEXICA and CC in a different way.

Consideration 1. Social and cultural aspects

> Creative computing's emphasis on cross-cultural inputs also stands out. Harrell's preface describes this field asking (singing, actually): Must computers always express the voice of the colonizer — could a computer instead express the voices of sovereign indigenous peoples, the oppressed, and the otherwise underrepresented? Could you place these particular questions in a broader cultural context in which we see, for instance, AI processes often absorbing racist biases circulating in U.S. culture, and then further entrenching and institutionalizing those biases? Does MEXICA work against such trends?(Fitch 2018)[5]

It is hard to find people in the CC community interested in studying and analyzing the cultural aspects of CC. A notable exception is the workshop in Computational Creativity and Social Justice, organized by Gillian Smith, Dan Brown and Anne Sullivan, that took place during ICCC17. The organizers wrote a report that is available online [6]. In this document they pointed out, among other things, the necessity of respecting audiences and cultures:

> Whose voices are represented in our current technologies, and how does this influence the design of CC technologies? How does the (Irani et al. 2010) notion of postcolonial computing relate to CC? How can we infuse our work with respect for the cultural roots of creativity?....

I believe these are very important questions that seem to be in harmony with Andy Fitch (and Fox Harrel)'s concerns about culture. However, unfortunately, during
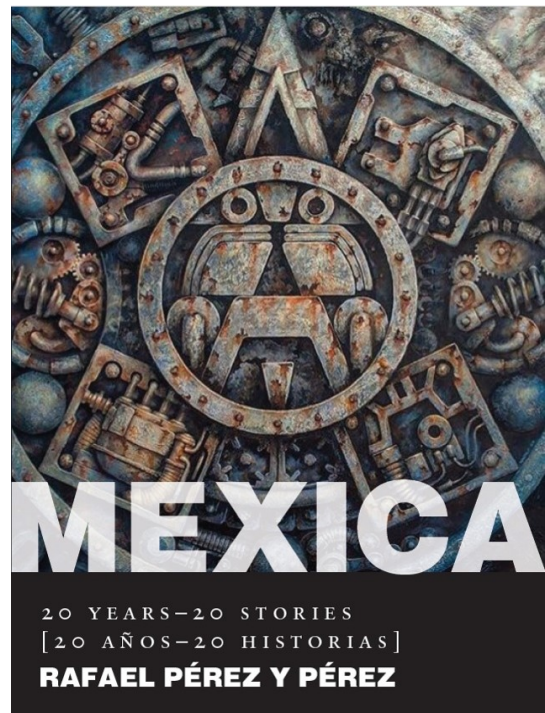

Figure 2: The cover of the book "MEXICA: 20 years, 20 stories"

ICCC18 there were no submissions to a similar workshop and, as a result, it was cancelled. In my experience, CC researchers and students rarely discuss such topics. By contrast, colleagues from the humanities and social sciences seem to believe that creative agents, like MEXICA, might contribute to the debate of these themes. As an example, I was invited to talk about the book and the MEXICA program in:*Crossborders: the aesthetic of immigration* organized by colleagues of the Department of English at the University of Colorado, Boulder, and Counterpath Press in Denver. The organizers describe the event as

> a collaborative transnational project that interrogates the cultural and artistic questions that develop from LatinX migration. CrossBorder includes leading scholars, artists, and writers from both sides of the Southern border who are creating work that directly deals with migration, both literal and cultural, of LatinX populations.[7]

CU Boulder Today advertised the meeting as follows: "The free public event series is meant to foster cross cultural dialogues about human migration from Latin America. It will offer attendees a unique combination of academic research and artistic expression" [8]

---

[5]http://blog.lareviewofbooks.org/interviews/computational-cognitive-social-talking-rafael-perez-y-perez/

[6]https://cs.uwaterloo.ca/conferences/ccsjw2017/CCSJW17 WorkshopReport.pdf

[7]http://counterpathpress.org/crossborders-the-aesthetics-of-migration-at-counterpath-and-cu-boulder-november-9-and-10-2018

[8]https://www.colorado.edu/today/2018/11/06/new-

MEXICA was part of the "artistic expression." Why was I invited? I do not have any idea. Nevertheless, the organizers seemed to be happy with my participation. As a result of all these experiences, I started to wonder how the CC community can contribute to the analyses and study of relevant social concerns such as those mentioned above. I believe that, as community, this is an issue we need to discuss.

