

Patterns of coordination in simultaneously and sequentially improvising jazz musicians

Matt Setzler (msetzler@iu.edu)

Cognitive Science Program, 1900 E. 10th St.
Bloomington, IN 47406 USA

Rob Goldstone (rgoldsto@indiana.edu)

Department of Psychological and Brain Sciences, 1101 E 10th St
Bloomington, IN 47405 USA

Abstract

In Joint Action (JA) tasks, individuals must coordinate their actions so as to achieve some desirable outcome at the group-level. Group function is an emergent outcome of ongoing, mutually constraining interactions between agents. Here we investigate JA in dyads of improvising jazz pianists. Participants' musical output is recorded in one of two conditions: a *real* condition, in which two pianists improvise together as they typically would, and a *virtual* condition, in which a single pianist improvises along with a "ghost partner" – a recording of another pianist taken from a previous *real* trial. The conditions are identical except for that in *real* trials subjects are mutually coupled to one another, whereas there is only unidirectional influence in *virtual* trials (i.e. recording to musician). We quantify ways in which the rhythmic structures spontaneously produced in these improvisations is shaped by mutual coupling of co-performers. Musical signatures of underlying coordination patterns are also shown to parallel the subjective experience of improvisers, who preferred playing in trials with bidirectional influence despite not explicitly knowing which condition they had played in. These results illuminate how mutual coupling shapes emergent, group-level structure in the creative, open-ended and fundamentally collaborative domain of expert musical improvisation.

Keywords: Joint Action; Music; Improvisation; Complex Dynamical Systems; Situated Cognition

Introduction

Joint action (JA) is a fundamental facet of human life. From the earliest infant-caregiver interactions to the subtle give and take of salsa dancers, we very often coordinate our actions with others (Sebanz, Bekkering, & Knoblich, 2006). In such endeavors, group success has less to do with individual efforts considered in isolation, and more to do with the ability of individuals to successfully coordinate with one another. Understanding behavior in these settings requires shifting the unit of analysis up from the individual to the group level, as collective behavior emerges out of the ongoing interactions among individual agents (Goldstone & Gureckis, 2009).

The past decade has seen a proliferation of research investigating JA in collaborative music performance (Palmer & Zamm, 2017). Music has long been recognized as a rich and meaningful domain for cognitive science. It is a central facet of all human cultures, and music performance demands the simultaneous engagement of a variety of cognitive, emotional and perceptual-motor processes (Pearce & Rohrmeier, 2012). The richness and complexity of music increases still further when we consider collaborative musical performance, where all of these intra-individual processes must be aligned and

coordinated amongst an ensemble of interacting musicians in service of a joint musical expression.

JA research has begun to elucidate how musicians meet these collaborative performance demands by examining the role of anticipatory auditory imagery in enabling performers to integrate their actions with one another, and how mutual coupling and leader-follower structures within ensembles facilitate musicians' ability to synchronize and fluidly change tempos (Chang, Livingstone, Bosnyak, & Trainor, 2017; Goebel & Palmer, 2009; Keller & Appel, 2010).

Joint Action in Improvised Music

Most of the work on music JA has taken place in the context of composed music, whereas very little has been done to examine JA in improvised music. JA in improvised music is a relatively neglected topic, and constitutes a uniquely rich and promising domain for examining joint action and complexity which is especially relevant to cognitive science.

When improvising musicians perform together they collectively generate abstract musical structures – rhythm, melody, harmony and sometimes even long-term song structures. In composed music, musicians must coordinate in terms of expressive parameters (like volume, tempo and articulation) but the abstract structure of the music is given *a priori* by the composer. The domain of interpersonal coordination in improvised music extends beyond these expressive parameters and into the formal architecture of the music. Abstract musical structures emerge out of ongoing interactions among improvisers. These interactions are nonlinear, mutually constraining and have the potential to evolve over time.

In many ways JA in improvised music is more closely aligned with other everyday JA situations than is performance of scored music. Improvisation is the norm in our daily life – group problem solving, scientific collaboration, and most of our conversations are improvised. It is actually quite rare that we perform scripted activities with others (composed music, religious ceremonies and theater performances are exceptional in this regard). Given the ubiquity of improvisation in everyday life, we might well expect some aspects of collaborative improvised music to generalize to other areas of cognition.

