
Mixed-Initiative Creative Interfaces

Sebastian Deterding**Jonathan Hook**

University of York
York, YO10 5GE, UK
sebastian.deterding@york.ac.uk
jonathan.hook@york.ac.uk

Rebecca Fiebrink**Marco Gillies****Jeremy Gow****Memo Akten**

Goldsmiths, University of London
New Cross, London SE14 6NW, UK
r.fiebrink@gold.ac.uk
m.gillies@gold.ac.uk
j.gow@gold.ac.uk
memo@memo.tv

Gillian Smith

Northeastern University
Boston, MA 02115, USA
gi.smith@neu.edu

Antonios Liapis

University of Malta
Msida, MSD2080, Malta
antonios.liapis@um.edu.mt

Kate Compton

UC Santa Cruz
Santa Cruz, CA 95064, USA
kcompton@soe.ucsc.edu

Abstract

Enabled by artificial intelligence techniques, we are witnessing the rise of a new paradigm of computational creativity support: *mixed-initiative creative interfaces* put human and computer in a tight interactive loop where each suggests, produces, evaluates, modifies, and selects creative outputs in response to the other. This paradigm could broaden and amplify creative capacity for all, but has so far remained mostly confined to artificial intelligence for game content generation, and faces many unsolved interaction design challenges. This workshop therefore convenes CHI and game researchers to advance mixed-initiative approaches to creativity support.

Author Keywords

Mixed-initiative interaction; creativity support tools; computational creativity; computer-aided design; interactive machine learning; artificial intelligence

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous; D.2.2 Design Tools and Techniques.

Introduction and Background

Creativity support is an important area of computing and human-computer interaction (HCI) [1]. Historically, work on such creative interfaces broadly clusters around two ends of a spectrum [2,3] (fig. 1): on the

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.

Copyright is held by the owner/author(s).

CHI'17 Extended Abstracts, May 06–11, 2017, Denver, CO, USA

ACM 978-1-4503-4656-6/17/05.

<http://dx.doi.org/10.1145/3027063.3027072>

□




Figure 1: The spectrum of agency and initiative distribution in creative interfaces (inspired by [2,3]).

one end, we find traditional creativity support tools like computer-aided design or search, visualization and collaboration software for creative work [1]. Here, the human is the initiating and deciding agent and the computer a mediating tool. On the opposite end sits computational creativity [4]: the computer is a 'heroic' [2] artificial intelligence (AI) agent that autonomously produces creative work, and the human its audience.

Recent years have seen the emergence of a third paradigm between the two: *mixed-initiative creative interfaces (MICI)* where human and computer interact as *collaborators* in a tight feedback loop [5,6]. As in human-human creative dialogue, both sides take turns constraining, suggesting, producing, evaluating, modifying, or selecting creative outputs in response to the other, such that creative agency and initiative cannot be easily ascribed to one side alone. While all creative practice is to some extent a 'dialogue' between creator and material [7], MICI literalize this metaphor by giving the computer the status of creative agency and initiative thanks to AI.

The vision of augmenting human problem-solving by sharing initiative within a larger human-computer "symbiont" reaches back through the history of HCI and AI to pioneers like Licklider, including early work on dialogue-based interaction with conversational agents and intelligent tutors [8-11]. Yet in creativity support, it




Figure 2: Tanagra, a MICI 2D game level generator: the designer adds, removes or modifies high-level 'beats' of gameplay and their length (like 'jump to kill an opponent'), the AI executes level geometry in detail, generates beats undefined by the designer, and continually tests whether the level remains playable [22].

is only today that we find this vision realized, chiefly in the field of procedural content generation (PCG) for games [6]. Here, a lively community of researchers and practitioners is currently exploring AI techniques such as evolutionary computation, heuristic search, or machine learning to (semi-)automatically create and evaluate art assets, game levels, or even entire games (see figures). Current mixed-initiative, human-AI co-creativity in PCG covers a fraction of possible MICI scenarios. Still, these systems already show a number of highly attractive features: they vastly accelerate the iterative exploration of solution spaces, which also enables divergent exploration that would be prohibitively time-consuming otherwise; they enable and accelerate learning and understanding creative practice through probing and hypothesis-testing; they provide rich lateral stimuli; and they can make creative