Consideration 2. Fresh perspectives on my research

MEXICA the computer system's calculations might all take place within the confines of a single story. But most readers of MEXICA the book probably absorb much more than one story per sitting... To what extent does the machinery of codex-book construction (with any story always placed alongside other stories, interacting with each other in ways MEXICA does not control, so that encountering the phrase "the princess" in one context might have quite different meanings for readers depending on where else they have encountered that phrase in other MEXICA stories)... seem crucial to the MEXICA project, or seem incidental — just random parts of present-day book circulation? (Fitch 2018)

Andy's comments made me ponder a novel perspective of the whole project. Because MEXICA's outputs are single stories, I have always pictured the book as a collection of unrelated narratives. But readers, at least Andy Fitch, seem to perceive the book's twenty stories as forming a unity. Thus, the idea of improving MEXICA in a way that it can represent the concept of "a collections of related stories", where characters and actions in one tale are somehow related to those in a different story, seems intriguing. This challenge requires figuring out how to build novel knowledge-structures capable of representing more abstract concepts and how to establish relations between diverse stories. This new perspective relates to Andy's next question:

You mentioned MEXICA's minimalist style. I actually thought of minimalist music, with its minimal-event horizons, where you might hear the same note many times, until a very slight change occurs, and this subtle shift suddenly feels like a big deal. By comparison, if MEXICA's princess has a bad mood for four straight stories, and then in the fifth story feels more ambivalent — even just that muddled mood can register as a significant tonal shift. And those types of structured variations can make MEXICA stories feel quite similar and quite distinct at the same time. (Fitch 2018)

Andy's thoughts suggested to me that it is not enough to establish links between events, characters, and scenarios in diverse stories in order to construct a

event-kick-series-exploring-immigration-through-art

Figure 3: Rafael Pérez y Pérez sharing readings from "MEXICA: 20 years, 20 stories" at the Guadalajara's International Book Fair, México.

coherent unity, but also to develop mechanisms that allow for the development of unified aesthetic intention, e.g. minimalism, through the book.

Discussion

We argue that CC systems and their products will benefit from being analyzed and evaluated by people outside of our community. We refer to this task as *computational creativity field work*. Although some colleagues have already moved in that direction (Twitter Bots (Veale 2015), DARCI (Norton, Heath, and Ventura 2015), The Painting Fool (Colton 2012), Improvisor (Keller and Morrison 2007), etc.) we have found it hard to come across papers that reflect on and share insights about such experiences; as a result, our community is missing relevant knowledge.

We have used our experiences with Alysia and MEXICA to illustrate how this field work might operate and the kind of knowledge that we can gain from it. In the following we point out some of the relevant aspects that surfaced from such practices:

- As we have showed in the previous sections, sometimes researchers see their projects only from one specific point of view; we refer to this as "researcher fixation." Perspectives from experts and artists in other fields about our systems and products might help to prevent such fixation and trace future research paths.

Making a system or its artifacts widely accessible naturally gives rise to expert feedback. For instance, we were surprised that some colleagues in social sciences saw a research potential of our systems in their field. This may well lead to exploring entirely novel applications for our models.

- A large number of users using a co-creative system can lead to effective means of evaluating the process, through both user feedback (solicited and unsolicited) and the analysis of user data. We notice that unsolicited feedback was a particularly useful source of information.

- Creative agents capable of representing social concerns might contribute to the study of creativity and to the study of such social phenomena. Furthermore, we strongly believe that, as scientists, we have a responsibility to at least reflect on the social implications of our projects. Similarly, we need to think how our systems might contribute to society. For instance, the reader might imagine how Alysia might help to preserve and spread (part of) the musical tradition of the original inhabitants of the south of México.

