

Frame-differencing methods for measuring bodily synchrony in conversation

Alexandra Paxton · Rick Dale

© Psychonomic Society, Inc. 2012

Abstract The study of interpersonal synchrony examines how interacting individuals grow to have similar behavior, cognition, and emotion in time. Many of the established methods of analyzing interpersonal synchrony are costly and time-consuming; the study of bodily synchrony has been especially laborious, traditionally requiring researchers to hand-code movement frame by frame. Because of this, researchers have been searching for more efficient alternatives for decades. Recently, some researchers (e.g., Nagaoka & Komori (*IEICE Transactions on Information and Systems*, 91(6), 1634–1640, 2008); Ramseyer & Tschacher, 2008) have applied computer science and computer vision techniques to create frame-differencing methods (FDMs) to simplify analyses. In this article, we provide a detailed presentation of one such FDM, created by modifying and adding to existing FDMs. The FDM that we present requires little programming experience or specialized equipment: Only a few lines of MATLAB code are required to execute an automated analysis of interpersonal synchrony. We provide sample code and demonstrate its use with an analysis of brief, friendly conversations; using linear mixed-effects models, the measure of interpersonal synchrony was found to be significantly predicted by time lag ($p < .001$) and by the interaction between time lag and measures of interpersonal liking ($p < .001$). This pattern of results fits with existing literature on synchrony. We discuss the current limitations and future directions for FDMs, including their use as part of a larger methodology for capturing and analyzing multimodal interaction.

Keywords Synchrony · Movement · Liking · Alignment · Frame-differencing methods · Image processing · MATLAB

Conversation is arguably one of the most common—and important—modes of social interaction. Combining a variety of intrapersonal and interpersonal mechanisms, conversation presents a rich source of data for researchers in numerous areas, from linguistics and affect to posture and gesture. Interpersonal synchrony research lies at the intersection of many of these areas, seeking to characterize the way that interlocutors (individuals involved in conversation) grow to have similar behavior, cognition, and emotion over time. Many areas of research in the behavioral sciences have approached the issue of synchrony, resulting in a scattered terminology: accommodation (Giles & Smith, 1979), alignment (Pickering & Garrod, 2004), coordination (Richardson & Dale, 2005), coupling (Shockley, Santana, & Fowler, 2003), entrainment (Brennan & Clark, 1996), mimicry (Chartrand & Bargh, 1999), and social tuning (Valdesolo, & DeSteno, 2011), among others.¹

The overwhelming growth of this research in recent years, including the diverse range of terms and concepts different researchers invoke, only adds to the importance of precise measurements and analytical methods. Research in this area has seen a recent push toward computer vision and computer-aided analysis techniques that can streamline objective measurement while significantly improving

A. Paxton (✉) · R. Dale
Cognitive and Information Sciences, School of Social Sciences,
Humanities, and Arts, University of California, Merced,
Merced, CA 95343, USA
e-mail: paxton.alexandra@gmail.com
URL: <http://graduatestudents.ucmerced.edu/aloan/>

¹ For the purpose of this article, we simply refer to these processes as *synchrony*, although additional research may help to determine relevant differences among these terms.

efficiency and cost effectiveness. Attempts to standardize methodology for this area of research may even lead to stricter definitions of terms, enriching the field while helping researchers communicate the precise nature of their work.

Issues with synchrony data collection and analysis

Interpersonal synchrony research faces problems not found in traditional research areas within experimental psychology. For example, synchrony research often centers on the dyad rather than the individual. This can lead to smaller sample sizes, since dyads are more costly and difficult to recruit. Conversations unfold over minutes, not milliseconds, and a single data collection session from one dyad may last several hours (e.g., Skuban, Shaw, Gardner, Supplee, & Nichols, 2006). Data collection for these studies, therefore, involves a significant investment of time by a researcher interested in interpersonal processes.

After collecting potentially dozens of hours of interaction for an experiment, researchers must spend even more time analyzing the data. Conversations must be transcribed for analyses of linguistic synchrony, but a single hour of dialogue may require over 10 h to transcribe (Kreuz & Riordan, 2011). Analysis of bodily synchrony, the primary focus of this article, has historically required researchers to meticulously hand-code limb movement in videotaped interactions frame by frame (e.g., Condon & Sander, 1974). Some postcoding automated techniques have been developed to detect repeated patterns of movement synchrony (e.g., THEME; Grammer, Kruck, & Magnusson, 1998); while inarguably helpful in identifying meaningful patterns of movement, these techniques do not mitigate the labor-intensive hand-coding process. Some believe that these issues have likely discouraged some researchers from studying interpersonal synchrony due to lack of funding or insufficient staffing (Bernieri, Davis, Rosenthal, & Knee, 1994).

Although significant, the challenges presented by data collection and analysis are not insurmountable. Researchers have been refining cost-effective and efficient methods of studying synchrony for decades, and research on interpersonal synchrony has unveiled new ways of exploring conversation through facilitated analysis (e.g., Bernieri et al., 1994). Improved research methods may minimize the restrictions imposed by minimal funding and can be combined with other methods to explore questions of cross-channel synchrony and interaction (e.g., affect and body movement; see the [General Discussion](#) section).

Existing alternatives to hand-coding for bodily synchrony

Holistic ratings

One of the most established alternatives to hand-coding involves holistic ratings by judges. The specific methods employed by each researcher vary, but all have a relatively similar general procedure. Judges may be completely naïve (e.g., Bernieri, Reznick, & Rosenthal, 1988) or strictly trained (e.g., Criss, Shaw, & Ingoldsby, 2003; Grammer, Honda, Jüette, & Schmitt, 1999), depending on the goals of the study. Judges are commonly instructed to watch videotaped interactions and provide a rating of the interaction, typically based on Likert scales of general interaction qualities (e.g., Bernieri et al., 1988; Criss et al., 2003). The interlocutors' dialogue may be muted (e.g., Bernieri et al., 1988) or included in the raters' materials (e.g., Criss et al., 2003); both have been established as equally effective as measures of bodily synchrony (Bernieri et al., 1994). Each interaction may be rated only once (e.g., Bernieri et al., 1988) or, to ensure high interrater reliability, by multiple raters (e.g., Criss et al., 2003).

Holistic ratings often require significantly less time than frame-by-frame analyses, but they are not without their own methodological problems. With the exception of event-based counting methods (e.g., Skuban et al., 2006), holistic ratings are almost entirely subjective. Intensive judge training and employing multiple raters may decrease subjectivity, but they increase the amount of time required for analysis (e.g., the 6-week training course used by Criss et al., 2003). Because synchrony is often based on measures with fewer than a dozen items, these methods often provide less within-subjects power for statistical tests.