Despite the paucity of research in this area, some efforts to understand JA in improvised music have begun. In a notable example, coordination in jazz piano duos was analyzed as a

function of musical context (Walton et al., 2018). In the experiment, dyads of jazz pianists were studied improvising together over a swing backing track and a drone (sustained tone with no rhythmic structure). The authors performed Cross Recurrence Quantification Analysis on recordings of musicians' body movements as well as recordings of their musical output. CRQA revealed that pianists spontaneously engaged in different patterns of interpersonal coordination depending on which musical setting they were performing in.

In the current study we directly examine the effects of mutual coupling in improvised music by experimentally manipulating interaction in dyads of professional jazz pianists. Specifically, we recorded pianists improvising in one of two conditions: a *real* condition, in which two pianists improvised together as they typically would, and a *virtual* condition, in which one pianist improvised along with a "ghost partner" – a recording of another pianist taken from a previous *real* trial. In the *real* condition pianists are mutually coupled in the sense that they have the ability to respond to one another in ongoing feedback loops. Such mutual coupling is absent in the *virtual* condition – live pianists have the ability to respond to the recording, but the recording will never respond to the live musician. This feature also makes virtual recordings a nice ground-truth for assessing leader-follower roles. Subjects were blind to which condition they played in, and their musical output was recorded in the form of isolated MIDI tracks¹.

How does the presence of mutual coupling influence the music jointly produced by an ensemble of improvising musicians? This question is addressed by quantitatively comparing rhythmic structures spontaneously generated in *real* performances against *virtual* performances. Notably, these performances were obtained from elite professional pianists from the New York City jazz scene. These are individuals who have dedicated their lives to mastering their instruments and the ability to fluidly interact with others in improvised performance. Our subjects improvised freely, without any specific instructions or musical constraints (other than the implicit constraints imposed by manipulating interaction). The current study thus represents an ecologically valid and scientifically grounded approach to studying JA and mutual coupling in the creative, open-ended and fundamentally collaborative domain of expert musical improvisation.

Methods

Participants

16 professional jazz pianists from the New York City music scene participated in this study. Participant age ranged from 23-35. On average participants had 22 years experience playing piano (sd=4) and 17 years experience improvising (sd=5). All participants received formal training in piano performance and jazz studies at elite conservatories. None

¹MIDI is a format for representing music on a computer. It symbolically records the pitch, volume and timing (onset and offset) of every note played

of our subjects had prior experience performing with one another.

Apparatus

Two MIDI-enabled keyboards were used: a Roland Juno-Di and Nord Electro 2, both of which had 61 semi-weighted keys. Both keyboards were used on every trial (i.e. virtual trials were arranged such that the live pianist played whatever keyboard their ghost partner did *not* play). Ableton Live 9 Lite (running on a MacBook Air) was used to collect isolated MIDI recordings for each musician. Ableton was also used to synthesize the audio participants heard, which allowed us to ensure time alignment of MIDI recordings, and that participants heard the same exact timbre for themselves and their partner, irrespective of condition. Participants were recorded at a music rehearsal studio in Brooklyn, NY. The studio was divided by a curtain such that participants could not see one another. Participants listened to themselves and their partners through Sony CH700N Noise Cancelling headphones. Thus, from the participants' perspective there was no visual or audible indication of their condition on a given trial.

Procedure

This study employed a within-subjects design, in which each musician played at least 3 trials² in both real and virtual conditions (Figure 1). Participants played with the same 'live' partner for each of their real trials and the same 'ghost' partner for each of their virtual trials. Altogether, 32 (128 minutes, 105,766 notes) trials were collected from 9 real pairs and 27 trials (108 minutes, 84,439 notes) were collected from 16 virtual pairs. To control for order effects, conditions were interleaved throughout the course of a session, and sessions were counterbalanced such that the order reversed every other session.

Participants were brought into the studio in pairs, and instructed to improvise a series of short (4-7 minute) duos. These improvisations were 'free', with no accompanying stimuli and no *a priori* musical template or constraints. Other than the suggested timeframe, the only instruction musicians were given was to do their best to improvise a compelling piece of music, as they would in a typical performance setting.

Subjects were told they would be improvising in one of the two conditions (real or virtual), but on any given trial they were not told their condition. At the start of each trial each participant was privately instructed to Play or Don't Play. At the conclusion of each trial (when the musicians had finished improvising) each player was asked to fill out a short questionnaire that had them rate the previous performance in terms of: (1) how easy it was to coordinate with their partner (2) how well coordinated they were with their partner (3) quality of the improvised piece and (4) to what degree they played a supporter or a leader role.

