



# AI as Social Glue: Uncovering the Roles of Deep Generative AI during Social Music Composition

Minhyang (Mia) Suh  
miasuh@google.com  
Google  
Mountain View, CA, USA

Michael Terry  
michaelterry@google.com  
Google Research  
Cambridge, MA, USA

Emily Youngblom  
eyoungblom@google.com  
Google  
Mountain View, CA, USA

Carrie J. Cai  
cjcai@google.com  
Google Research  
Mountain View, CA, USA

## ABSTRACT

Recent advances in deep generative neural networks have made it possible for artificial intelligence to actively collaborate with human beings in co-creating novel content (e.g. music, art). While substantial research focuses on (individual) human-AI collaborations, comparatively less research examines how AI can play a role in human-human collaborations during co-creation. In a qualitative lab study, we observed 30 participants (15 pairs) compose a musical phrase in pairs, both with and without AI. Our findings reveal that AI may play important roles in influencing human social dynamics during creativity, including: 1) implicitly seeding a common ground at the start of collaboration, 2) acting as a psychological safety net in creative risk-taking, 3) providing a force for group progress, 4) mitigating interpersonal stalling and friction, and 5) altering users' collaborative and creative roles. This work contributes to the future of generative AI in social creativity by providing implications for how AI could enrich, impede, or alter creative social dynamics in the years to come.

## CCS CONCEPTS

• **Human-centered computing** → *Empirical studies in HCI*.

## KEYWORDS

machine learning, human-AI co-creation, music composition

### ACM Reference Format:

Minhyang (Mia) Suh, Emily Youngblom, Michael Terry, and Carrie J. Cai. 2021. AI as Social Glue: Uncovering the Roles of Deep Generative AI during Social Music Composition. In *CHI Conference on Human Factors in Computing Systems (CHI '21)*, May 8–13, 2021, Yokohama, Japan. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3411764.3445219>

## 1 INTRODUCTION

Major innovations in deep learning have made it conceivable for humans to partner with artificial intelligence (AI) in co-creating

novel content. With the advent of deep generative networks that are now capable of synthesizing creative new text, images, and artwork, a user might ask an AI to auto-complete a paragraph [12], flesh out a hand-drawn sketch [59], or even compose whole sections of music [54]. However, despite substantial research aimed at improving the performance capabilities of these deep generative algorithms, recent research suggests that it can be challenging for individuals to create music [36, 54] or write [10] in partnership with AI, even when using state-of-the-art algorithms.

People's creative needs can be even more nuanced and complex in real-world contexts, where creative practices are often experienced socially. For example, it is common for musicians to leverage each other's expertise by co-composing a song in pairs, or for script writers to create dialogue in groups. Among everyday musicians, music 'jam sessions' and round-robin creative writing serve important social roles, enabling novices and aspiring hobbyists alike to connect with one another and grow their community. In these collaborative, creative practices, prior research reveals a number of nuances to these collaboration dynamics [62]. For example, a group may feel pressure to achieve premature consensus, resulting in sub-optimal and non-creative solutions [39]. On the other hand, a group can also enable the brainstorming of more novel and diverse ideas, leading to innovation [60]. Given the increasing abilities of generative models to create novel content, along with the nuances of multi-human creative collaborations, there is an opportunity to examine how AI may assist in collaborative, creative practices involving two or more individuals.

In this paper, we investigate how pairs of people collaboratively create artistic content with a deep generative AI, and how the social dynamics can be affected by the presence of AI. We ground our investigation within the music domain, where not only are collaborative practices common, but a wide range of deep generative models has also been made available to the public [15, 16, 34, 64].

In a qualitative lab study, we observed 30 participants (15 pairs) of people compose a musical phrase in pairs, both with and without AI. Our findings reveal that AI may play important roles in influencing social dynamics during creativity by: 1) implicitly seeding common ground at the start of collaboration, 2) acting as a psychological safety net in creative risk-taking, 3) providing a force for group progress, 4) mitigating interpersonal stalling and friction, and 5) altering users' collaborative and creative roles. Together, these findings lay a groundwork for understanding the future of



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs International 4.0 License.

CHI '21, May 8–13, 2021, Yokohama, Japan

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8096-6/21/05.

<https://doi.org/10.1145/3411764.3445219>

generative AI in social creativity, suggesting implications for how AI could enrich, impede, or alter creative social dynamics in the years to come. In sum, our major contributions are:

- (1) An understanding of the various roles of AI in mediating human-human co-creation (as enumerated above)
- (2) Descriptions of how collaborators develop strategies together when co-creating with AI
- (3) Implications and recommendations for what to consider and how to utilize AI in supporting co-creative processes

## 2 RELATED WORK

This paper draws on prior work at the intersection of human-AI co-creation, AI and multi-human collaboration, and technologies for collaborative music creation.

### 2.1 Human-AI Co-creation

Modern machine learning (ML) has made it conceivable for AI to serve as not only a creative tool, but as an active collaborator in *co-creating* novel content with a user. Human-AI co-creation has recently been investigated in domains such as drawing [18, 20, 42, 43, 59], music composition [54], creative writing [26, 37], design ideation [46, 51, 69], visual metaphors [11], video game content generation [28, 47], improvisation [38], public displays [52], and dance [53]. For example, prior studies demonstrated how an AI can collaborate with a person by generating inspirational sketches to assist with ideation [51, 59], writing text to continue a half-written story [12], or creating characters for storyboarding [69].

Within the music domain, ML-powered interfaces have traditionally allowed users to harmonize a melody [71], find adventurous chord progressions [33], or generate novel sounds via gestures [22]. More recently, progress in *generative deep neural networks* (DNNs) has dramatically expanded the possibilities of human-AI co-composition. Beyond small segments of music, generative DNNs can now produce full, high-quality scores that automatically mesh well with local and distant regions of music [29, 34, 35, 64], making it conceivable for users to auto-complete long stretches of music, coherently fill in any missing parts of a song, or ask for new music to be generated from scratch. In examining how these models can aid a person in composing music, two recent studies uncovered the dual opportunities and challenges of co-creating with these powerful generative models. Specifically, Louie et. al., [54] and Huang et. al., [36] illuminated challenges around establishing a human sense of control and creative authorship when co-composing with deep generative AI. Huang et al. [36] illuminated the strategies teams leverage in their attempt to control and decompose AI. However, less attention has been placed on AI's role in affecting the *interpersonal dynamics* within a team itself.

Taken together, the body of research described above forms a basis for understanding human-AI co-creation, with comparatively less known about how this new frontier of ML could and should support co-creation among multi-human partnerships or the *interpersonal dynamics* within teams. Building on this prior work in human-AI co-creation, we provide a foundational perspective on human-AI co-creation in contexts where multiple people co-create together, illuminating how AI may alter the interpersonal dynamics within a team itself.

### 2.2 AI and Multi-Human Collaboration

While integrating AI into the human workflow has always been an important area of research, there is increasing value in introducing Computer-Supported Cooperative Work (CSCW) perspectives in Human-AI collaboration research. Examples of this style of research include emerging research agendas centered on AI as collaborative teammates [67, 73], and AI as a mediator of teamwork and interpersonal communication itself [30, 65]. Along these lines, recent studies have examined how AI influences social, collaborative activities. Many of these focus on how chatbots or bots affect social collaboration in personal [55] and professional [49, 75] contexts, and within online [1, 44, 56, 68, 74, 77] and local communities [50, 76]. Others consider how AI-powered virtual assistants play roles in family contexts [6, 25, 61]. These studies inform how AI could affect collaborative dimensions such as interpersonal dynamics, impression formation, trust, and conflict [3, 30, 41, 74]. Building on this growing body of work examining AI in human-human collaboration, our work contributes to the broader literature by investigating AI's potential role in multi-human *co-creation* of novel, creative content.