Researchers have improved holistic ratings, but these methods remain unable to objectively quantify bodily synchrony. Ratings may be effective for studying how individuals perceive synchrony, but their inherent subjectivity limits the degree to which researchers can parse apart the mechanisms behind synchrony. While these methods are significantly more efficient, judges' holistic ratings lose the precision of Condon and Sander's (1974) hand-coding methods.

Automated video analysis

Other researchers have begun attempts at automating analyses of bodily synchrony. Although these new methods are accompanied by new difficulties, they provide significant advantages over other methods proposed to date. Many of the methods blend computer vision techniques with psychological research to create rater-free, coding-free analytical techniques. Computer-driven techniques minimize researcher

interaction with raw data, thereby removing the subjectivity of holistic ratings and the labor of hand-coding. These analyses are intended to be efficient, content-free methods for assessing bodily synchrony during interaction.

Motion-tracking systems appear to be an ideal candidate for tracking interlocutors' body movement over time. Other areas of research have already begun utilizing these systems' automated collection of data and computation of movement-related variables for the whole body and individual body parts (e.g., Battersby, Lavelle, Healey, & McCabe, 2008; Lu & Huenerfauth, 2010). However, existing motion-tracking systems are almost as restrictive in their own right as hand-coding methods. Current systems are expensive and can present methodological concerns (Welch & Foxlin, 2002). For example, specialized motion-tracking suits are often tight-fitting; participants may feel discomfort, impacting naturalistic movement. Once these systems become cheaper and less restrictive, motion tracking may become a standard tool for bodily synchrony research. Nevertheless, for researchers facing limitations in funding and for those whose questions are not compatible with the high-tech motion capture requirements, body-suit motion capture still poses significant challenges.

Meservy et al. (2005) have pioneered another appealing method. Their methodology—similar to attempts at automated blob analysis (e.g., Lu, Tschepnakis, Metaxas, Jensen, & Kruse, 2005)—is intended to automatically track patterns of head and hand movement in videos captured in moderately high quality. However, as presented in their article, the program is only partially automated; it currently requires a significant investment of time at the beginning of analysis and almost constant guidance throughout the process. It also poses restrictions for interaction researchers: Videos must be shot head-on with a single participant in the image, creating problems for applications in naturalistic conversation between two interlocutors. While interesting, Meservy et al.'s paradigm is not yet feasible for interaction research.

We believe that the most promising and effective interpersonal bodily synchrony techniques to date are what we will call *frame-differencing methods* (FDMs). Rather than tracking specific body parts, FDMs are grounded in research showing that interlocutors synchronize in overall body movement in addition to posture and gesture (e.g., Shockley et al., 2003). Encompassing several existing named (e.g., motion energy analysis [MEA]; Grammer et al., 1999; Ramseyer & Tschacher, 2008, 2011) and unnamed (e.g., Nagaoka & Komori, 2008) methods, FDMs track the changes in pixels from one frame to the next. These methods require the background of an image to remain static, so the only pixel changes from frame to frame will likely be caused by interlocutors' movement. FDMs generally analyze movement quantitatively by strictly measuring pixel changes

between frames (e.g., Nagaoka & Komori, 2008), but some FDMs utilize qualitative analyses of movement (e.g., judges' ratings of movement FDM-derived visualizations; Grammer et al., 1999).

FDM data collection setups (e.g., Nagaoka & Komori, 2008; Ramseyer & Tschacher, 2008, 2011) generally have similar requirements. They often require only one or two unmoving video cameras and stable ambient lighting, making FDMs highly cost effective. Prior to analysis, the video data are often transformed—manually or automatically—into grayscale images and normalized for brightness. Existing FDMs are indifferent to many movement characteristics (e.g., direction), and they have been used successfully in several studies to date, primarily in clinical (e.g., Kupper, Ramseyer, Hoffmann, Kalbermatten, & Tschacher, 2010; Nagaoka & Komori, 2008; Ramseyer & Tschacher, 2008, 2011) and ethological (e.g., Grammer et al., 1999) domains. Our goal in this article is to present FDM to those interested in basic experimental research on conversation and to develop a simple version of an FDM that can be run with minimal programming experience in MATLAB code, which we supply in [Appendix 2](#).

By presenting a template for a very basic but nevertheless powerful FDM, we hope to provide experimental researchers with a tool that can be easily modified according to their research needs. We also point researchers to existing methods in nonexperimental fields to explore additional ways of implementing similar analyses (e.g., Grammer et al., 1999; Ramseyer & Tschacher, 2011).

A simple frame-differencing method

In this article, we showcase a highly simplified, MATLAB-based method of extracting overall body movement between two people engaged in conversation. Interpersonal synchrony is a highly diverse research topic, comprising researchers from various fields and technical backgrounds. The FDM presented here is based on modifications to existing methods and may provide an affordable, efficient, yet robust source of data to explore how bodily synchrony relates to conversation. A script only a few lines in length provides the basic measures, and analyses can be performed quickly and with very little effort by the researcher. While some semi-automated analyses require researchers to specify individual areas of interest to be analyzed (regions of interest [ROIs] for MEA; e.g., Ramseyer & Tschacher, 2011), many FDMs—including the one presented here—analyze overall body movement (e.g., Boker, Xu, Rotondo, & King, 2002). By combining existing FDM-based techniques and contributing some additions, the FDM offered here provides an analytical method for researchers equipped only with a moderate-to high-quality digital video recorder, standard analysis

software (e.g., MATLAB), and very modest programming skills. After we detail an example FDM, we will demonstrate its use and effectiveness in a study on conversational interaction.

Data collection and preparation

The presented FDM has been designed to require as little direct supervision and specialized equipment as possible. For data collection, researchers will only require a digital video camera (preferably, high definition; no specific codec required),² mounted to be completely stable throughout the recording. Although the light source need not remain completely stable, it should not be subject to large fluctuations. In order to provide time-locked images, we recommend using a single camcorder to record both participants (see Fig. 1). These analysis methods can be adjusted to accommodate multiple cameras, so long as the sequences can be accurately synchronized in time. However, the description of our methods is written under the assumption that the researcher is using a single camera.

The videos must then be uploaded to a computer and segmented into image sequences; higher-quality image formats (e.g., PNG) are preferable, although not required. Assuming a high-quality recording, the images may remain in the native camcorder resolution and do not require rescaling. This can be done with a number of commercially available video processing programs, including Apple's QuickTime or iMovie. For researchers using Apple products, we have included a sample AppleScript to automate the image segmentation of videos in iMovie (see Appendix 1).³ Researchers may also use MATLAB and VideoReader to import video directly, but we chose to use out-of-the-box software to get image sequences in order to further minimize programming requirements.