²Subjects played more trials if time permitted.

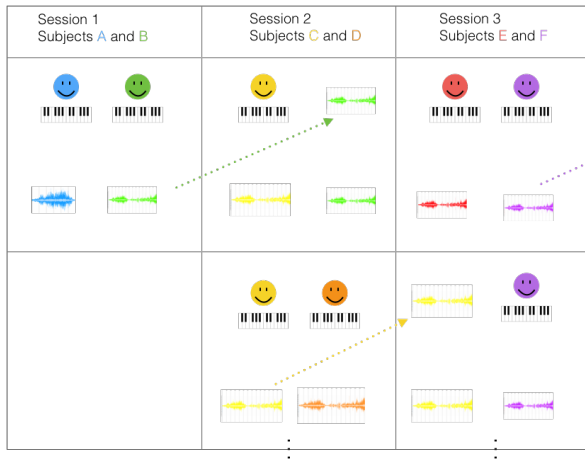


Figure 1: Sequencing of real and virtual trials. Each subject played multiple trials in each condition (repetition of trials not shown in abbreviated figure). Subjects were paired with same partner in all real trials, and a separate ghost partner for all virtual trials. Ghost recordings were taken from real trials of previous sessions, ensuring that the live musician had never heard them before.

Results and Discussion

MIDI data was collected for 32 real trials and 27 virtual trials. Each trial consists of two MIDI recordings, one for each individual (the same MIDI recordings in real trials were used as ghosts in virtual trials). 105,766 improvised note onsets were collected in real trials and 84,439 improvised note onsets were collected in virtual trials. Over 11 hours of music was collected in total. We also collected subjective ratings of participants after every trial they performed.

Figure 2 shows participants' responses to the questionnaire they were given at the conclusion of each trial. Despite the fact that participants were blind to which condition they were in in a given trial, their ratings differed systematically as a function of condition. Overall, subjects rated real trials to be of higher quality than virtual trials (paired $T(df)=15$, $p<.01$). Real trials were also generally rated as being characterized by better inter-musician coordination (paired $T(df)=15$, $p<.01$) and ease of collaboration (paired $T(df)=15$, $p<.01$).

Subjects were also asked to rate the degree to which they felt they played a leader or supporter role, which also revealed a main effect of condition (paired $T(df)=15$, $p<.05$). As expected, participants felt they mostly played a supporter role in virtual trials (in which they were playing with an unresponsive recording), whereas participants neither identified with leader or follower roles in real trials. This last result could indicate multiple things. One possibility is that musicians felt they played an equally leading and supporting role throughout the course of the performance. Alternatively, it could be that leadership roles shifted throughout the course of improvised performances. More data would be needed to differentiate between these possibilities, but in informal conversa-

tion with subjects they often alluded to the latter. At the very least, time-evolving leadership dynamics were achievable in real trials characterized by mutual coupling, but not in virtual trials characterized by unidirectional influence.

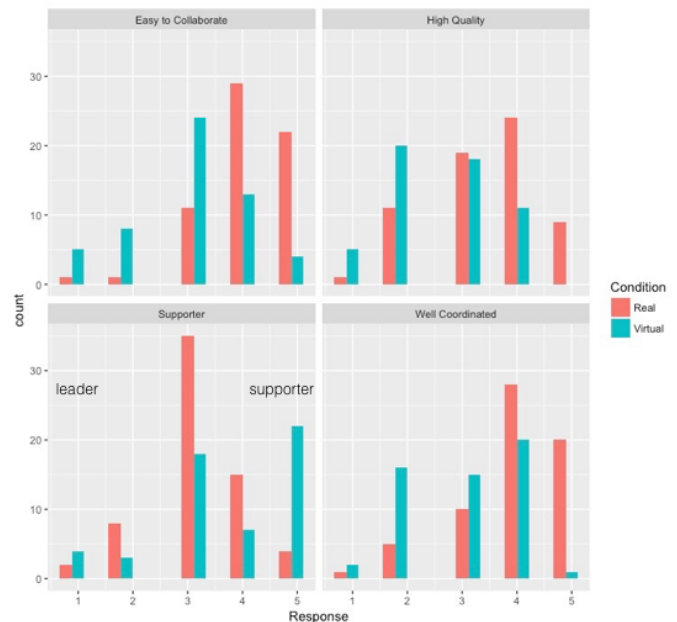


Figure 2: Subjective ratings by condition. Despite being blind to condition, participants generally rated *real* trials to be of higher quality (top right) and characterized by better inter-musician coordination (bottom right) and ease of collaboration (top left) as opposed to *virtual* trials. Participants felt they played more of a supporter role in *virtual* trials, and generally did not identify with either a leader or supporter role on *real* trials (bottom left).