Figure 3: The Wekinator, an interactive machine learning MICI that allows designers to create new controllers, instruments and the like: Designers demonstrate desired input-output pairings to the AI, which learns to translate them into a controller scheme, which the designer can then interactively adapt with new sample demonstrations. [21]

practice accessible and enjoyable to non-professional and even disabled user groups, with rich benefits to personal wellbeing and societal inclusion [5,6,12-14].

Yet although MICI and their advantages easily generalize across creative practice, they have so far remained largely disconnected from creative interface researchers in HCI. Examples of MICIs outside games demonstrate the relevance of the approach to a broad range of creative contexts under active consideration by HCI researchers, including: urban design [16], sketching [15], interface design [13], prototyping musical instruments [14], or data visualization [18]. Indeed, with the current rise of AI, (semi-)autonomous systems, and conversational agents [19,20], human-AI mixed-initiative presents a generally valuable and underexplored interface paradigm.

While HCI could richly benefit from current MICI work in game AI, MICIs in turn present a number of formidable interface and interaction design challenges: enabling creators with little computational literacy to readily express ideas, constraints, or criteria in a formal(izable) and thus, computable manner; avoiding user fatigue and making interaction with the machine engaging; rendering computer design decisions legible and transparent to users; identifying visualization patterns for computational design evaluations and modification and selection suggestions [6]. Since MICIs in games have been chiefly designed by AI researchers for similarly computationally literate user groups, they would greatly benefit from the design expertise of HCI researchers.

Workshop Goal and Questions

In summary, mixed-initiative creative interfaces, as currently developed in PCG for games, hold rich opportunities for creative and (semi-)autonomous interfaces far beyond games. At the same time, they present unsolved interface and interaction design challenges that would immediately benefit from HCI and interaction design. This workshop therefore brings together game AI and CHI communities to advance the MICI paradigm by exploring its opportunities for HCI beyond game design and address its unsolved interface design challenges. We wish to facilitate collaboration around questions such as:

- How do people experience and interact with MICIs?
- How can we evaluate the quality of human-AI co-creativity?
- How do MICIs challenge and advance creativity support tools, computational creativity, and human-human creative collaboration?
- What interaction and interface design challenges do MICIs hold? What established design methods, principles, and patterns can address these?
- What AI challenges of MICIs have HCI solutions, what HCI challenges have AI solutions?
- How can MICIs up-skill not de-skill end users?
- What new kinds of human-AI co-creativity can we envision across and beyond creative practice?
- What opportunities and challenges do different genres and forms of creativity hold for MICIs?
- What are human-AI creative relations beyond collaboration, such as competition (e.g. DJ battles), rivalry or disruption?
- How might MICIs further accessibility, inclusion, and participation?




Figure 4: Sentient Sketchbook, a MICI game map generator: the designer creates a low-detail map, the AI evaluates its playability and evolutionarily generates better alternative suggestions, which the designer can then choose and modify. The AI can turn any low-detail into a high-detail map [23].

Organizers

Sebastian Deterding (contact author) is a senior research fellow at the Digital Creativity Labs at the University of York. His work explores the use of game design methods and principles beyond games, more recently in intelligent, data/AI-driven interfaces. He has chaired workshops on gamification at CHI'11, '13, and '15, and on embarrassing interactions at CHI'15.

Memo Akten is an artist and researcher at Goldsmiths, University of London exploring collaborative co-creativity between humans and machines for artistic expression. His collaboration with Quayola, 'Forms', received the Prix Ars Electronic Golden Nica in 2013. Past exhibitions and performances include the Victoria & Albert Museum, London; Royal Opera House, London; Garage Center for Contemporary Culture, Moscow; La Gaîté Lyrique, Paris; Holon Design Museum, Israel and the EYE Film Institute, Amsterdam.