### 2.3 Technologies for Collaborative Music Creation

Co-creation is a social creativity process “leading to the emergence and sharing of creative activities and meaning in a socio-technical environment” [48]. Co-creation has long been of interest within the HCI and CSCW communities [2, 23, 24], with music one of the domains studied [24]. Prior work on collaborative technologies for music-making has examined how users use tools and systems while co-composing [7, 31, 57, 66] or co-curating [5, 17] music. For example, Benford et al., [7] observed how Irish musicians utilized various digital resources and social networks to develop their mutual musical knowledge while collaboratively co-creating music. Another set of research has focused on developing novel technologies to support or augment co-creative experiences in music, such as creating new instruments for co-playing (e.g., Jam-O-Drum [8], Reactable [40], and MuSeeCol [27]), or new tools that allow users to collaboratively influence musical output (e.g. volume, timbre) given pre-produced musical materials [21, 58].

Across this range of prior work, a core agreement is that tools and technologies play important roles in collaborative music-making by enhancing group creativity and facilitating social activities that can be engaging and playful [21, 27]. With the advent of machine learning and deep generative neural networks, modern technology has the potential to serve not only as a tool for social music creation, but also as an active team member, collaborator, or mediator. This new wave of generative AI calls for a re-examination of how technology could enrich, impede, or transform collaborative music creation processes.

## 3 METHOD

To understand the potential for AI to assist with co-creation, we conducted a qualitative laboratory study with 30 participants (15 pairs). Using a previously developed interface to a deep generative AI, Cococo [54], we examined the social dynamics and co-creation process of pairs of users co-composing a musical phrase, with and without AI.



**Figure 1: Study version with AI (Cococo): a music editor web-interface for human-AI co-creation with AI**

### 3.1 Cococo: Overview

Cococo (collaborative co-creation) is a web-based music editor for human-AI co-creation that augments standard generative music interfaces with a set of AI steering tools [54]. Built on top of Coconet [32], a deep generative model trained on 4 part harmony that takes incomplete music as input and outputs complete music, Cococo supports an iterative co-creation process by providing the following main components and capabilities (see Figure 1):

- (1) Users can manually draw and edit notes
- (2) Voice Lanes, in which users specify what voices to compose for (e.g. soprano, alto, tenor, bass), and the time range in which the AI-generates music
- (3) A “Generate” button to ask the AI to fill in music given the existing notes on the page
- (4) Semantic Sliders (e.g. more surprising or conventional, or content more similar to or different from a selected segment of music) to adjust musical qualities of the AI-generated notes
- (5) Multiple Alternatives of the AI-generated music, from which users can audition and choose
- (6) While auditioning the AI-generated options, users see the temporarily substituted notes in the music score, shown as glowing notes
- (7) An Infill Mask, which users can use to crop an existing section of notes to be infilled using AI

For the condition without the AI, participants used a version that is visually similar to Cococo, but without any of the AI-related features. Specifically, participants could manually draw and edit


notes, but did not have access to features (2) to (7) on the list/figure above (see Figure 2).

Because the study took place during the Covid-19 pandemic, all studies were conducted remotely over video-conference. During the study, participants in pairs collaborated using the Chrome Remote Desktop extension, which allows one participant to remotely access another participant’s computer through the Chrome browser, and to share a cursor, screen, and audio during their collaborations.

### 3.2 Procedure

Each participant first completed an online tutorial of the interfaces on their own (30 minutes). Then in pairs, they composed two pieces, one with and one without Cococo, with the order of the conditions counterbalanced (about 20 minutes each). As a prompt, users were provided a set of images from the card game Dixit [14] and were asked to compose music that reflected the character and mood of one image of their choosing, similar to prior studies [33, 54]. Users were observed while composing using a think-aloud procedure. Finally, participants compared and contrasted their experiences co-composing with and without Cococo via a semi-structured interview (about 20 minutes).

All studies were screen recorded and transcribed. To analyze the data, following [9], two researchers individually familiarized themselves with the data and created initial codes, then collaboratively discussed these codes to create an initial code book together. They then worked systematically through the data set, attaching codes to sections of text, and refined the codes in a collaborative, iterative process. Then, four researchers discussed the codes and data to collectively create higher-level categories and themes, iteratively



**Figure 2: Study version without AI: the same music editor interface as Cococo (see Figure 1 above), but without the AI-related functionality. Users can still manually draw and edit notes using this interface.**

converging on the final themes, which we describe in the next section.

### 3.3 Participants

The 15 pairs of participants (30 in total) included 15 females, 14 males, and 1 non-binary, ages 23 to 73 years old (mean = 36). Participants were required to know how to read traditional music notes to ensure they would understand the interfaces. Participants were also required to have a reliable internet connection to participate in the remote study.

Because the social composition process and interaction with AI may vary depending on musical skill level, we sought to recruit and observe a broad cross-section of skill levels. To do so, we included an equal number of novices (no experience composing or improvising), dabblers (e.g., played in jazz band in high school; tried making tunes in Garageband, played improvisation occasionally), and serious composers/improvisers (e.g., formal music composition/theory training; music major in college; wrote multiple songs for *a capella* groups; regular jam sessions). In total, the study included 5 pairs of novices, 5 pairs of dabblers, and 5 pairs of serious improvisers/composers. We also balanced the gender composition, resulting in 5 sets of female-female, male-male, and mixed-gendered

pairs. We paired participants who did not know each other to minimize the potential influence that prior social relationships could have on the collaborative process. (The strength of a social relationship is known to influence teamwork [45]. For example, people with stronger ties (e.g., friendship) are better able to co-construct knowledge and coordinate tasks, are more adaptive to each other's behavior, and are less reliant on explicit communication [63].) Participants were recruited through mailing lists at our institution and came from a variety of professional backgrounds (e.g., technical writer, administrator, manager and designer). Each received a \$40 gift credit for their time.

## 4 RESULTS

In this section, we describe how the 15 pairs of participants collaborated in composition without and with AI, and what kinds of roles AI played in their co-creative processes.

### 4.1 Co-creation Processes

Before putting any music down on the page, participants usually started composing by jointly establishing a creative strategy or goal (e.g., “What kind of overall feeling do you want to have for this one [picture]” (P3) or by agreeing on an implicit “social contract” that

would make it easier to be consistent with each other later on (e.g. e.g., “Any specific key that you want?” (P13)). In many cases, the pairs then divided their roles in composition (e.g., splitting by voice (P4, P7, and P15), or by bar (P9)), and took turns from time to time (5 pairs). In other cases, one took the lead by suggesting the notes, and the other provided feedback while constantly asking for confirmation from each other (6 pairs). Throughout the collaboration, participants repeatedly listened to their music together, and mutually decided when it was complete (15 pairs). A few participants wanted a clearer visual delineation of who produced which parts to see their own contributions (P1) or where to edit later (P9), as music is subjective and overwriting the other’s contributions often felt rude (P7).