Onset Analysis

MIDI recordings contain a wealth of musical information: rhythmic structure (timing of note onsets and offsets), volume, and tonal structure (sequential pitch information). In expert improvisation, interpersonal coordination occurs in each of these musical dimensions. However, we initially analyzed one clear and unambiguous aspect of the data – timing of note onsets. Timing of note onsets is a good starting point for investigating inter-musician coordination because it is simple to analyze but encapsulates an essential musical component – rhythmic structure.

Synchronization A central challenge in collaborative music making is synchronization. Musicians playing together often need to align their note onsets to occur simultaneously. Previous work has demonstrated that in composed musical settings, piano dyads' synchronize more effectively when they are mutually coupled to another another than in experimental manipulations in which auditory feedback was removed (Demos, Carter, Wanderley, & Palmer, 2017; Goebel & Palmer, 2009). It has also been demonstrated that musical

leaders play onsets of nominally simultaneous notes (notes occurring at the same metrical positions in a written score) slightly before followers (Goebel & Palmer, 2009).

To what degree does mutual coupling facilitate synchronization in improvised music? Without a written score it is difficult to assess this question, because there is no ‘ground truth’ for when and whether improvisers are trying to synchronize. Nonetheless, we approached the question by identifying near-simultaneous onsets, those occurring within 100ms of one another, played by co-performers. Degree of synchrony can be assessed by looking at the magnitude of asynchronies (henceforth ‘asyns’) by which near-simultaneous onsets were displaced from one another. While we cannot be certain whether improvisers were explicitly trying to synchronize, this metric gives us insight into how precisely synchronization occurred spontaneously, as a joint outcome of our subjects’ sensibilities and the affordances of inter-musician coupling.

Figure 3 displays the magnitude of onset asyns colored by experimental condition. Asyns are more peaked around 0 for real trials compared to virtual trials – indicating that when co-performers did synchronize, they did so more precisely in real conditions compared to virtual conditions. A Kolmogorov-Smirnov test confirmed a significant difference between asyn distributions in each condition ($p < .01$). This reproduces the result of past work showing that mutual coupling promotes greater synchronization in piano dyads in an improvised context (Demos et al., 2017; Goebel & Palmer, 2009).

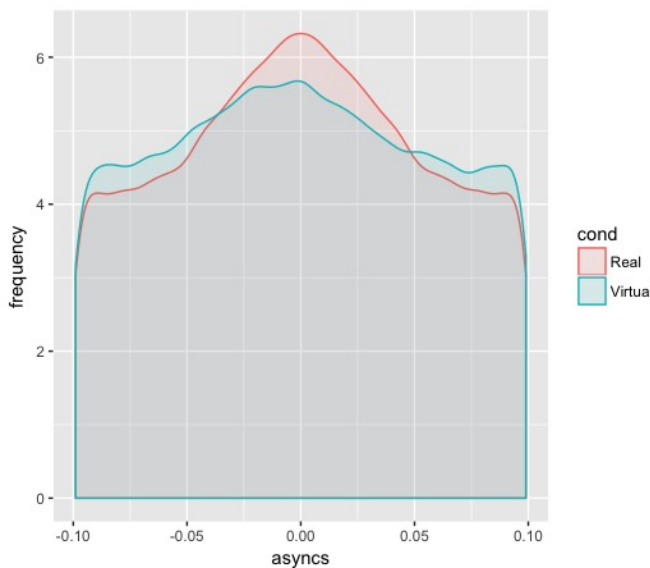


Figure 3: Mutual coupling promotes better synchronization. Density plot of asynchronies between co-performers’ near-simultaneous (occurring within 100 ms of one another) note onsets. Asyns are more tightly clustered around 0 seconds in *real* trials, in which mutual coupling is present.