Kate Compton is a game designer, artist, and PhD Studio. Her work focuses on the development of AI tools to augment user creativity, especially in casual or playful audiences. She specializes in designing and implementing systems that assist users in quickly moving through the possibility space of a creative problem, a genre she calls Casual Creators.

Rebecca Fiebrink is a Lecturer in Computing at Goldsmiths, University of London. Her research focuses on using machine learning as a tool for designing interactive systems, especially systems for musical expression and embodied interaction. She co-organised the CHI'16 workshop on human-centered machine learning and co-chaired the 2014 conference on New Interfaces for Musical Expression.

Marco Gillies is a Senior Lecturer in Computing at Goldsmiths, University of London. His work concerns tools for designing movement based and body language based interaction, particularly for virtual reality. He co-organized the CHI'16 workshop on human-centered machine learning and will be conference co-chair of the Movement & Computing conference (MOCO) in 2017.

Jeremy Gow is a Lecturer in Computing at Goldsmiths, University of London. His research explores computational creativity and game design: how can artificial agents understand play and participate in making games, collaboratively or autonomously? He co-organised the CHI'10 workshop on video games as research instruments, and AISB symposia on game AI (2014) and computational creativity (2014, 2015).

Jonathan Hook is a Lecturer in the Department of Theatre, Film and Television at the University of York. His research draws on human-centered design methods to develop novel interactive technologies for creative practitioners. He has chaired workshops on interaction design for creativity at CHI'12, CHI'13 and DIS'14.

Antonios Liapis is a Lecturer at the Institute of Digital Games, University of Malta. His research sits at the crossroads of game design, artificial intelligence and




Figure 5: Funky Ikebana, a MICI art generator: the AI evolves a number of flower species, the designer selects a preferred one to direct the AI to evolve the next set of species from the DNA of the selected species [12].

computational creativity. More specifically, he explores the limits of computational input to the human-driven design process in computer-aided design tools such as Sentient Sketchbook. He has chaired workshops on game aesthetics (AIIDE'13), PCG (FDG'15), experimental AI (AIIDE'15-16) and computational creativity in games (ICCC'16).

Gillian Smith is an Assistant Professor in Art+Design and Computer Science at Northeastern University. Her research sits at the intersection of HCI and AI, with a focus on computer creativity, computational craft, and how humans interact with generative systems. Recently, she has also been studying the history of procedural content generation and mixed-initiative creativity. She has organized several workshops on artificial intelligence and design, including Artificial Intelligence in the Design Process (AIIDE'11 & '13), procedural content generation (FDG'13), and {craft|game}play (FDG'15).

Workshop Organisation

Website

We will advertise the workshop and host its papers and videos at <http://mici2017.org>.

Pre-Workshop Plans

We expect to draw interest from researchers working on creative interfaces, procedural content generation, interactive machine learning, and games across the CHI, game AI, Intelligent User Interfaces and Creativity and Cognition communities. To ensure a fruitful mix of participants, we will distribute the CFP through our extensive personal networks spanning said communities, social media, as well as relevant mailing lists such as CHI Announcements or Gamesnetwork.

Acceptance to the workshop will be based on a juried review of short (2-6pp.) position papers, aiming for relevance, quality, diversity, and complementarity of overall contributions. We aim for a total of 15-25 participants; in shared papers, we expect only one author to participate. Participants are expressly invited to bring interactive demos, and are required to read all accepted papers ahead of the workshop.

Workshop Structure

Based on past workshop experiences, we aim to avoid fatigue and maximize productive discussion and work with decent breaks and a shorter workshop. Depending on turnout, participants will present their work as a 10 min. panel presentation/demo or coffee break poster/demo. To prime fruitful debate, the first part of the workshop will alternate thematic panels of 3-4 presentations each and a break with thematic poster/demo islands. Each will be followed by a discussion teasing out particular challenges and opportunities that will be collected and clustered in parallel by the organizers on shared boards.