Though participants’ co-creation process with AI present was largely similar to their process without AI, they had additional decision-making points on when, where, and how to use AI in their collaboration: “If it’s just the two of us, we kind of decide what parts we do each, but now we have to decide on top of that.” (P15) These additional decisions included: 1) deciding where and when to request that the AI generate content (voices and time segments), 2) deciding what quality of music they wanted in that segment by adjusting the sliders, 3) evaluating and selecting which of the *N* AI-generated options sounded best, and 4) deciding which partner would be the main driver of these decisions above, each time they occurred. Whereas participants sometimes yielded the first two types of decisions (where to generate, which slider settings) to each other (e.g. “Sure, go for it” (P1)), they almost always jointly evaluated the *N* AI-generated options together (e.g. “I like the first one, what do you think?” (P3)).

In addition to these decision points, one can also consider how the AI was used as a composition tool in the writing process. We observed three primary ways that AI was employed. First, some participants used AI for open-ended inquiries: “Let’s see what the AI says” (P5). Second, some participants had goal-directed requests for AI during the composing process (e.g., “Let’s do it even more surprising and super minor” (P15)). In another example, when P13 needed to fix specific parts, they requested the AI to generate content for those spots to get some ideas on how to proceed. Third, AI was used to flesh things out once they had completed a basic skeleton of their parts. For example, participants (P6) took turns creating the main structure of the music by composing the soprano and bass voices, then used AI to fill in the rest of the voices (alto and tenor voices). P8 described a similar process: “We draw the creative outline and the AI helps fill it in.” In these cases, where AI was used to flesh out the music by ‘gluing’ together various human-created pieces, participants usually had the AI generate musical components that would have been more tedious (or required more music theory knowledge) for humans to complete (e.g. harmonizing given a melody, or filling in a transition between the beginning and end). When introducing AI in co-composition processes, pairs’ division of labor became somewhat blurred, which we elaborate on in more detail in the next section.

## 4.2 Roles of AI in Social Collaboration

In this section, we describe ways in which the AI appeared to influence social dynamics during co-composition, taking into account

users’ reported perceptions of AI as well as their observed behaviors during the co-composition process. Here, we frame our findings around the five primary themes that emerged: 1) AI as establishing common ground, 2) AI as a psychological safety net, 3) AI as a force for progress, 4) AI as a social lubricant, and 5) AI as a force that alters the creative and collaborative roles of humans.

**4.2.1 AI as common ground.** At the start of the co-composition process, pairs often sought to establish common ground as a way of familiarizing themselves with each other and to build a shared, mutual understanding of each other (5 pairs). Finding common ground is a technique used by people to facilitate interpersonal relationships, and is often a critical ingredient for establishing mutual respect and trust in interpersonal collaborations and collective actions [13]. When co-composing, participants asked a number of questions to each other (e.g., preference, feedback or confirmation on suggestions) to establish common ground:

P11a: Do you have any preference? like a chord progression or...

P11b: Well, I was thinking about maybe picking a key first. How does that sound?

P11a: Sure.

P11b: What are our key options? I feel like...it’s probably a major picture. Does that seem right?

P11a: That works for me!

...

P11a: How does E flat sound?

P11b: E flat sounds great.

P11a: Great.

In the presence of AI, however, pairs often found common ground through discussing the AI-generated music (7 pairs). In their reflections, pairs found it helpful to be able to anchor their conversation around the AI-generated music, rather than needing to rely on conversing solely about self-generated ideas. As one member of P10 described, “Our collaboration, without AI, required coming up with some kind of basis like a core progression, and then building from there. Whereas with AI, we use that [the AI generated options] as our basis.” As indicated in this quote, the AI suggestions can provide an immediate platform for discussion. Discussions about the AI-generated music also enabled participants to get to know each other’s preferences. For example, P9a was able to figure out P9b’s musical taste while discussing the AI generated options.

P9a : We listen to this new one ((*playing the first option*))<sup>1</sup>

P9b : ((*Positively surprised while listening to the song*))  
Woooo!

P9a : That was nice

P9b : Yeah, that was nice

P9a : Let’s listen to this one ((*playing another option*))

P9b : That is nice too, but I think I like the previous one a little better

P9a : This one ((*pointing to the first option*)) ... Let’s listen to a final ((*playing the second option again*)). I’m guessing more movement is better, but okay...

<sup>1</sup>Italic text in double parentheses denotes annotation of non-verbal activity, following Jefferson (1984) “Transcription Notation,” in J. Atkinson and J. Heritage (eds), *Structures of Social Interaction*, New York: Cambridge University Press

P9b : ((*Listening to the second option again*)) Yeah, you're right. It's a little too slow or it's a little...

P9a : All right, that was my inclination or my assumption about why you like this one more

When initiating collaboration without AI present, participants sometimes encountered a *cold-start problem*: “*Jumping into the pool and proposing something...that was not easy*” because “*There's like a first mover thing where whoever proposes the first thing gets to dictate the rest of it. I didn't want to dictate anything.*” (P15) Some pairs found it easier to jump into a collaborative processes when AI was present (7 pairs), mainly because it allowed them to ground themselves on existing content without first having to coordinate who would be creating which parts: “*The AI helps a lot for working together, otherwise it would be really tricky to coordinate who's working on what line and what part, whereas here we're able to listen to think together and then decide together what sounds right and then make adjustments.*” (P14)

In sum, AI's generated output provided a seed for discussion that participants found helpful for finding common ground. It also provided initial momentum for getting the composition started.

**4.2.2 AI as a psychological safety net.** Beyond common grounding, AI also provided a psychological safety net during co-composition. A psychological safety net refers to a belief that one is not at risk of embarrassment or rejection in particular setting or role [19], and it is known to increase information sharing [72], group learning [19] or group performance [4]. While composing with each other, both novices and serious composers were sometimes wary of feeling embarrassed or judged as they created notes. For example, even a participant with extensive musical training stated, “*Oh god, what is this going to sound like?*”, adding that they would “*be slightly embarrassed after my first attempt*” (P10). Participants felt that this self-consciousness may stem from the pressure of having a human collaborator watching them; they indicated that, when composing alone, they'd usually feel more free to explore and experiment: “*If I were alone, there would be less of a feeling of being tentative or worried that something is going to sound weird. You can be more creative and messy without the social pressure. But when you collaborate with another human...you interpret pictures differently, and one person might see hope where someone else sees fear.*” (P2)

In contrast, some participants indicated they were able to feel more relaxed and playful when AI was present (5 pairs). Comparing their experiences when AI was present versus it not being present, participants stated that there was less “*responsibility to make something that sounded cool... there's less weight on you*” (P10) when AI was present. This sentiment also tied into how participants felt about their ability to be playful, both with each other and with the music. For example, P1 said, “*there was kind of a feeling of playfulness. Encountering what the AI did, it gave me some thoughts about, 'oh, we could also do this.' It's like having a third person come in and play a little riff and go. 'Oh, that is nice. Let's try this'... so it kind of opened up more ideas, so that was great.*” Similarly, P13 described AI as a collaborator, introducing fresh ideas: “*We asked for it like, 'Hey, oddballs! Line yourselves up and give us like the most odd surprising thing,' and we got exactly what we asked for...I would never ask a musician who walked in the room and just be like off the wall, and go nuts, right?*”

A few participant felt they were more willing to try new things, since “*knowing we had AI felt like we could always smooth it out later.*” (P4) This willingness to explore new territory was most frequently observed among pairs who had less musical background. Novices generally believed that the AI's training surpassed their own music theory abilities, and that it would—at the very least—keep their music from sounding incoherent: “*I just have this image of the AI being...very intelligent and far more experienced than [me]. You'd assume the AI has been written and processed to understand proper music theory and things like that. So they're not going to spit out something super crazy.*” (P5) Interestingly, the unique idiosyncracies of AI may have also helped introduce a more playful atmosphere. One pair felt that working with each other initially felt like “*treading a lot more...on the careful side,*” whereas with AI, “*we laughed a lot more...because of this like unexpectedness...this random person came and just joined us.*” (P14)

Overall, AI may help provide psychological safety by offloading the burden of generating acceptable content, acting as a safety net that re-mediate or smooths over human-created rough edges, and introducing an air of playfulness through its surprises and idiosyncracies.