Async frequency is symmetric around 0 for real trials because the same asyn was computed once for each partner (and thus represented twice with opposite signs). To assess asymmetries in virtual trials, asyns were only computed by subtracting onset timestamps of ghost partners from the timestamps of live musicians. Thus positive asyns indicate that the ghost led the live musician and negative asyns indicate the reverse. Given past work which demonstrated musical leaders in composed settings play onsets slightly before other ensemble members, we were interested to see if ghost recordings (*de facto* leaders in virtual conditions) would lead the live players (Goebel & Palmer, 2009). However, the mean asyn across all virtual trials was less than 1 millisecond, indicating a symmetry between how often and how much live players led ghosts and vice versa.

Onset Density Given the lack of musical constraints, the improvised performances in our dataset exhibited high variability in rhythmic structure. Such variability could be found not just between subjects and trials, but even within particular performances. Tempos sped up and slowed down, and dyads moved in and out of “time” – sometimes playing *rubato* sections that lacked any steady pulse. Even within a given tempo, improvisers had the freedom to play more or fewer notes. To index all of this rhythmic variety, we compute onset density for each performer as the number of note onsets occurring within a given time window. Onset density was computed for each trial using a sliding window of 2 seconds and step size of 0.2 seconds, resulting in one onset density time series per subject-trial (Figure 4A).

How is inter-musician rhythmic coordination influenced by the presence or lack of mutual coupling? This question was approached by looking at cross-correlations between co-performers’ onset density throughout the course of each trial. Cross-correlation was computed across a range of lags (± 5 seconds) to test for longer-term system memory and directional influence from one musician to another (Figure 4). Overall there was significantly greater cross-correlation in onset density in real trials as opposed to virtual (Figure 4B). This was confirmed with a Mann-Whitney test performed on the distributions collapsing over all time lags for all real trials (mean=.535, sd=.237) and virtual trials (mean=.356, sd=.287); $p < .01$.

Figure 4C displays how cross-correlation varies across lags as a function of condition. At each lag, the mean cross-correlation was obtained for all trials in each condition. In virtual trials we see greater cross-correlation at positive time lags, indicating that onset density of live musicians was more correlated with onset density of ghosts in previous time steps, as opposed to the other way around. This reflects the unidirectional influence inherent in virtual trials, whereby live musicians were influenced by their ghost partners but not the other way around. Onset density is symmetric around 0 for real trials, because data from each co-performer in a given trial was included. But it is interesting to note the dip