In the second part, we will break up into small groups. Each group will be facilitated by two designated leads (one HCI, one game AI) and develop (a) approaches to an identified cluster of design challenges or (b) uses to a cluster of opportunities, captured in a shared online document and illustrated with one-page mockups. A final plenum session will share group results, discuss emergent issues, and identify follow-up plans facilitated with gamestorming methods.

Post-Workshop Plans

Accepted papers will be published online as CEUR workshop proceedings (ceur-ws.org) and linked on the website ahead of the workshop. Videos of all

presentations will likewise be shared on the workshop website. Results will be written up as an article for *ACM Interactions* or *Communications of the ACM*. We plan to run a mirror event at a relevant game AI conference (e.g. FDG, CIG, or Dagstuhl seminar); based on interest and discussion during both workshops, we will consider further dissemination options.

Call for Participation

Enabled by artificial intelligence techniques, we are witnessing the rise of a new paradigm of computational creativity support: *mixed-initiative creative interfaces* (MICIs) put human and computer in a tight interactive loop where both take turns suggesting, producing, evaluating, modifying, selecting creative outputs. This paradigm could broaden and amplify creative capacity for all, but is today mostly developed and studied in game AI. This workshop therefore convenes CHI and game researchers to advance mixed-initiative approaches to creativity support.

Topics of Special Interest

- User experience and engagement in MICIs
- Interaction, design and evaluation challenges of MICIs
- Future scenarios of human-AI co-creativity and MICIs in and beyond creative practice
- MICIs for accessibility, inclusion, participation, and user-generated content
- MICIs and models of creativity and collaboration

We invite participants to submit 2-6pp. position papers in the CHI Extended Abstract format via EasyChair. We will review submissions and select 15-25 based on relevance, quality, and diversity of inputs. At least one author of each accepted paper needs to register for the workshop and for one or more days of the conference.

We expressly invite participants to bring system demos to the event.

Important dates

- December 11, 2016: Early acceptance submission deadline
- December 21, 2016: Notification of early acceptance
- January 27, 2017: Regular submission deadline
- February 15, 2017: Notification of regular acceptance
- May 6/7, 2017 (tbd.): Workshop

For further information, consult the workshop site mici2017.org or write to mici2017@easychair.org.

Acknowledgements

This work was conducted in the Digital Creativity Labs (digitalcreativity.ac.uk), jointly funded by EPSRC/AHRC/Innovate UK under grant no. EP/M023265/1.

References

1. Ben Shneiderman. 2007. Creativity support tools: accelerating discovery and innovation. *Communications of the ACM* 50, 12: 20–32. <https://doi.org/10.1145/1323688.1323689>
2. Jon McCormack and Mark D’Inverno. 2014. On the Future of Computers and Creativity. *AISB14 Symposium on Computational Creativity*, 1–4.
3. Todd Lubart. 2005. How can computers be partners in the creative process: Classification and commentary on the Special Issue. *International Journal of Human Computer Studies* 63, 4–5: 365–369. <https://doi.org/10.1016/j.ijhcs.2005.04.002>
4. Simon Colton, R Lopez De Mantaras, and O Stock. 2009. Computational Creativity: Coming of Age. *AI Magazine* 30, 3: 11–14. <https://doi.org/10.1609/aimag.v30i3.2257>