**4.2.3 AI as a force for progress.** AI also acted as a force for progress, helping participants' co-composing process move forward “*easier*” (P12) and “*quicker*” (P5). Specifically, some felt the AI was useful for helping them get started quickly and build momentum (8 pairs). For example, participants described it as “*a way to accelerate through writer's block*” (P8), or “*it jump-started us...like having a bass line to work on*” (P11). Without AI, users felt that the process of generating and writing down ideas from scratch sometimes slowed down iteration: “*Manually going through that ideation phase, the time to manually pick out how to express ...takes so long to get to a point where you can even react to it.*” (P8). Users largely agreed that it was far easier to be given some initial content to use that served as a starting point for ideas, rather than having to generate new content from scratch.

While some felt the AI helped jump-start their initial progress, others saw the AI as helping flesh out or continue half-completed work (7 pairs). For example, the AI offloaded work that would normally be difficult and time-consuming, such as harmonizing chords given a melody, or filling in the ‘glue’ between human-created components (e.g., filling in the middle voices, or bridging between two sections across time). P10 said, “*We got a lot further than I would have expected in a shorter amount of time..I listen to music a lot, but would never be able to do something like this anywhere near as quickly, so the AI serves a huge function.*” Others felt that the AI served as a “*helpful assistant*” (P8) because it could auto-complete content based on a rough seed provided by the user: “*I felt like we were able to go faster because like we could put in the basis of an idea and then the AI could fill in the rest.*” (P13) Because users could pick empty regions of any size and request the AI to automatically fill them in, users felt a rapid sense of progress, reinforced through seeing notes fill more and more empty space on the page. As an unintended side effect of this, some pairs who were running out of time asked the AI to auto-complete the rest of the missing sections.

Aside from producing content, AI also helped participants move forward in their creative decision making. It did so by both narrowing a large space of decisions to a few concrete decision points,

and by *broadening* the scope of ideas being produced when users were stuck. Under non-AI circumstances, users were sometimes faced with a broad decision space when deciding what to do next. For example, during the collaboration, P3 frequently had silent moments where both were stalled because they could not decide on next steps:

P3a: So should we do the other lines?  
 P3b: Yeah.  
 P3a: Ummm...Which one [voice] do you want to do?  
 P3b: ((*Quietly pondering*))  
 P3a: Tenor?  
 P3b: ((*Silence*))  
 P3a: Do you want to do bottoms up or do you want to...  
 P3b: Maybe we can do like soprano?  
 P3a: ((*Reluctantly*)) Ummm...Okay  
 P3a: ((*Without confidence*)) mmm...Okay

Conversely, with the AI present, users primarily needed to decide *where* and *how much* content to request that the AI generate; this reduced the scope of decisions to be made: “*With the AI, you set the parameters/constraints, but AI does everything else.*” (P2). Paradoxically, AI’s production of multiple alternatives also reduced the decision space. Whereas judging the merits of a single musical snippet may feel open-ended without another point of comparison, picking one’s favorite between three AI-produced musical phrases made the decision inherently more bounded: “*It gave us the options to decide to agree or disagree on them.*” (P9)

Some felt that, aside from narrowing the scope of decisions, AI also helped them make progress by expanding the space of ideas. For example, one participant felt the generated notes “*gave us something that we wouldn’t have thought of on our own,*” and that “*it would be super useful for anyone who runs out of inspiration, they would get unlimited ideas.*” (P2) Another described the AI as substantially increasing the number of collaborators and volume of ideas generated: “*it was like having a third person come in and play a little riff and go, ‘oh, that is nice, let’s try this’...So it was instead of two collaborators, we have five for example...three other sequences presented to us.*” (P1) During composition, some also thought of new creative ideas (e.g., P13: happy/sad; P3: folksy/mid-western) while considering how to use the AI-steering sliders (4 pairs): “*I am curious how it could be some type of more emotional aspect to it...even if you kept this like conventional or surprising [slider] but it [emotion] is something a little different.*” (P5)

While most participants felt the AI was helpful for accelerating progress (14 pairs), one pair remarked that the AI added an additional overhead of decision making that may deter progress: “*There’s an additional thing that we have to decide on. So if it’s just the two of us, we kind of decide what parts do we do each one, but now we have to decide on top of that...there are just more things to decide and coordinate.*” (P15) In this case, the participants were experienced composers, and found it more work to steer the AI towards their creative goals than to create the music themselves: “*AI didn’t really enhance from that [the partner’s initial] idea...[it] diverged out of the team so to speak. so it became disjointed...it didn’t capture our initial idea at all.*”

Overall, AI helped the co-composition process move forward more quickly, providing initial content as a starting point, narrowing a large decision space to a few concrete decision parameters, generating alternative ideas when users were stuck, and filling in large chunks of unfinished work, all of which enabled participants to speed up their progress.

**4.2.4 AI as social lubricant.** Although participants rarely had explicit disagreements, some indicated in post-study interviews that they felt hesitant to freely comment on their partner’s work when opinions diverged: “*I wouldn’t say that to [partner] like if he did something like that...because like...that’s rude.*” (P9) Some explained that it seemed especially inappropriate to freely express judgments due to the intrinsic subjectivity of creativity: “*There is no right or wrong thing [in music] as long as it aligns with the basic chords.*” (P7)

In cases when opinions differed, participants found AI helpful as it provided various options that they could branch from and further explore beyond the points of disagreement (7 pairs). For example, P3 said, “*it makes it easier especially...for if people don’t have same ideas,*” because it provides “*choices right through points of conflict and it also provides options like, you know...we can discuss all those options and you can make changes on those options as well.*” (P3) Others also appreciated that AI’s Multiple Alternatives feature helped scope these implicit debates down to a few concrete options: “*It gave us the options to decide to agree or disagree on them.*” (P9)

Importantly, participants felt that AI made it easier to express differences in opinion, because its feelings as a non-human could not be hurt (7 pairs). For example, participants said they were “*not worried about hurting the computer’s feelings*” (P11), and “*it’s not like the AI is going to get upset that we didn’t like either choice*” (P5). In a sense, the presence of AI as a third collaborator shifted one’s judgment from being directed at one another, to being directed toward a non-human third party: “*It takes away that tension of ‘this is how I see the world and I want to make you see it like I do’, it’s more ‘let’s be playful and see if we can converge on something that a third party suggests.’*” (P10) This was in large part because participants did not feel the need to maintain the same level of politeness towards AI-generated music (“*it eliminated any possibility of offending each other because we weren’t the ones creating it.*” (P9)) or AI-generated decisions (“*there’s that other added element that you avoid dissonance because it’s not you or the other person making the decision.*” (P7))

The presence of AI—especially its surprises or mistakes—also led some participants to feel closer to their partners in their shared human-ness: “*It [AI] puts you and the other person on the same team. You’re both humans, so you can kinda laugh at this thing, or think ‘Wow, this thing is brilliant, what it created.’ Team human!*” (P10) When pairs mutually felt dissatisfied with the options that AI generated, AI felt like “*the bad team member that doesn’t show up for practice*” (P15), making the human-human partnership feel more positive in comparison.