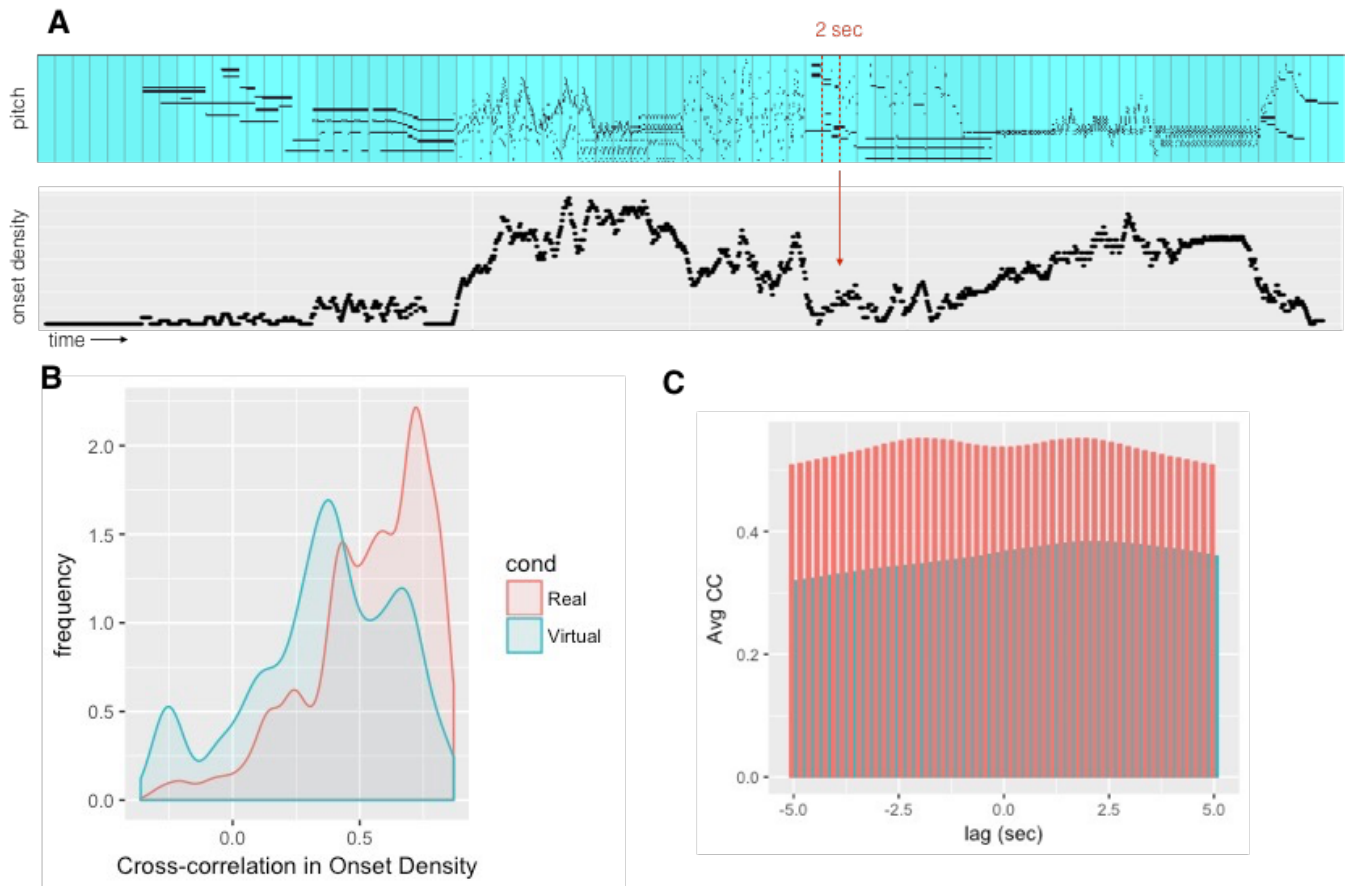


Figure 4: Mutual coupling promotes tighter coordination in onset density. (A) Time series of onset density (lower plot) are obtained by tallying the number of note onsets that occur within a 2 second sliding window of the MIDI recording (upper plot). (B) Cross-correlation in co-performer onset density, collapsing over a range of lags (± 5 sec). Overall there is greater cross-correlation in real trials. (C) Cross-correlation in co-performer onset density by lag, averaged across all trials in each condition. Greater cross-correlation in positive lags for virtual conditions reflects ground truth asymmetry inherent in this condition (i.e. live musician can respond to the recording, but not vice versa).

in cross-correlation around lag 0, surrounded by increased cross-correlation around lags ± 2 sec. In real performances, co-performers' onset density is less correlated right at simultaneous time points, and more correlated at small lags of around 2 seconds. More analysis would be needed to elucidate this pattern, but it could be a result of "call and response" interplay between improvisers, whereby they exchange musical gestures in an interleaved manner.

General Discussion

In this work we have quantitatively demonstrated ways in which inter-musician mutual coupling in improvising jazz ensembles influences the music they spontaneously produce. Specifically, we showed that mutual coupling facilitates more effective coordination in jointly producing rhythmic structure. Musicians synchronized more precisely in performances where they were mutually coupled, and exhibited tighter coupling (greater cross-correlation) in onset density – a metric that captures tempo change and overall rhythmic activity.

Subjects were coupled the most not at simultaneous times, but at small lags of about 2 seconds, suggesting a natural timescale of interaction. We also observed a quantitative artifact of musical leadership, as the onset density of improvisers in virtual trials was more correlated more with onset density of ghost partners at previous time points (again at about a 2 second lag). These objective results parallel the subjective intuitions of our performers, who rated trials with mutual coupling to be of higher quality and characterized by better coordination which was easier to achieve than in conditions with unidirectional influence.

When one listens to a great jazz combo, they are not merely listening to the sounds produced, but also to the complex underlying patterns of interaction which give rise to those sounds. This work provides the first controlled investigation of quantitatively measured coordination patterns demonstrated by freely interacting jazz musicians. It builds on prior research that studied the affordances of mutual coupling amongst co-performers by experimentally manipulating in-

teraction to reduce coupling in control conditions (Goebel & Palmer, 2009; Demos et al., 2017). But whereas past work focused on fine-temporal structure and movement dynamics exhibited by pianists, studying improvisers provided the opportunity to examine how mutual coupling influences larger scale musical structures (such as onset density) that are spontaneously generated in joint performance.