The sampling rate may vary according to researcher needs and storage space. We have experimented with a number of sampling rates and have found that 8 Hz affords a great deal of detail without generating unwieldy amounts of data. In contrast with existing methods, the FDM we present can utilize full-color image sequences and does not require the images to be transformed into grayscale or normalized grayscale brightness. We calculate the frame differences using the RGB code in MATLAB's image arrays. By analyzing images in full color, this FDM is able to detect movement of an object of one color against a

² Although these methods are likely to capture movement effectively even when given provided lower-quality recordings, lower resolutions may be less sensitive to smaller body movements (e.g., postural sway, facial expressions).

³ The AppleScript code is also available for download from the first author's Web site: http://graduatestudents.ucmerced.edu/aloan/Miscellany_files/imovie_segmentation.scipt.

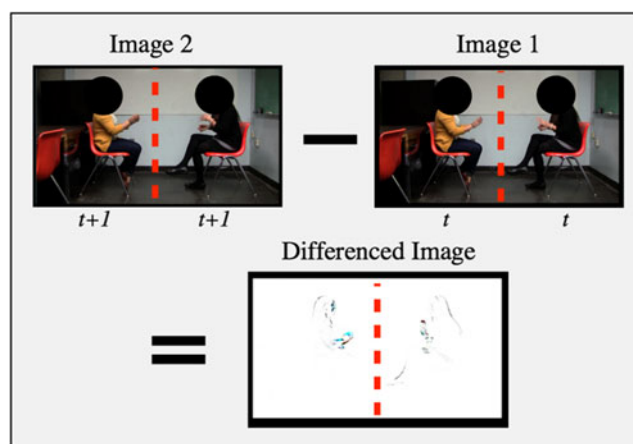


Fig. 1 Visualization of differencing process and sample participant setup. In order to establish time-locked capture of both interlocutors' movements, we recommend that participants be recorded facing one another, in profile to the camera. Additionally, no moving objects (e.g., clocks) should be included in the frame. Each image is halved, and its pixels are compared with those of its predecessor. Movement between the consecutive images is registered as changes in pixels from one image to the next. This image demonstrates the nonlagged comparison of time series (i.e., movement of participant A at time t compared with movement of participant B at time t)

background of a different color that may have the same intensity. This means that we can track changes in colors that may have the same overall intensity (i.e., the same summed color codes, but differently distributed over the red, green, and blue spectra). Any differences in intensity of a person's clothing and background will be captured with this approach.


Data analysis with MATLAB

We have written and combined a series of MATLAB scripts to create a single, short script (see Appendix 2)⁴ to automatically analyze the bodily synchrony between interlocutors in a single video frame. Using a "for loop," the MATLAB script sequentially loads each image of a given frame sequence. The images are halved so that each interlocutor's movement is on only one half of the frame; if all dyads do not have the same halfway point, the researcher must designate the halfway points for any exceptions. The script then compares the pixels of the current half-frame with the pixels of the previous half-frame, yields a raw pixel change score between the images, and then creates a standardized difference score between them (see Fig. 1 for visualization; see Fig. 2 for sample sequence).

A second-order Butterworth low-pass filter is then applied to each sequence of half-images in lieu of a threshold

⁴ The MATLAB code is also available for download from the first author's Web site: http://graduatestudents.ucmerced.edu/aloan/Miscellany_files/sample_FDM.m.

Fig. 2 Sample sequence of differenced images. The image sequence begins at the top of the left-hand column, proceeds down, then begins again at the top of the right-hand column. This represents the differencing process applied to 1.5 s of interaction



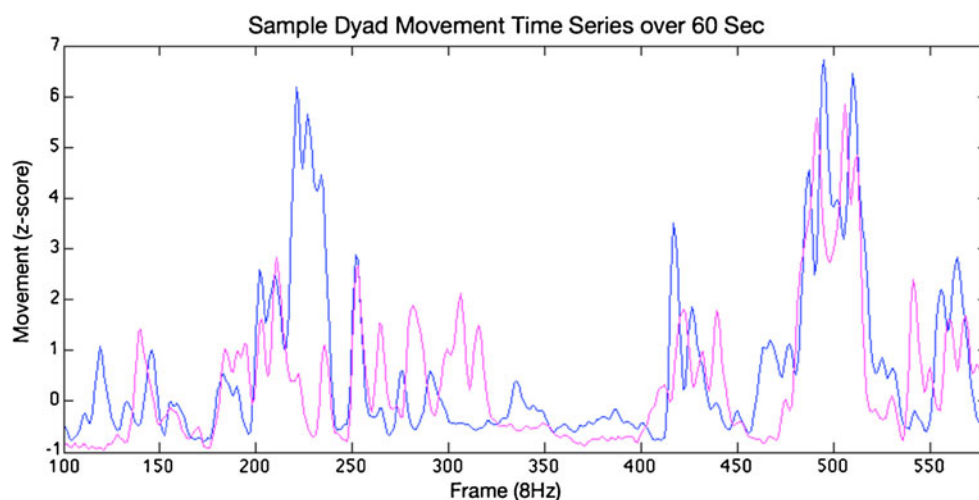
(e.g., Grammer et al., 1999) or band-pass filter (e.g., Nagaoka & Komori, 2008). A powerful but relatively simple filter, the Butterworth filter is characterized by normalized cutoff frequency, a maximally flat passband, and a stopband that slopes down to zero. By standardizing the images and applying the low-pass filter, the script is able to control for slight fluctuations in light sources (e.g., high-frequency fluctuations of fluorescent lighting) while remaining sensitive to slight movements (e.g., shifts in posture). Without a filter, co-occurring sources of fluctuation across the images may lead to false detection of synchronized movement, since these fluctuations will occur in time for both image sequences. All of this is done with a few simple lines of code in MATLAB.

The script then combines the standardized scores for the two sequences of half-images (i.e., the movement of each individual within the dyad; see Fig. 3 for a sample time series of standardized image scores) to derive cross-correlation coefficients, the measure for interpersonal synchrony, at various

time lags. In other words, a correlation coefficient is calculated for each *relative time lag* between the two interlocutors' time series. A lag of 0 would reflect the Pearson correlation coefficient between the two sequences of movements, pairing participant A's movement at time t with participant B's movement at time t . A lag of -1 would shift one time series by one step (i.e., pairing participant A's movement at t with participant B's movement at $t + 1$) and carry out a correlation again. A lag of $+1$ would then shift in the other direction (i.e., pairing participant A's movement at $t + 1$ with participant B's movement at t) and calculate r .