In sum, AI helped uncover and mitigate latent friction, both by offering a path for moving forward, and by shifting social dynamics away from human-human differences towards human-AI differences.

**4.2.5 AI altering the creative and collaborative roles of humans.** The presence of AI also affected participants’ creative roles in the composition process, as well as their collaborative roles with each other.



While most participants noted that AI was like a “third party collaborator” (10 pairs), many felt that their own roles shifted from being composers to “producers” (P9), “advisors” (P7) or “museum curators” (P6): rather than generating ideas from scratch, participants felt that they were “patch-working” or organizing little pieces together (3 pairs). As one user put it, “When AI was present, our roles were more of choosing the ones that sound best and not necessarily building on top of it or creating something of ourselves.” (P15) Another described their own role as one of giving feedback to the AI composer, rather than directly composing themselves: “I felt like it was the composer and we were the listeners giving feedback and choosing. I’m telling the composer, ‘...here’s our feedback for how you can do better’ and let it compose something else.” (P6). Some felt that this shift in role could make composition skill-level less pertinent: “It’s much less about our personal skill in composition, and rather our ear and what our ear thinks is better.” (P9)

Alongside this shift in creative roles, participants also felt a shift in the depth of their creative engagement with each other. Some felt they were more creatively engaged with each other in sessions where AI was absent: “When we couldn’t rely on the AI, then we were definitely more involved [with each other].” (P2) For example, in one session without AI, a pair took turns building on top of each other’s notes to create a sense of tension in the music (P13). P13b carefully chose a note in the alto voice that would complement what P13a had created in the soprano voice: “Let me put something without breaking that [the notes on soprano].” P13a then provided positive feedback, and subsequently added notes in the tenor voice to build on top of the song they had composed together: “((after listening to the soprano and alto voices together)) Yeah, okay. ((with a big smile)) That’s really cool. I like that! Okay, then I’ll do tenor.” The pair was deeply engaged with each other during this co-creative process.

Conversely, with the AI generating content, there was less of a need to think deeply about what one’s partner had composed and consider how best to build on top of it: “I think when AI is present, we get a kind of feeling like ‘all right. yeah, you do the work, I will let you know what we like and don’t like.’ Versus when AI is not present, ‘Oh we both got to put in a little bit more work.’” (P5) One participant summarized this as a trade-off between the ease and depth of collaborative co-creation: “The AI made the collaboration easier, but we collaborated deeper without it.” (P12)

Finally, intrinsic constraints of the AI algorithm itself may have limited users’ creative scope and freedom. Because the machine learning algorithm was trained on Bach, and generated music in classical genres, participants realized they needed to write music that would maximize compatibility with the algorithm. For example, a participant who picked up on the AI generation’s similarity to Bach stated: “It always has a Bach flair on the melodies it suggests. So, if I wanted to go a totally different stylistic way, I might not be able to count on it to produce the results I want” (P6). Conversely, in the non-AI condition: “We knew this time we weren’t using AI, so I actually felt less restricted. This time, we could get a funkier color and break more rules that AI wouldn’t like.” (P14) Indeed, multiple participants asked if the AI controls could be broadened in scope to enable a greater diversity of musical styles and genres: “I would like the AI options to be labeled with a different style. It didn’t have that many customization built-ins as you’re automating it. I would like the sliders to all have different options, like a Gregorian chant”

(P4) and “If there were more flavors of the AI besides Bach, it would be more interesting.” (P6)

## 5 DISCUSSION

Our findings bring to light the meaningful roles that AI may play in mediating social dynamics during collaborative co-creation. Given these findings, we now discuss the broader implications of this research.

### 5.1 AI as Social Glue

We observed a wide range of nuanced social dynamics at play during co-composition: participants needed to 1) express opinions about their partners’ work while maintaining an air of cordiality; 2) feel a sense of psychological safety while taking creative risks; and 3) engage actively with each other’s ideas while sustaining forward progress at a reasonable pace. Though the participants in our study did not previously know each other, and the task was relatively low-stakes, these implicit tensions are likely to arise in a wide range of collaborative settings, at different levels of extensiveness: generating novel content with a human partner will often involve striking a delicate balance between respecting and building on another person’s ideas, without losing one’s own creative voice.

Beyond functioning as yet another collaborator, AI appeared to play a meaningful role in mediating these nuanced social dynamics. This raises the question of how AI could be designed more intentionally to facilitate social collaboration. For example, AI could generate solutions that are “middle-grounds” or compromises between the ideas of multiple collaborators. Or, it could produce something explicitly divergent from both, to help collaborators break out of the situation and more broadly consider their options. Interestingly, even though the semantic sliders were ostensibly designed to enable *individual* control of AI [54], a side effect was that they also occasionally helped mitigate social friction, by allowing participants to change the slider settings as a way of quickly demonstrating an idea to their partner in a low-cost way. Likewise, the “multiple alternatives” feature of Cococo enabled collaborators to “hedge” with each other by debating a set of available options in a temporary sandbox, rather than editing anything in the final composition head-on. Along these lines, AI could further help a person propose an idea without stepping on another’s toes, by generating multiple variants of that idea, thereby giving one’s collaborator a greater sense of choice. As such, AI could act as a “social glue,” allowing collaborators to build fluidly on each other’s ideas, maintaining cordiality without losing momentum.

### 5.2 Potential AI Shortcomings: Easing vs. Deepening Collaboration

While we observed many benefits regarding AI functioning as social glue, our study also uncovered several potential shortcomings of AI. First, participants felt that AI could limit their *depth* of collaboration, despite it enhancing their *ease* of collaborating together. For example, rather than evaluating and directly building on each other’s ideas, participants naturally shifted towards jointly evaluating the AI’s outputs. Some felt they were acting more as co-producers and less as co-composers (see Section 4.2.5). Second, the presence of



AI also gave participants a way to jointly judge a third party’s creative content, without fear of offending it. While this had a positive effect on the human collaboration (Section 4.2.4), strengthening the bond between human participants and mitigating interpersonal friction, it is also possible that, if taken to the extreme, AI’s role as an ostracized third party may contribute to an “us vs. them” mentality. Finally, though most pairs felt that AI helped them make progress faster, a few who were already experts in music felt it also introduced some additional decision making overhead, especially when AI output was difficult to control (Section 4.2.3). Designers of collaborative systems may want to carefully consider these possibilities in multi-agent systems where AI plays a role as a third party.

Given the observed trade-offs between AI easing vs. deepening collaboration, future work could investigate how the *depth* of collaboration and interpersonal engagement can be preserved, or even enriched, when AI is present. Or, future designs of deep generative AI may consider explicitly allowing collaborators to control how deep they would like the collaboration to be. For example, to deepen engagement with other (human) collaborators, an AI could go beyond generating coherent raw material (e.g., musical notes, raw pixels, or brush strokes), to generating higher-level strategies specifically designed to support a user in complementing, contrasting, or furthering their partner’s work (e.g. creating variations on a colleague’s motif). Conversely, if collaborators merely seek to quickly complete a task together, or are pressed for time, the AI could take on more of the creative work, allowing users to primarily act as a panel of judges or producers. Future AI could enhance both the efficiency and depth of collaboration by automatically generating not only low-level notes and content, but also higher-level templates for how to do the social coordination work itself (e.g., learning workflows for dividing up creative responsibility).