In the future we plan to delve deeper into the dynamics of how musicians pick up on and respond to the melodic and rhythmic offerings of their partner. One of the mesmerizing capabilities of expert jazz musicians, such as the pianists recorded in the current work, is their ability to improvise music with coherent and compelling tonal structure (melody and harmony). In ensemble performance, this melodic and harmonic structure emerges out of the ongoing interactions and musical negotiations taking place between ensemble members. In the future we plan to use this same (and expanding) dataset to delve deeper into how mutual coupling affords the emergence of stable tonal structure. Information theory offers a promising framework for inferring synergy and causality in multivariate time series of discretized, non-ordinal data, such as the musical pitches used in our data (Williams & Beer, 2010; Runge, 2018).

We also plan to extend our analyses to investigate the dynamical structure of our performances. Improvised music is essentially dynamic: the ensemble-generated musical structures evolve over time, as do the patterns of interaction between ensemble members. This is immediately evident observing the exemplar MIDI recording in Figure 4A, which appears to transition between regimes of sparse, sustained tones, and pointilistic sections characterized by short punctuated notes. Indeed, such dynamical structure is a central component of what makes improvised music so compelling.

To take another example, it could be the case that cross-correlation in onset density changes throughout performances. Imagine a performance in which there is initially negative cross-correlation between co-performers' onset density. This may transition to a period of positive cross-correlation, which could then be followed by yet another segment exhibiting no cross-correlation, in which performers go off on independent trails of rhythmic exploration. Such time-evolving interpersonal coordination would be lost on our current analysis (in fact it would obscure our results), but could be identified by computing cross-correlation over a sliding window and analyzing how it varies over the course of a performance.

It is also likely the case that leadership roles shift throughout improvised performances. Without a well-established social structure (as exists in many forms of composed music), the distribution of leadership in ensembles is free to evolve in a self-organized fashion. Given our finding that musical leadership is associated with increased lagged cross-correlation of onset density (where followers are more influenced by the prior rhythmic activity of leaders), a sliding window analysis of onset density cross-correlation may also provide insight

into the dynamical patterns governing time-evolving social structures. These kinds of higher-order analysis are a promising avenue towards contributing to the joint (mutually coupled!) efforts of empirical joint action studies and modeling of complex dynamical systems (Richardson, Dale, & Marsh, 2014).

References

- Chang, A., Livingstone, S. R., Bosnyak, D. J., & Trainor, L. J. (2017). Body sway reflects leadership in joint music performance. *Proceedings of the National Academy of Sciences*, 201617657.
- Demos, A. P., Carter, D. J., Wanderley, M. M., & Palmer, C. (2017). The unresponsive partner: roles of social status, auditory feedback, and animacy in coordination of joint music performance. *Frontiers in psychology*, 8, 149.
- Goebel, W., & Palmer, C. (2009). Synchronization of timing and motion among performing musicians. *Music Perception: An Interdisciplinary Journal*, 26(5), 427–438.
- Goldstone, R. L., & Gureckis, T. M. (2009). Collective behavior. *Topics in cognitive science*, 1(3), 412–438.
- Keller, P. E., & Appel, M. (2010). Individual differences, auditory imagery, and the coordination of body movements and sounds in musical ensembles. *Music Perception: An Interdisciplinary Journal*, 28(1), 27–46.
- Palmer, C., & Zamm, A. (2017). Empirical and mathematical accounts. *The Routledge Companion to Embodied Music Interaction*, 370.
- Pearce, M., & Rohrmeier, M. (2012). Music cognition and the cognitive sciences. *Topics in cognitive science*, 4(4), 468–484.
- Richardson, M. J., Dale, R., & Marsh, K. L. (2014). Complex dynamical systems in social and personality psychology. *Handbook of research methods in social and personality psychology*, 253.
- Runge, J. (2018). Causal network reconstruction from time series: From theoretical assumptions to practical estimation. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 28(7), 075310.
- Sebanz, N., Bekkering, H., & Knoblich, G. (2006). Joint action: bodies and minds moving together. *Trends in cognitive sciences*, 10(2), 70–76.
- Walton, A. E., Washburn, A., Langland-Hassan, P., Chemero, A., Kloos, H., & Richardson, M. J. (2018). Creating time: Social collaboration in music improvisation. *Topics in cognitive science*, 10(1), 95–119.
- Williams, P. L., & Beer, R. D. (2010). Nonnegative decomposition of multivariate information. *arXiv preprint arXiv:1004.2515*.