If two individuals' movements are synchronized, r will be highest closer to a lag of 0, reflecting that changes in their movement coincide in time. Unlike other channels of communication (e.g., speech), both interlocutors are able to move simultaneously without impeding the flow of the interaction. Individuals spontaneously synchronize in dyadic rhythmic movement tasks (e.g., Miles, Lumsden, Richardson, & Macrae, 2011; Richardson, Marsh, Isenhower, Goodman, &

Fig. 3 Sample dyad movement sequence over 60 s after the Butterworth filter has been applied to the time series. Movement scores (y -axis) are standardized, and each frame (x -axis) is 125 ms. The blue line tracks the movement of the person on the right; the purple line records the movement of the person on the left



Schmidt, 2007; Schmidt, Carello, & Turvey, 1990). These findings suggest that interlocutors may exhibit some forms of synchronous—rather than time-lagged—body movement even in naturalistic contexts.

In addition to providing an objective quantification of bodily synchrony, the cross-correlation coefficient across time lags allows for a greater exploration of trends of lagging and leading in bodily synchrony (for a discussion, see Boker et al., 2002). Because we do not have any explicit hypotheses about leading or following in this “role symmetrical” conversation in our sample study, we take the mean r between -1 and $+1$ lag, -2 and $+2$ lag, and so on.⁵ The MATLAB script then records all coefficients for analysis. The entire analysis for a 10-min dyadic interaction requires approximately 6 min on a 3.1-GHz Intel Core i5 Apple iMac computer with 4-GB 1333-MHz DDR3 memory.

Quantifying synchrony

After retrieving the cross-correlation coefficients, researchers may use them in a variety of statistical tests. Researchers may use the entire time series, a portion of the time series, an average synchrony score, or the highest/lowest synchrony scores, according to the research questions and statistical tools available (see, e.g., Caucci, 2011, for some discussion on the use of interpersonal synchrony scores in various analyses). In order to confirm that this method works to measure synchrony of body movement rather than co-occurring artifacts, we ran two validation analyses, shown

⁵ Both negative and positive raw correlations were used. The data in Fig. 3 reflect these raw correlations. We did not apply Fisher’s Z -transformation to these data because the correlations were too low to be affected (i.e., correlations of magnitude less than .5). As is discussed later, we standardized the correlations before using them in the linear mixed-effects model in order to obtain beta weights instead of raw change values.

in Appendix 3. In the next section, we present a larger study that explores how synchrony is organized in naturalistic interaction.

Interaction study

Much research has shown a possible link between affiliation and synchrony (e.g., Bernieri et al., 1994; Chartrand & Bargh, 1999; Lakin & Chartrand, 2003; Ramseyer & Tschacher, 2008). However, this synchrony–affiliation link can be modulated under various group circumstances (e.g., Miles et al., 2011). Here, we exemplify our methods in a study that shows gross-body movement synchrony during conversational interaction and tests a correlation of this synchrony with liking between interlocutors.⁶

As a proof of concept, we investigated whether individuals involved in naturalistic conversations with a broad affiliative prompt achieve bodily synchrony detectable by the FDM outlined above. The correlation coefficient should be higher closer to a lag of 0, because this correlation reflects the closest match in time. As lag increases, the time series are being correlated at a wider relative lag and are, therefore, further apart in time; synchrony would predict a *drop* in the correlation coefficients as time lag increases.

Existing literature suggests that synchrony should be positively correlated with liking. Rather than using simple correlation, the present study uses linear mixed-effects models for data analysis. We hypothesize, therefore, that the model will predict an increase in r as levels of interpersonal liking increase.

⁶ The research we present is part of a larger study we are conducting on differences in conversation types. Here, we focus on analysis of the conversations that involved the basic goal of affiliation.

Method

Participants

Participants were 40 undergraduate students at the University of Memphis (mean age = 22.08 years; females = 32) and 22 undergraduate students at University of California, Merced (mean age = 19.36 years; females = 17).⁷ All were awarded extra credit for participating. All were fluent in English. They participated in pairs, as 31 conversational dyads (19 female, 1 male, 11 mixed-sex), according to individual availability via the participant pool's online scheduler. Only two dyads (one mixed-sex, one female) reported knowing one another prior to participation in the study and were retained for all analyses. One female dyad was removed from all analyses due to experimenter error. Although a seemingly small sample size, this exceeds established dyadic research sample sizes by a moderate (e.g., 21 dyads; Ramseyer & Tschacher, 2008) or wide (e.g., 4 dyads; Boker et al., 2002; Nagaoka & Komori, 2008) margin.

Materials and procedure

After individually completing a brief questionnaire packet and signing informed consent forms, participants were brought into a private room. They were seated facing one another and were recorded in profile (see Fig. 1) to ensure that their movements were captured in a time-locked fashion. Interactions were recorded using a Canon Vixia HF M31 high-definition digital video camcorder mounted on a tripod to ensure stability. To ameliorate the potential awkwardness of interacting with a stranger, participants were allowed 3 min to introduce themselves and briefly get to know one another without the experimenter present. Following the introductory period, the experimenter prompted the participants to discuss entertainment media (e.g., movies, music) that they both enjoyed. They were instructed to talk for 10 min with the experimenter present but outside their line of sight. Experimenters ensured that all conversations lasted at least 8 min and issued additional brief prompts to keep participants on topic (mean additional prompts per conversation = 0.54). Frames during which prompts occurred were excluded from analysis. The participants were then brought to separate rooms and rated how much they liked their partner on a 1–6 Likert scale. After completing the measure, participants were brought together, debriefed, and thanked for their participation.

⁷ Previous research has shown significant differences between the interaction styles of same-sex and mixed-sex dyads, and such composition may have important ethological implications (see Grammer et al., 1998). However, we exemplify our method by showing aggregate synchrony across dyad types and reserve an analysis of gender for a later study, since it is not an immediate goal of this methodological presentation.

Data handling and analysis

The participant videos were processed and analyzed using the methods described in the preceding section. We extracted a time series of body movement at 8 Hz for each person, applied a second-order Butterworth filter to each time series, calculated the cross-correlation coefficients at each lag within a window of ± 3 s (recommended by Richardson, Dale, & Tomlinson, 2009), and standardized the resulting coefficients. These standardized coefficients served as the dependent variables in the following analyses.

Results

The standardized coefficients were predicted with a series of linear mixed-effects models to investigate basic questions of synchrony (as in Baayen, Davidson, & Bates, 2008; Boker et al., 2002). Using the standardized cross-correlation coefficients derived from the MATLAB script, bodily synchrony was defined as concurrent movement in time. Therefore, when absolute time lag is included as a predictor, r should *go down* as lag increases (from a lag of 0—matching in time—to lags reflecting greater temporal disparity). In addition, we tested whether there is a relationship between affiliation and r : We predicted that the more participants reported affiliation, the higher the standardized r would be overall. To test these questions, we included fixed factors of time lag and affiliation. All models used dyad and participant as random effects.