### 5.3 Human-ness of Deep Generative AI

Through its surprises and idiosyncrasies, we observed that AI provided a psychological safety net while introducing a level of playfulness during the co-creative process. This allowed users to feel more comfortable taking creative risks with a partner watching. Psychological safety may be particularly valuable in collaborations between unfamiliar partners, at the start of collaboration, or for novices with less artistic experience [19, 70]. Future systems and studies could investigate how to leverage the playful properties of AI as a psychological safety net, particularly in “icebreaker” phases of collaboration where collaborators have yet to build up an adequate level of mutual trust, or when creative risk-taking is detected to be low.

Though not by design, knowing that the AI was non-human also played a central role in mitigating latent tension. Whereas participants experienced some reservation about openly commenting on each other’s work, they felt much more comfortable judging and debating the AI’s outputs. This lack of reservation both enabled them to move forward and created the space for common grounding and discovering each other’s musical tastes.

These observations raise important implications around how to define objective functions for deep generative AI, particularly given that machine learning models are currently optimized to imitate

human-created training data. On the one hand, an AI’s human-like generative capabilities may provide psychological safety (by allowing users to experiment and move quickly without fear of failure) and create a force for progress (by adeptly filling in missing parts). On the other hand, knowledge of the AI’s non-human-ness may encourage playfulness, enabling users to openly express opinions, and facilitate social bonding through the establishment of a “common foil.” Hence, disentangling the multiple dimensions of AI’s “human-ness” (or lack thereof), and identifying which aspects enrich vs. impede the quality of collaboration, deserves further investigation.

### 5.4 Scope of Collaboration as a function of AI Training Data

While Cococo gave users compositional “superpowers,” its Bach-like nature also implicitly limited the creative scope of their compositions, and likewise, the scope of their collaboration. Whereas musicians may, under normal circumstances, lean on their collaborators’ stylistic or instrumental expertise, it is less likely that, when composing with AI, one would be able to fully leverage a collaborator’s skill set in an area that is sufficiently different from the scope of the AI’s training set. This difficulty of “breaking out” of an AI’s training set may suggest that the value one reaps from a human collaboration may depend somewhat on an algorithm’s inherent scope—whereas it may be beneficial to partner with someone whose skill set diverges slightly from that of the AI, a skill set that departs dramatically from the AI may not be fully expressible.

### 5.5 Limitations/Future Work

Some limitations should be noted. As discussed in section 5.1, some level of creative tension is likely to arise across a wide range of tasks. However, the extensiveness of interpersonal tension and one’s desire for creative agency may vary depending on context. Though the task in our study was relatively low-stakes, we acknowledge that participants may be more or less forthcoming depending on how high-stakes the task is, or to what extent they are already familiar with each other when conflicts arise. While we found that AI served as a mitigator of interpersonal friction, the nuances of those dynamics could be affected by the specific context and nature of the task. Thus, future research should further examine the role of AI across a greater variety of human-human collaboration contexts (e.g., different tasks, social relationships among human collaborators).

## 6 CONCLUSION

Through a study examining how pairs of individuals compose music with and without AI, this research demonstrates ways AI can play meaningful roles in mediating social dynamics during collaborative composition. We also show how people adapt and utilize AI as an additional collaborator in their co-creation practices. Noting the potential of AI as social glue, future systems could be explicitly designed to help *deepen* collaborations (as opposed to simply automating work). There is also an opportunity to experiment with how “human” an AI feels, given the benefits and drawbacks we observed for human-human collaboration (e.g., allowing partners to openly express opinions to each other). Taken together, this work contributes to an understanding of how to design AI to assist social

collaboration, leveraging AI to help people be more creative and successful in their human-human collaborations.

## ACKNOWLEDGMENTS

The authors thank Anna Huang, Andy Coenen, Ian Simon, and Dan Russell for their support and feedback on this research.