5. Antonios Liapis, Georgios N Yannakakis, Constantine Alexopoulos, and Phil Lopes. 2016. Can Computers Foster Human User's Creativity? Theory and Practice of Mixed-Initiative Co-Creativity. *Digital Culture & Education* 8, 2: 136–153.
6. Antonios Liapis, Gillian Smith, and Noor Shaker. 2016. Mixed-initiative Content Creation. In *Procedural Content Generation in Games: A Textbook and an Overview of Current Research*, Noor Shaker, Julian Togelius and Mark J. Nelson (eds.). Springer, 195–214.
7. Donald A, Schön. 1992. Designing as reflective conversation with the materials of a design situation. *Research in Engineering Design* 3, 3: 131–147.
8. Joseph Carl Robnett Licklider. 1960. Man-Computer Symbiosis. *IRE Transactions on Human Factors in Electronics* 1, 1: 4–11.
<https://doi.org/10.1109/THFE2.1960.4503259>
9. Jaime R Carbonell. 1971. AI in CAI: An Artificial-Intelligence Approach to Computer-Assisted Instruction. *IEEE Transactions on Man-Machine Systems* 11, 4: 190–202.
<https://doi.org/10.1109/TMMS.1970.299942>
10. Eric Horvitz. 1999. Principles of mixed-initiative user interfaces. In *Proceedings of the ACM CHI 99 Human Factors in Computing Systems Conference (CHI'99)*, 159–166.
<https://doi.org/10.1145/302979.303030>
11. Marti Hearst, James F. Allen, Eric Horvitz, and Curry I. Guinn. 1999. Mixed-initiative interaction. *IEEE Intelligent Systems and their Applications* 14, 5: 14–23.
12. Kate Compton and Michael Mateas. 2015. Casual Creators. *Proceedings of the Sixth International Conference on Computational Creativity*, 228–235.
13. Simon Katan, Mick Grierson, and Rebecca Fiebrink. 2015. Using Interactive Machine Learning to Support Interface Development Through Workshops with Disabled People. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI'15)*, 251–254.
<https://doi.org/10.1145/2702123.2702474>
14. Hugo Scurto and Rebecca Fiebrink. 2016. Grab-and-Play Mapping: Creative machine Learning Approaches for musical Inclusion and Exploration. In *International Computer Music Conference*.
15. Nicholas Davis, Chih-pin Hsiao, Kunwar Yashraj Singh, Lisa Li, Sanat Moningi, and Brian Magerko. 2015. Drawing Apprentice: An Enactive Co-Creative Agent for Artistic Collaboration. In *Proceedings of the 2015 ACM SIGCHI Conference on Creativity and Cognition*, 185–186.
16. Artem M. Chirkin and Reinhard Koenig. 2016. Concept of Interactive Machine Learning in Urban Design Problems. In *Proceedings of the SEACHI 2016 on Smart Cities for Better Living with HCI and UX*, 10–13.
17. Saleema Amershi, Maya Cakmak, William Bradley Knox, and Todd Kulesza. 2014. Power to the People: The Role of Humans in Interactive Machine Learning. *AI Magazine* 35, 4: 105–120.
<https://doi.org/10.1609/aimag.v35i4.2513>
18. Stephen Makonin, Daniel Mcveigh, Wolfgang Stuerzlinger, Khoa Tran, and Fred Popowich. 2016. Mixed-Initiative for Big Data: The Intersection of Human + Visual Analytics + Prediction. In *Proceedings of the 49th Hawaii International Conference on System Sciences*, 1427–1436.
19. Jelena Stajic, Richard Stone, Gilbert Chin, and Brad Wible. 2015. Rise of the Machines. *Science* 349, 6245: 248–249.
20. Ewa Luger and Abigail Sellen. 2016. "Like Having a Really Bad PA": The Gulf between User Expectation and Experience of Conversational Agents. In *Proceedings of the SIGCHI Conference on Human*

Factors in Computing Systems (CHI'16), 5286–5297. <https://doi.org/10.1145/2858036.2858288>

21. Rebecca Fiebrink, Dan Trueman, and Perry R. Cook. 2009. A metainstrument for interactive, on-the-fly machine learning. In *Proceedings of the 9th New Interfaces for Musical Expression Conference*.
22. Gillian Smith, Jim Whitehead, Michael Mateas. 2011. Tanagra: Reactive planning and constraint solving for mixed-initiative level design. *IEEE Transactions on Computational Intelligence and AI in Games* 3, 3: 201–215.
23. Antonios Liapis, Georgios N. Yannakakis, Julian Togelius. 2013. Sentient Sketchbook: Computer-Aided Game Level Authoring. In *Proceedings of the 8th Conference on the Foundations of Digital Games (FDG'13)*, 213–220.