In the first model, we focused on synchrony as a function of time lag. This basic model tested whether individuals are more likely to move together in time. The model was found to be significant, $p < .001$ [$t(1842) = 27.6$],⁸ and predicts a *drop* in the cross-correlation coefficient with each successive time lag (i.e., 125 ms) away from 0 ($\beta = -.22$). This indicates that interpersonal synchrony is highest toward a time lag of 0, or that interlocutors' movements coincide at relatively the same amplitude in time. Put simply, individuals synchronize their body movements during conversation.

Importantly, the average peak of this function seems to be closest to 0, rather than peaking at a lag greater than 0 (see Fig. 4). Such a pattern suggests that interlocutors do not, on

⁸ Degrees of freedom are not easily defined for mixed models; t -values for mixed models, therefore, are often not included when reporting results (e.g., Boker et al., 2002). Some (e.g., Bates, 2006) have argued that reporting degrees of freedom can be misleading, given differences in techniques for obtaining them. However, there are several sources available for those who wish to report them (e.g., Baayen, 2008; Baayen et al., 2008). Degrees of freedom reported here were estimated using the LME function described therein. The t -values reported here are based on MCMC sampling using the "pvals.fnc" function in R, as described in Baayen et al., which also includes an excellent introduction to MCMC methods.

average, lead or follow in body movement patterns and that body movement is synchronized at the same relative phase.

We ran a second basic model to test whether reported levels of interpersonal liking would significantly predict the correlation coefficient. In this model, we included all the data (each r at each lag) and participants' ratings of interpersonal liking. The model was not found to be significant, $p = .84$ [$t(1842) = .102$], suggesting that interpersonal bodily synchrony is not predicted by self-report levels of liking alone.

Finally, we combined these two fixed factors into a single model, using both lag and liking (centered) to predict the correlation coefficient at each time lag. In this model, the interaction term was significant, $p < .001$ [$t(1840) = 9.37$, $\beta = -.07$]. The significance of the interaction term implies that, although liking alone does not predict r values, it can moderate the effects of time lag. To illustrate this, we split our participants into two groups, high and low liking. As can be seen in Fig. 4, individuals who like their partner more achieve higher r near lag 0 than those who do not.

To confirm that the full model was the best-fitting one, we compared the Akaike information criterion (AIC) for each model. We observed that the AIC for the first model (predicting synchrony as a function of lag; AIC = 1,401) and the second model (predicting synchrony as a function of liking; AIC = 2,140). The AIC for the saturated model (predicting synchrony as a function of lag and liking) showed it to be the model best fitted to the data (AIC = 1,355).

Discussion

In this brief study, we found that interlocutors synchronize with their partners during affiliative conversations. The results of this FDM analysis conform to patterns of results from previous FDM-analyzed research (e.g., Nagaoka & Komori, 2008; Ramseyer & Tschacher, 2008). In fact, our analyses have extended these other naturalistic studies and present novel insights into synchrony: We found that body movement tended to be synchronized *in time*, such that synchrony is greatest at a lag of 0. Thus, as a behavioral channel used during conversation, gross body movement may be patterned *in-phase* between two interlocutors. This means that body movement synchrony has properties that differ from synchrony in speech, which cannot be carried out in-phase during conversation, due to turn-taking. Other forms of movement have been demonstrated to have in-phase synchrony between individuals (e.g., Miles et al., 2011; Richardson et al., 2007; Schmidt, Morr, Fitzpatrick, & Richardson, *in press*; Schmidt et al., 1990), and FDM analyses have revealed bodily alignment in brief windows of time (e.g., 1-min windows and 5-s time-lags, Ramseyer & Tschacher, 2011; 10-min windows and 5-s time-lags,

Nagaoka & Komori, 2008). However, this is the first FDM-based analysis demonstrating millisecond-to-millisecond synchrony between interlocutors' broad body movements.⁹

Although no main effect for liking was found, levels of liking moderated interpersonal bodily synchrony: The more participants liked one another, the more closely synchronized their movements tended to be. Despite the lack of main effect, the interaction effect fits with previous research linking affiliation and body movement patterns (e.g., Chartrand & Bargh, 1999; Miles et al., 2011).

General discussion


We describe FDMs as promoting objective quantification of interpersonal (bodily) synchrony, even in small labs with minimal funding. Although several studies on interpersonal interaction have used FDMs, there is little work showing its direct relation to holistic ratings, and there is no detailed methodological presentation of it in the experimental literature. This article serves as an introduction for experimental researchers to FDMs generally and to one simplified version (see Appendix 2).

Using similar methods to existing FDMs (e.g., Nagaoka & Komori, 2008; Ramseyer & Tschacher, 2008), we have provided MATLAB code for a simple automated version, intended to minimize the required amount of postrecording editing (see Appendix 2). This simplified FDM provides researchers with added flexibility in recording setups and, in conjunction with AppleScripts to automate data preparation (see Appendix 1), allows for an almost completely automated analysis of multiple interactions at a time. By broadening data collection conditions and automating data analysis, researchers will be able to spend more time collecting dyads, leading to larger sample sizes. The use of cross-correlation coefficients as an indicator of interpersonal bodily synchrony, rather than generalized rating scores, will give statistical analyses greater power.

We hope to expand this method to include ways of parsing out the movement of individual body parts to promote more fine-grained analysis of interpersonal synchrony (e.g., posture, gesture). By combining these and other automated methods (e.g., blob analysis; Lu et al.,

⁹ We did not explore synchrony relative to baseline, but methods are available to do so. For example, Ramseyer and Tschacher (2008) offer an elegant technique of window-wise shuffling. Shockley, Baker, Richardson, and Fowler (2007) recommend using a "virtual pair" analysis in which the researcher forms baseline dyads from individuals of separate dyads in the experiment. These are relatively straightforward time-series methods that are outside the methodological scope of this article, but we point the reader to these studies in case this is of interest.

Fig. 4 Interaction effect for time lag and liking on cross-correlation coefficients. The y -axis is the mean r , and the x -axis is time lag in 125-ms segments. Self-report levels of liking have been divided into high- and low-liking groups, using a median split. The green reflects the high-liking group; the red reflects the low-liking group



2005), researchers may continue to refine the flexibility and utility of FDM-based methodologies.