## REFERENCES

- [1] Nuha Albadi, Maram Kurdi, and Shivakant Mishra. 2019. Hateful People or Hateful Bots? Detection and Characterization of Bots Spreading Religious Hatred in Arabic Social Media. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–25.
- [2] Cecilia R. Aragon and Alison Williams. 2011. Collaborative Creativity: A Complex Systems Model with Distributed Affect. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Vancouver, BC, Canada) (CHI '11). Association for Computing Machinery, New York, NY, USA, 1875–1884. <https://doi.org/10.1145/1978942.1979214>
- [3] Zahra Ashktorab, Mohit Jain, Q. Vera Liao, and Justin D. Weisz. 2019. Resilient Chatbots: Repair Strategy Preferences for Conversational Breakdowns. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3290605.3300484>
- [4] Markus Baer and Michael Frese. 2003. Innovation is not enough: Climates for initiative and psychological safety, process innovations, and firm performance. *Journal of Organizational Behavior: The International Journal of Industrial, Occupational and Organizational Psychology and Behavior* 24, 1 (2003), 45–68.
- [5] Jared S. Bauer, Aubury L. Jellenek, and Julie A. Kientz. 2018. Reflektor: An Exploration of Collaborative Music Playlist Creation for Social Context. In *Proceedings of the 2018 ACM Conference on Supporting Groupwork* (Sanibel Island, Florida, USA) (GROUP '18). Association for Computing Machinery, New York, NY, USA, 27–38. <https://doi.org/10.1145/3148330.3148331>
- [6] Erin Beneteau, Ashley Boone, Yuxing Wu, Julie A Kientz, Jason Yip, and Alexis Hiniker. 2020. Parenting with Alexa: Exploring the Introduction of Smart Speakers on Family Dynamics. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [7] Steve Benford, Peter Tolmie, Ahmed Y. Ahmed, Andy Crabtree, and Tom Rodden. 2012. Supporting Traditional Music-Making: Designing for Situated Discretion. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work* (Seattle, Washington, USA) (CSCW '12). Association for Computing Machinery, New York, NY, USA, 127–136. <https://doi.org/10.1145/2145204.2145227>
- [8] Tina Blaine and Tim Perkin. 2000. The Jam-O-Drum Interactive Music System: A Study in Interaction Design. In *Proceedings of the 3rd Conference on Designing Interactive Systems: Processes, Practices, Methods, and Techniques* (New York City, New York, USA) (DIS '00). Association for Computing Machinery, New York, NY, USA, 165–173. <https://doi.org/10.1145/347642.347705>
- [9] Virginia Braun and Victoria Clarke. 2012. Thematic Analysis. In *APA handbook of research methods in psychology*, D. L. Long A. T. Panter D. Rindskopf & K. J. Sher H. Cooper, P. M. Camic (Ed.). American Psychological Association, 51–74.
- [10] Alex Calderwood, Katy Ilonka Gero, and Lydia B. Chilton. 2020. How Novellists use Generative Language Models: An Exploratory User Study. In *IUI Workshops*.
- [11] Lydia B Chilton, Savvas Petridis, and Maneesh Agrawala. 2019. VisiBlends: A flexible workflow for visual blends. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [12] Elizabeth Clark, Anne Spencer Ross, Chenhao Tan, Yangfeng Ji, and Noah A Smith. 2018. Creative writing with a machine in the loop: Case studies on slogans and stories. In *23rd International Conference on Intelligent User Interfaces*. 329–340.
- [13] Herbert H Clark and SE Brennan. 1991. Grounding in communication. LB Resnick, RM Levine, SD Teasley, eds. *Perspectives on Socially Shared Cognition* (1991).
- [14] Wikipedia contributors. 2020. *Dixit (card game)*. Retrieved September 8, 2020 from [https://en.wikipedia.org/w/index.php?title=Dixit\\_\(card\\_game\)&oldid=908027531](https://en.wikipedia.org/w/index.php?title=Dixit_(card_game)&oldid=908027531)
- [15] Wikipedia contributors. 2020. *Emily Howell*. Retrieved September 17, 2020 from [https://en.wikipedia.org/wiki/Emily\\_Howell](https://en.wikipedia.org/wiki/Emily_Howell)
- [16] David Cope. 2005. *Computer models of musical creativity*. MIT Press Cambridge.
- [17] D. M. Cunningham, S. J. & Nichols. 2009. Exploring social music behaviour: An investigation of music selection at parties. In *Proceeding of 10th International Society for Music Information Retrieval Conference*. 747–752.
- [18] Nicholas Davis, Cih-Pin Hsiao, Kunwar Yashraj Singh, Lisa Li, and Brian Magerko. 2016. Empirically studying participatory sense-making in abstract drawing with a co-creative cognitive agent. In *Proceedings of the 21st International Conference on Intelligent User Interfaces*. 196–207.
- [19] Amy Edmondson. 1999. Psychological safety and learning behavior in work teams. *Administrative science quarterly* 44, 2 (1999), 350–383.
- [20] Judith E Fan, Monica Dinculescu, and David Ha. 2019. collabdraw: an environment for collaborative sketching with an artificial agent. In *Proceedings of the 2019 on Creativity and Cognition*. 556–561.
- [21] Robin Fencott and Nick Bryan-Kinns. 2010. Hey Man, You're Invading my Personal Space! Privacy and Awareness in Collaborative Music.. In *NIME*. Citeseer, 198–203.
- [22] Rebecca Anne Fiebrink. 2011. *Real-time human interaction with supervised learning algorithms for music composition and performance*. Citeseer.
- [23] Gerhard Fischer and Frank Shipman. 2013. Collaborative design rationale and social creativity in cultures of participation. In *Creativity and Rationale*. Springer, 423–447.
- [24] Jonas Frich, Michael Mose Biskjaer, and Peter Dalsgaard. 2018. Twenty Years of Creativity Research in Human-Computer Interaction: Current State and Future Directions. In *Proceedings of the 2018 Designing Interactive Systems Conference*. 1235–1257.
- [25] Marco Furini, Silvia Mirri, Manuela Montangero, and Catia Prandi. 2020. On the Usage of Smart Speakers During the Covid-19 Coronavirus Lockdown. In *Proceedings of the 6th EAI International Conference on Smart Objects and Technologies for Social Good*. 187–192.
- [26] Katy Ilonka Gero and Lydia B Chilton. 2019. Metaphoria: An Algorithmic Companion for Metaphor Creation. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [27] Wooi Boon Goh, Lijia Yang, and Kian Leong Ong. 2016. MuSeeCol: A See-through Multi-Touch Surface for Face-to-Face Musical Collaboration. In *Proceedings of the 13th International Conference on Advances in Computer Entertainment Technology* (Osaka, Japan) (ACE '16). Association for Computing Machinery, New York, NY, USA, Article 15, 6 pages. <https://doi.org/10.1145/3001773.3001788>
- [28] Matthew Guzdial, Nicholas Liao, Jonathan Chen, Shao-Yu Chen, Shukan Shah, Vishwa Shah, Joshua Reno, Gillian Smith, and Mark O Riedl. 2019. Friend, collaborator, student, manager: How design of an ai-driven game level editor affects creators. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [29] Gaëtan Hadjeres, François Pachet, and Frank Nielsen. 2017. Deepbach: a steerable model for bach chorales generation. In *International Conference on Machine Learning*. PMLR, 1362–1371.
- [30] Jeffrey T Hancock, Mor Naaman, and Karen Levy. 2020. AI-Mediated Communication: Definition, Research Agenda, and Ethical Considerations. *Journal of Computer-Mediated Communication* 25, 1 (2020), 89–100.
- [31] Michael T Hopkins. 2015. Collaborative composing in high school string chamber music ensembles. *Journal of Research in Music Education* 62, 4 (2015), 405–424.
- [32] Cheng-Zhi Anna Huang, Tim Cooijmans, Adam Roberts, Aaron Courville, and Douglas Eck. 2019. Counterpoint by convolution. In *Proceedings of the 18th International Society for Music Information Retrieval Conference, ISMIR 2017*.
- [33] Cheng-Zhi Anna Huang, David Duvenaud, and Krzysztof Z Gajos. 2016. Chordriple: Recommending chords to help novice composers go beyond the ordinary. In *Proceedings of the 21st International Conference on Intelligent User Interfaces*. 241–250.
- [34] Cheng-Zhi Anna Huang, Curtis Hawthorne, Adam Roberts, Monica Dinculescu, James Wexler, Leon Hong, and Jacob Howcroft. 2019. The bach doodle: Approachable music composition with machine learning at scale. *arXiv preprint arXiv:1907.06637* (2019).
- [35] Cheng-Zhi Anna Huang, Ashish Vaswani, Jakob Uszkoreit, Ian Simon, Curtis Hawthorne, Noam Shazeer, Andrew M Dai, Matthew D Hoffman, Monica Dinculescu, and Douglas Eck. 2018. Music transformer: Generating music with long-term structure. In *International Conference on Learning Representations*.
- [36] Koops Hendrik Vincent Huang, Cheng-Zhi Anna, Ed Newton-Rex, Monica Dinculescu, and Carrie J. Cai. 2020. AI Song Contest: Human-AI Co-creation in Songwriting. (2020).
- [37] Mikhail Jacob and Brian Magerko. 2015. Interaction-based Authoring for Scalable Co-creative Agents.. In *ICCC*. 236–243.
- [38] Mikhail Jacob and Brian Magerko. 2018. Creative arcs in improvised human-computer embodied performances. In *Proceedings of the 13th International Conference on the Foundations of Digital Games*. 1–6.
- [39] Irving L Janis. 2008. Groupthink. *IEEE Engineering Management Review* 36, 1 (2008), 36.
- [40] Sergi Jordà, Günter Geiger, Marcos Alonso, and Martin Kaltenbrunner. 2007. The ReacTable: Exploring the Synergy between Live Music Performance and Tabletop Tangible Interfaces. In *Proceedings of the 1st International Conference on Tangible and Embedded Interaction* (Baton Rouge, Louisiana) (TEI '07). Association for Computing Machinery, New York, NY, USA, 139–146. <https://doi.org/10.1145/1226969.1226998>
- [41] Malte F. Jung, Nikolas Martelaro, and Pamela J. Hinds. 2015. Using Robots to Moderate Team Conflict: The Case of Repairing Violations. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction* (Portland, Oregon, USA) (HRI '15). Association for Computing Machinery, New York, NY, USA, 229–236. <https://doi.org/10.1145/2696454.2696460>
- [42] Pegah Karimi, Mary Lou Maher, Nicholas Davis, and Kazjon Grace. 2019. Deep Learning in a Computational Model for Conceptual Shifts in a Co-Creative Design System. *arXiv preprint arXiv:1906.10188* (2019).