- We need to provide broader audiences with artefacts that they feel comfortable with. This is a vital point that needs to be taken seriously. We cannot expect broader audiences to interact with systems that lack accessible user interfaces, or to engage with information or artifacts published through traditional academic channels. The channel should ideally represent the manner in which broad audiences are used to receiving artifacts in the domain, such as books for stories and narrative, gallery showing for art, and concerts for music. Feedback on the process of our co-creative systems is best enabled by making them available through easily-accessible channels (website, App store, etc.) paired with a simple user interface.

- Our systems received different kinds of reactions. We found that Alysia's feedback was more specific, targeting precise features of the agent. Users wanted the system to provide further assistance and to be better tailored to their needs. The co-creative aims for Alysia make this type of feedback most beneficial to further development of the system. By contrast, MEXICA's feedback was more general. Expert feedback helped to frame MEXICA in a social context as well provide valuable insight for directions for future work.

In order to make the most of practices such as those described here, our community needs to develop methodology for its field work. It is out of the scope of this paper to define such methodology; however, we would like to contribute with four initial ideas.

1. Classifications. The methodology should classify the human actors participating in field work. We distinguish the following categorizations:

   - The interdisciplinary team working on a specific project.

   - Experts from other disciplines that are unfamiliar with the goals and methods of CC.
   - Lay audiences outside of academia.

   The last two categories conform what we refer to as "broad audiences." We expect to obtain different information from the last two types of audiences.

   In the same way, we suggest to classify CC artefacts into at least three categories:

   - Co-creative systems and their products
   - Independent systems and their products
   - Products of creative systems

   This classification can be used in conjunction with the Computational Creativity Continuum (Pérez y Pérez 2018) in order to organize the information obtained from broad audiences. For instance, we can compare the feedback that mathematical/engineering oriented co-creative systems receive against the comments that cognitive-oriented co-creative systems receive. Of course, there are other possible classifications.

2. Collection of data. There are several ways to collect data. In this text, log files and unsolicited feedback stand out. Implicit feedback is particularly useful due to its utility in identifying potential pitfalls in the co-creative process, discover which components are most used (and as such potentially most successful) and identify whether the system achieves its objective of improving human creativity in an engaging manner through retention metrics.

   Unsolicited feedback proved to be a valuable source of information. A subject that takes the initiative of providing comments about her experience with CC artefacts clearly is a motivated person that has engaged with such an artifact. Unsolicited comments might come, for instance, from direct messages from a user (email, social messaging, etc). But they also can take the form of interviews, essays, reviews, and so on, pondering a system and/or its products. In this case, besides having an inspired subject, we also need a person capable of making sense, from their own perspective, of the creative agent and its outputs. A good methodology should consider these and other possibilities in order to gain as much insight as possible.

3. How to present a system and/or its products to a broader audience requires a detailed study. We need to create situations where people feel comfortable; encourage the development of products designed to work as bridges with other disciplines; analyze which outputs are better suited to represent particular outputs. The main drawback is the resource-consuming nature of making systems and artifacts broadly accessible. How do we make CC systems and/or their artifacts broadly accessible in a repeatable, effective manner? Is there a minimum accessibility metric that is sufficient to gain broad audience input? Is it

possible to create a joint mechanism for making CC systems (perhaps across a similar domain) broadly accessible without the associated overhead for each individual researcher?

4. Evaluation. As part of the methodology, we suggest to break down evaluation techniques into internal and external. Internal evaluation criteria would include surveys we design and run on our students or CC researchers, and other evaluation methods that occur within the community. External criteria could include the opinion of domain experts out of CC (literature expert, painter, musician, etc.), or the utility of a co-creative system, or artifact created by a CC system, for the general population.