Limitations

Of course, FDMs are not without their limitations. The FDM outlined here is intended to minimize cost and to automate as much of the data analysis as possible. In doing so, it loses detail afforded by other methods (e.g., movement direction and velocity, limb tracking). Other FDMs offer researchers the ability to track movement in designated areas (e.g., ROIs in MEA); these allow researchers to manually designate specific areas in which to track movement (e.g., limbs). However, even these FDMs are generally blind to movement direction and velocity. All FDMs, by using varying degrees of automated methods to detect movement, tend to underestimate participants' movements toward the camera(s). Because fewer pixels change with medial movement (relative to camera position), FDMs are far more sensitive to lateral movement.

For researchers interested in finer-grained movement characteristics, hand-coding techniques and motion tracking may prove to be worth their respective investments. Hand-coding techniques have been widely used and broadly accepted in inter- and intrapersonal synchrony research. The significant time and training required to chart each movement from frame to frame may be useful to researchers interested in tracking even participants' smallest movements.

Motion tracking may be a viable alternative to both FDMs and hand-coding for those with ample funding and strong data management resources. Currently, few researchers employ these methods for synchrony research, but these systems have unique capabilities that would help to investigate other movement-related questions, as mentioned earlier. Motion-tracking systems would permit an investigation of temporal movement dynamics more precisely than hand-coding permits. However, researchers should weigh the impact of such an invasive technique against its sensitivity to movement: Participants may be less likely to exhibit naturalistic movement patterns while wearing a tight-fitting motion-capture suit than when being filmed, which is relatively noninvasive in comparison.

Here, we have also not discussed the issue of stationarity. This is an important issue in any time series analysis using regression-based methods. Inspecting our data, we have mostly found evidence of relative stationarity (i.e., relatively unchanged mean and variance throughout each 10-min conversation). For further discussion of this issue and potential methods to manage it, see Boker et al. (2002) and Ramseyer and Tschacher (2008, 2011).

Future directions

Researchers are beginning to find evidence of interpersonal synchrony across a number of channels (Louwerse, Dale, Bard, & Jeuniaux, 2012). We believe that cross-channel

questions—for example, the relation between body movement and verbal turn-taking—are an essential next step for this research area and will promote a deeper understanding of the general and channel-specific mechanisms of synchrony. Although our efforts are presently in the area of bodily synchrony, we plan to incorporate other methods for studying additional channels of interpersonal synchrony.

Using new and established methods, we have endeavored to assemble a cost-effective and efficient methodology that facilitates research into multimodal questions. All items used in the setup are commercially available and highly regarded by reviewers on commercial Web sites (e.g., Amazon). As was noted above, conversations were recorded using a Canon Vixia HF M31 HD digital video camera, mounted on a Sunpak PlatinumPlus 6000PG tripod. To facilitate linguistic analyses, each interlocutor's audio was recorded on a separate audio channel (using an Audio-Technica ATR3350 Omnidirectional Condenser Lavalier Microphone), attached to an Azden CAM-3 On-Camcorder Mini Audio Mixer. The setup as described above costs less than \$800; however, researchers may readily substitute less expensive items (e.g., a webcam for the camcorder) as needed.

We believe that this setup and methodology are flexible enough to capture a number of modes of communication. For example, researchers interested in questions of linguistic alignment (e.g., priming; Brennan & Clark, 1996; Cleland & Pickering, 2003; Kousidis & Dorran, 2009; Niederhoffer & Pennebaker, 2002; Reitter, Moore, & Keller, 2010) will find the two-channel recording method amenable to their research (e.g., transcription; Kreuz & Riordan, 2011). Additionally, by combining the FDM with pre- or postinteraction questionnaires, researchers interested in affective synchrony (e.g., Chartrand & Bargh, 1999; Lakin & Chartrand, 2003; Miles et al., 2011; Sadler et al., 2009; Valdesolo & Desteno, 2011) may begin to investigate questions of affective alignment in conjunction with other channels of communication. By combining research into these and other channels, the field can better understand the functions of interpersonal synchrony. Further investigations into cross-channel questions will serve to complement the findings of early efforts in these issues (e.g., Louwerse et al., 2012).

Acknowledgements The authors wish to thank Sidney D'Mello for his advice on filters. We also wish to thank undergraduate research assistants Will Dunbar and John James for their help in hand-coding for the second validation study presented in [Appendix 3](#). This work was supported in part by NSF Grant HSD-0826825.

Appendix 1 AppleScript for automating image segmentation

```
(* Note: You must have first set the defaults on the image segmentation process before running this script. To do so, open iMovie and select "Export using QuickTime" under the "Share" menu. Select "Movie to Image Sequence," and set your preferences by clicking "Options." *)
```

```
repeat 1 times # once per video
  tell application "System Events"
    tell application "iMovie" to activate
      key code 125
      key code 120 using control down
      repeat 7 times
        key code 124
      end repeat
      repeat 12 times
        key code 125
      end repeat
      repeat 2 times
        key code 36
      end repeat
    end tell
    delay 420
  end repeat
```

Appendix 2 MATLAB code for simple frame-differencing method

```

% PROCESSING LOOP % PROCESSING LOOP % PROCESSING LOOP % PROCESSING LOOP %

% insert appropriate directory below
cd('directory');

% basic variables
h = gcf;
lag_size = 150;

% fetch images
imgpath = 'img*.jpg';
imgfiles = dir(imgpath);
disp(['Found ' int2str(length(imgfiles)) ' image files.'])

% create vectors for differenced image z-scores and L/R movement scores
image_z_diffs = [];
pLms = [];
pRms = [];

% begin loop through images
for j=2:length(imgfiles)
    disp(['Processing image: ' int2str(j) '.']);

    % prep the files
    file_name = imgfiles(j).name;
    image_2 = imread(file_name);
    file_name = imgfiles(j-1).name;
    image_1 = imread(file_name);

    % collapse images across color
    image_2 = mean(image_2,3);
    image_1 = mean(image_1,3);

    % turn images into pixel z-scores
    image_2 = (image_2 - mean(image_2(:)))./std(double(image_2(:)));
    image_1 = (image_1 - mean(image_1(:)))./std(double(image_1(:)));

    % difference, standardize, and store difference vectors
    image_diff = abs(image_2 - image_1);
    image_z_diffs = [image_z_diffs ; mean(image_diff(:))];

    % split images into L/R
    pLm = mean(mean(mean(image_diff(:,1:320,:)))); % change pixels as needed
to half image; first colon in image_diff designates y-axis range, second range
designates x-axis range
    pRm = mean(mean(mean(image_diff(:,321:end,:)))); % see above

    % store split vectors
    pLms = [pLms ; pLm];
    pRms = [pRms ; pRm];
end

% apply Butterworth filter to results
[bb,aa] = butter(2,.2);
pLms = filter(bb,aa,pLms);
pRms = filter(bb,aa,pRms);

% save workspace
save sample_FDM.mat;
disp('Frame-Differencing for Sample Dyad Complete.')
```

```

% CALCULATE CORRELATIONS % CALCULATE CORRELATIONS % CALCULATE CORRELATIONS %

% create matrix for correlations
dy_xcorrs = [];
disp('Creating Correlations for Sample Dyad.')

% cross-correlate and fill matrix
dy_xcorr = xcov(pLms,pRms,lag_size,'coeff');
dy_xcorrs = [dy_xcorrs dy_xcorr];
disp('Cross-Correlation for Sample Dyad Complete.')

% save workspace
save sample_FDM.mat
disp('MATLAB Workspace Saved.')

% GENERATE TEXT FILE % GENERATE TEXT FILE % GENERATE TEXT FILE %

% create csv file
delete('sample.FDM.csv');
data_out = fopen('sample.FDM.csv ','w');
disp('Text File Created.')

% fill the file with data
for corr=1:301
    % insert cross-correlation coefficient calculated above
    fprintf(data_out,'%d,',eval(['dy_xcorrs(' int2str(corr) ')']));

    % add time slice (i.e., where along +/- 3s each coefficient was derived)
    fprintf(data_out,'%d,',corr);
end

% close the data file
fclose(data_out);
disp('Text File Complete.');
```