- [43] Pegah Karimi, Jeba Rezwana, Safat Siddiqui, Mary Lou Maher, and Nasrin Dehbozorgi. 2020. Creative sketching partner: an analysis of human-AI co-creativity. In *Proceedings of the 25th International Conference on Intelligent User Interfaces*. 221–230.
- [44] Charles Kiene and Benjamin Mako Hill. 2020. Who Uses Bots? A Statistical Analysis of Bot Usage in Moderation Teams. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–8.
- [45] Michael Klug and James P Bagrow. 2016. Understanding the group dynamics and success of teams. *Royal Society open science* 3, 4 (2016), 160007.
- [46] Janin Koch, Andrés Lucero, Lena Hegemann, and Antti Oulasvirta. 2019. May AI? Design ideation with cooperative contextual bandits. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [47] Max Kreminski, Melanie Dickinson, Michael Mateas, and Noah Wardrip-Fruin. 2020. Why Are We Like This?: The AI Architecture of a Co-Creative Storytelling Game. In *Proceedings of the Fifteenth International Conference on the Foundations of Digital Games*.
- [48] Kenneth Lee. 2004. *Principles of Metadesign: Processes and Levels of Co-Creation in the New Design Space*. Ph.D. Dissertation. University of Plymouth, Plymouth, UK. Advisor(s) Giaccardi, Elisa. CAiiA-STAR.
- [49] Yi-Chieh Lee, Naomi Yamashita, and Yun Huang. 2020. Designing a Chatbot as a Mediator for Promoting Deep Self-Disclosure to a Real Mental Health Professional. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW1 (2020), 1–27.
- [50] Zhuying Li, Yan Wang, Wei Wang, Stefan Greuter, and Florian 'Floyd' Mueller. 2020. Empowering a Creative City: Engage Citizens in Creating Street Art through Human-AI Collaboration. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–8.
- [51] Yuyu Lin, Jiahao Guo, Yang Chen, Cheng Yao, and Fangtian Ying. 2020. It Is Your Turn: Collaborative Ideation With a Co-Creative Robot through Sketch. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [52] Duri Long, Mikhail Jacob, and Brian Magerko. 2019. Designing co-creative ai for public spaces. In *Proceedings of the 2019 on Creativity and Cognition*. 271–284.
- [53] Duri Long, Lucas Liu, Swar Gujranja, Cassandra Naomi, and Brian Magerko. 2020. Visualizing Improvisation in LuminAI, an AI Partner for Co-Creative Dance. In *Proceedings of the 7th International Conference on Movement and Computing*. 1–2.
- [54] Ryan Louie, Andy Coenen, Cheng Zhi Huang, Michael Terry, and Carrie J. Cai. 2020. Novice-AI Music Co-Creation via AI-Steering Tools for Deep Generative Models. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3313831.3376739>
- [55] Kai Lukoff, Taoxi Li, Yuan Zhuang, and Brian Y Lim. 2018. TableChat: mobile food journaling to facilitate family support for healthy eating. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 1–28.
- [56] Michal Luria, Joseph Seering, Jodi Forlizzi, and John Zimmerman. 2020. Designing Chatbots as Community-Owned Agents. In *Proceedings of the 2nd Conference on Conversational User Interfaces*. 1–3.
- [57] Mao Mao, Alan F. Blackwell, Johanna M. Lukate, and David A. Good. 2016. Supporting Retirement Socially and Musically by Technology: An Ethnographic Study of Local Community Musicians. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems* (San Jose, California, USA) (CHI EA '16). Association for Computing Machinery, New York, NY, USA, 2886–2892. <https://doi.org/10.1145/2851581.2892285>
- [58] Evandro Manara Miletto, Luciano Vargas Flores, Marcelo Soares Pimenta, Jerome Rutily, and Leonardo Santagada. 2007. Interfaces for Musical Activities and Interfaces for Musicians Are Not the Same: The Case for Codes, a Web-Based Environment for Cooperative Music Prototyping. In *Proceedings of the 9th International Conference on Multimodal Interfaces* (Nagoya, Aichi, Japan) (ICMI '07). Association for Computing Machinery, New York, NY, USA, 201–207. <https://doi.org/10.1145/1322192.1322228>
- [59] Changhoon Oh, Jungwoo Song, Jinhan Choi, Seonghyeon Kim, Sungwoo Lee, and Bongwon Suh. 2018. I lead, you help but only with enough details: Understanding user experience of co-creation with artificial intelligence. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [60] Alex Osborn. 2012. *Applied imagination-principles and procedures of creative writing*. Read Books Ltd.
- [61] Sunjeong Park and Youn-kyung Lim. 2020. Investigating User Expectations on the Roles of Family-Shared AI Speakers. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3313831.3376450>
- [62] Paul B Paulus and Bernard A Nijstad. 2003. *Group creativity: Innovation through collaboration*. Oxford University Press.
- [63] Maria B Rasmussen, Martin G Tolsgaard, Peter Dieckmann, Doris Østergaard, Jonathan White, Pernille Plenge, and Charlotte V Ringsted. 2020. Social ties influence teamwork when managing clinical emergencies. *BMC medical education* 20, 1 (2020), 1–7.
- [64] Adam Roberts, Curtis Hawthorne, and Ian Simon. 2018. Magenta.js: A javascript api for augmenting creativity with deep learning. (2018).
- [65] Aleksandra Sarcevic, Ivan Marsic, and Randall S. Burd. 2020. Challenges for AI Approaches to Activity Recognition in Teamwork. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–8.
- [66] Frederick Seddon. 2006. Collaborative computer-mediated music composition in cyberspace. *British Journal of Music Education* 23, 3 (2006), 273–283.
- [67] Isabella Seeber, Eva Bittner, Robert O Briggs, Triparna de Vreede, Gert-Jan De Vreede, Aaron Elkins, Ronald Maier, Alexander B Merz, Sarah Oeste-Reiß, Nils Randrup, et al. 2020. Machines as teammates: A research agenda on AI in team collaboration. *Information & management* 57, 2 (2020), 103174.
- [68] Joseph Seering, Juan Pablo Flores, Saiph Savage, and Jessica Hammer. 2018. The social roles of bots: evaluating impact of bots on discussions in online communities. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 1–29.
- [69] Yang Shi, Nan Cao, Xiaojuan Ma, Siji Chen, and Pei Liu. 2020. EmoG: Supporting the Sketching of Emotional Expressions for Storyboarding. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [70] Enno Siemsen, Aleda V Roth, Sridhar Balasubramanian, and Gopesh Anand. 2009. The influence of psychological safety and confidence in knowledge on employee knowledge sharing. *Manufacturing & Service Operations Management* 11, 3 (2009), 429–447.
- [71] Ian Simon, Dan Morris, and Sumit Basu. 2008. MySong: automatic accompaniment generation for vocal melodies. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 725–734.
- [72] Renee Tynan. 2005. The effects of threat sensitivity and face giving on dyadic psychological safety and upward communication 1. *Journal of Applied Social Psychology* 35, 2 (2005), 223–247.
- [73] Dakuo Wang, Elizabeth Churchill, Pattie Maes, Xiangmin Fan, Ben Shneiderman, Yuanchun Shi, and Qianying Wang. 2020. From Human-Human Collaboration to Human-AI Collaboration: Designing AI Systems That Can Work Together with People. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–6.
- [74] Mairieli Wessel, Bruno Mendes De Souza, Igor Steinmacher, Igor S Wiese, Ivanilton Polato, Ana Paula Chaves, and Marco A Gerosa. 2018. The power of bots: Characterizing and understanding bots in oss projects. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 1–19.
- [75] Christine Wolf and Jeanette Blomberg. 2019. Evaluating the promise of human-algorithm collaborations in everyday work practices. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–23.
- [76] Deepika Yadav, Prerna Malik, Kirti Dabas, and Pushpendra Singh. 2019. Feedpal: Understanding opportunities for chatbots in breastfeeding education of women in india. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–30.
- [77] Lei Zheng, Christopher M Albano, Neev M Vora, Feng Mai, and Jeffrey V Nickerson. 2019. The roles bots play in Wikipedia. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–20.