   Internal evaluation has the indisputable advantage of scientific rigor. We have established criteria for what constitutes a legitimate survey, how questions should be phrased, and mechanisms to avoid biasing the subjects. By contrast, external evaluation can be unpredictable and difficult to control. This can help push the researcher outside of the "Hans Horse" phenomenon[9], where we unintentionally overestimate the capabilities of systems we create.

We hope this paper will encourage the CC community to participate in the development of a methodology and engagement in field work. Ultimately, research is about venturing out into the unknown and discovering that which has never been previously found through whatever means necessary. This means different things for different research projects. There are many different dimensions and methods for exploration, where taking CC outside of CC is just one of many possibilities.

References

Ackerman, M., and Loker, D. 2017. Algorithmic songwriting with Alysia. In *International Conference on Evolutionary and Biologically Inspired Music and Art*, 1–16. Springer.

Colton, S., and Ventura, D. 2014. You can't know my mind: A festival of computational creativity. In *ICCC*, 351–354.

Colton, S.; Llano, T.; Hepworth, R.; Charnley, J.; Gale, C.; Baron, A.; Pachet, F.; Roy, P.; Gervás, P.; Collins, N.; et al. 2016. The beyond the fence musical and computer says show documentary.

Colton, S. 2012. The painting fool: Stories from building an automated painter. In *Computers and creativity*. Springer. 3–38.

Cope, D., and Mayer, M. J. 1996. *Experiments in musical intelligence*, volume 12. AR editions Madison.

Fitch, A. 2018. The computational and the cognitive-social: Talking to Rafael Pérez y Pérez. Blog Los Angeles Review of Books.

Irani, L.; Vertesi, J.; Dourish, P.; Philip, K.; and Grinter, R. E. 2010. Postcolonial computing: a lens on design and development. In *Proceedings of the SIGCHI conference on human factors in computing systems*, 1311–1320. ACM.

Jordanous, A. 2012. A standardised procedure for evaluating creative systems: Computational creativity evaluation based on what it is to be creative. *Cognitive Computation* 4(3):246–279.

Keller, R. M., and Morrison, D. R. 2007. A grammatical approach to automatic improvisation. In *Proceedings, Fourth Sound and Music Conference, Lefkada, Greece, July.*"Most of the soloists at Birdland had to wait for Parker's next record in order to find out what to play next. What will they do now.

Norton, D.; Heath, D.; and Ventura, D. 2015. Accounting for bias in the evaluation of creative computational systems: An assessment of darci. In *ICCC*, 31–38.

Pérez y Pérez, R., and Sharples, M. 2001. Mexica: A computer model of a cognitive account of creative writing. *Journal of Experimental & Theoretical Artificial Intelligence* 13(2):119–139.

Pérez y Pérez, R., and Sharples, M. 2004. Three computer-based models of storytelling: Brutus, minstrel and mexica. *Knowledge-based systems* 17(1):15–29.

Pérez y Pérez, R. 2014. The three layers evaluation model for computer-generated plots. In *ICCC*, 220–229.

Pérez y Pérez, R. 2015a. A computer-based model for collaborative narrative generation. *Cognitive Systems Research* 36:30–48.

Pérez y Pérez, R. 2015b. From mexica to mexica-impro: The evolution of a computer model for plot generation. In *Computational Creativity Research: Towards Creative Machines*. Springer. 267–284.

Pérez y Pérez, R. 2015c. Mexica-impro: Generación automática de narrativas colectivas. In y Pérez, R. P., ed., *Creatividad Computacional*. México City: Grupo Editorial Patria.

Pérez y Pérez, R. 2017. *Mexica: 20 Years-20 Stories*. Counterpath.

Pérez y Pérez, R. 2018. The computational creativity continuum. In *Proceedings of the Ninth International Conference on Computational Creativity ICCC*, 177–184.

Veale, T. 2015. Hallando la creatividad en conceptos en conflicto. In y Pérez, R. P., ed., *Creatividad Computacional*. México City: Grupo Editorial Patria.

---

[9]The parable of a man who fooled himself into believing that his horse could do arithmetic, when in fact it was merely attuned to its owner's facial cues.