Appendix 3 Two validation analyses of presented FDM

We wished to test that our frame-differencing method was doing what we conceptualized—capturing the synchrony of body movement during interaction. To do this, we completed two brief conceptual validation analyses, one using an artificial scenario and another using a small subset of the data analyzed in the experimental portion of the article.

Artificial scenario

Two people (the second author and another human, unfamiliar with the study) sat across from each other as described in the Method section. They moved in a variety of directions for 60 s and agreed to “attempt to synchronize movements together.” Movements involved a variety of pointing, nodding, and rhythmic motions, including the Y-M-C-A dance. Interspersed within these bouts of coupling were moments of nonmovement before engaging in the next bout of synchrony. The video was deliberately designed to produce synchrony in body motion.

When running the FDM (described in the main text of the article), the cross-correlation profile demonstrates




Fig. 5 Plot of cross-correlation coefficients derived from the FDM. Mean r is plotted along the y -axis; time lag in 125-ms segments is plotted along the x -axis. The spike around lag 0 indicates in-phase synchrony of movements between interlocutors

what is expected: a very sharp rise of r at a lag of 0 ($p < .0001$). The maximum r is much higher than our conversational sample, as would be expected from the artificial nature of the activity. Yet the shape of the function is the same, reflecting a significant drop in r as lag increases from 0 (see Fig. 5).

In order to test that the FDM values are based on *body movement* and not some other artifact, the authors *separately* analyzed the video in two different ways. The first author analyzed the video using FDM. The second author carried out a coarser second-by-second coding of the video using a 1–7 Likert scale. The scale was used to reflect, by the naked eye, how much overall body movement was present (at 1 Hz). This process and time scale are akin to holistic ratings of body movement described earlier in the article. This was done separately for the left and right persons in the video. Crucially, when one participant's movement was coded, it was done *blindly* to the movement of the other participant (i.e., only one half of the video was seen when coding). The coding was straightforward given the artificial nature of the task. The cross-correlation function for these human judgments matches closely with that obtained with the FDM (see Fig. 6).

Because the 8-Hz FDM-derived time series has considerably more data points than the 1-Hz hand code, we down-sampled the FDM time series to the size of the hand-coded body change values. We then compared them with a simple Pearson correlation and obtained a strong correlation, $r = .68$, $p < .0001$. The figure of the scatterplot for the movement of both people (blue = right, black = left) is shown in Fig. 7.




Fig. 6 Plot of cross-correlation coefficients derived from Likert ratings. Mean r is plotted along the y -axis; time lag in 1-s segments is plotted along the x -axis. The spike around lag 0 indicates in-phase synchrony of movements between interlocutors




Fig. 7 Plot of correlations between movement registered by the FDM and Likert ratings. The FDM-derived movements lie along the y -axis, and the Likert ratings lie along the x -axis. The blue circles indicate movement by the person on the right, while the black crosses indicate movement by the person on the left

Subset of experimental data

In order to verify that similar patterns hold when looking at experimental data, 2 min of two separate dyads' conversations were analyzed using the methods outlined above. Two undergraduate research assistants (blind to the study and results) were recruited to code the movement. Each assistant rated one of two 2-min subsets of participant dyads' conversations, chosen semirandomly. Both were instructed to rate the second-by-second (1 Hz) movement of each participant (again, separately and blind to the movement of the other participant), using a 1–7 Likert scale. The only guideline given to raters was to remain consistent in their subjective evaluations of the movement.

The down-sampled FDM time series of the two conversation subsets were compared with the holistic ratings using simple Pearson correlations. Again, we found a strong correlation for each, $r = .66$, $p < .001$, and $r = .67$, $p < .001$.

Summary

These validations are deliberately simple. The first demonstrates that, in deliberately synchronized video clips, synchrony produces a marked cross-correlation peak at a lag of 0 and that human judgment of the video corresponds with a separate analysis based on the FDM. The second confirms that, in experimental data, the FDM provides a measure of actual movement, rather than spurious co-occurring phenomena (e.g., light fluctuations). By executing these analyses as straightforwardly and simply as possible, we attempted to confirm our methodology's effectiveness with intuitive holistic ratings (e.g., Bernieri et al., 1988).

References

- Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, 59(4), 390–412.
- Baayen, R. H. (2008). *Analyzing linguistic data: A practical introduction to statistics using R*. Cambridge: Cambridge University Press.
- Bates, D. (2006). lmer, p-values, and all that. *The R-help archives*. Retrieved June 6, 2012, from <https://stat.ethz.ch/pipermail/r-help/2006-May/094765.html>
- Battersby, S. A., Lavelle, M., Healey, P. G., & McCabe, R. (2008). Analysing interaction: A comparison of 2D and 3D techniques. Paper presented at the Multimodal Corpora Workshop in the Sixth International Language Resources and Evaluation (LREC'08), Marrakech, Morocco.
- Bemieri, F. J., Davis, J. M., Rosenthal, R., & Knee, C. R. (1994). Interactional synchrony and rapport: Measuring synchrony in displays devoid of sound and facial affect. *Personality and Social Psychology Bulletin*, 20(3), 303–311.
- Bemieri, F. J., Reznick, J. S., & Rosenthal, R. (1988). Synchrony, pseudosynchrony, and dissynchrony: Measuring the entrainment process in mother-infant interactions. *Journal of Personality and Social Psychology*, 54(2), 243–253.
- Boker, S. M., Xu, M., Rotondo, J. L., & King, K. (2002). Windowed cross-correlation and peak picking for the analysis of variability in the association between behavioral time series. *Psychological Methods*, 7(3), 338–355.
- Brennan, S. E., & Clark, H. H. (1996). Conceptual pacts and lexical choice in conversation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22(6), 1482–1493.
- Caucci, G. (2011). *When I move, you move: Coordination in conversation*. Unpublished dissertation.
- Chartrand, T. L., & Bargh, J. A. (1999). The chameleon effect: The perception–behavior link and social interaction. *Journal of Personality and Social Psychology*, 76(6), 893–910.
- Cleland, A. A., & Pickering, M. J. (2003). The use of lexical and syntactic information in language production: Evidence from the priming of noun-phrase structure. *Journal of Memory and Language*, 49(2), 214–230.
- Condon, W. S., & Sander, L. W. (1974). Neonate movement is synchronized with adult speech: Interactional participation and language acquisition. *Science*, 183(4120), 99–101.
- Criss, M. M., Shaw, D. S., & Ingoldsby, E. M. (2003). Mother–son positive synchrony in middle childhood: Relation to antisocial behavior. *Social Development*, 12(3), 379–400.
- Giles, H., & Smith, P. (1979). Accommodation theory: Optimal levels of convergence. In H. Giles & R. St. Clair (Eds.), *Language and social psychology* (pp. 45–65). Oxford: Blackwell.
- Grammer, K., Kruck, K. B., & Magnusson, M. S. (1998). The courtship dance: Patterns of nonverbal synchronization in opposite-sex encounters. *Journal of Nonverbal Behavior*, 22(1), 3–29.
- Grammer, K., Honda, M., Jüette, A., & Schmitt, A. (1999). Fuzziness of nonverbal courtship communication unblurred by motion energy detection. *Journal of Personality and Social Psychology*, 77(3), 487–508.
- Kousidis, S., & Dorran, D. (2009). *Monitoring convergence of temporal features in spontaneous dialogue speech*. Paper presented at the First Young Researchers Workshop on Speech Technology, Dublin, Ireland.
- Kreuz, R. J., & Riordan, M. A. (2011). The transcription of face-to-face interaction. In W. Bublitz & N. R. Norrick (Eds.), *Handbooks of pragmatics* (Vol. 1, pp. 657–680). Berlin: De Gruyter Mouton.
- Kupper, Z., Ramseyer, F., Hoffmann, H., Kalbermatten, S., & Tschacher, W. (2010). Video-based quantification of body movement during social interaction indicates the severity of negative symptoms in patients with schizophrenia. *Schizophrenia Research*, 121(1–3), 90–100.
- Lakin, J. L., & Chartrand, T. L. (2003). Using nonconscious behavioral mimicry to create affiliation and rapport. *Psychological Science*, 14(4), 334–339.
- Louwerse, M. M., Dale, R., Bard, E. G., & Jeuniaux, P. (2012). Behavior matching in multimodal communication is synchronized. *Cognitive Science*.
- Lu, P., & Huenerfauth, M. (2010). Collecting a motion-capture corpus of American Sign Language for data-driven generation research. *Proceedings of the NAACL HLT 2010 Workshop on Speech and Language Processing for Assistive Technologies*, 89–97.
- Lu, S., Tsechpenakis, G., Metaxas, D. N., Jensen, M. L., & Kruse, J. (2005). *Blob analysis of the head and hands: A method for deception detection*. Paper presented at the annual Hawaii International Conference on System Sciences (HICSS '05), Hawaii.
- Meservy, T. O., Jensen, M. L., Kruse, J., Burgoon, J. K., Nunamaker, J. F., Jr., Twitchell, D. P., Tsechpenakis, G., et al. (2005). Deception detection through automatic, unobtrusive analysis of nonverbal behavior. *Intelligent Systems, IEEE*, 20(5), 36–43.
- Miles, L. K., Lumsden, J., Richardson, M. J., & Macrae, C. N. (2011). Do birds of a feather move together? Group membership and behavioral synchrony. *Experimental Brain Research*, 211(3–4), 495–503.
- Nagaoka, C., & Komori, M. (2008). Body movement synchrony in psychotherapeutic counseling: A study using the video-based quantification method. *IEICE Transactions on Information and Systems*, 91(6), 1634–1640.
- Niederhoffer, K. G., & Pennebaker, J. W. (2002). Linguistic style matching in social interaction. *Journal of Language and Social Psychology*, 21(4), 337–360.
- Pickering, M. J., & Garrod, S. (2004). Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences*, 27(02), 169–190.
- Ramseyer, F., & Tschacher, W. (2008). Synchrony in dyadic psychotherapy sessions. In S. Vrobel, O. E. Roessler, & T. Marks-Tarlow (Eds.), *Simultaneity: Temporal structures and observer perspectives* (pp. 329–347). Singapore: World Scientific.
- Ramseyer, F., & Tschacher, W. (2011). Nonverbal synchrony in psychotherapy: Coordinated body-movement reflects relationship quality and outcome. *Journal of Consulting and Clinical Psychology*, 79(3), 284–295.
- Reitter, D., Moore, J. D., & Keller, F. (2010). *Priming of syntactic rules in task-oriented dialogue and spontaneous conversation*. Paper presented at the annual conference of the Cognitive Science Society, Vancouver, BC.
- Richardson, D. C., & Dale, R. (2005). Looking to understand: The coupling between speakers' and listeners' eye movements and its relationship to discourse comprehension. *Cognitive Science*, 29(6), 1045–1060.
- Richardson, D. C., Dale, R., & Tomlinson, J. M. (2009). Conversation, gaze coordination, and beliefs about visual context. *Cognitive Science*, 33(8), 1468–1482.
- Richardson, M. J., Marsh, K. L., Isenhower, R. W., Goodman, J. R. L., & Schmidt, R. C. (2007). Rocking together: Dynamics of intentional and unintentional interpersonal coordination. *Human Movement Science*, 26(6), 867–891.
- Sadler, P., Ethier, N., Gunn, G. R., Duong, D., & Woody, E. (2009). Are we on the same wavelength? Interpersonal complementarity as shared cyclical patterns during interactions. *Journal of Personality and Social Psychology*, 97(6), 1005–1020.
- Schmidt, R. C., Carello, C., & Turvey, M. T. (1990). Phase transitions and critical fluctuations in the visual coordination of rhythmic

- movements between people. *Journal of Experimental Psychology*, 16(2), 227–247.
- Schmidt, R. C., Morr, S., Fitzpatrick, P., & Richardson, M. J. (in press). Measuring the dynamics of interactional synchrony. *Journal of Nonverbal Behavior*.
- Shockley, K., Baker, A. A., Richardson, M. J., & Fowler, C. A. (2007). Articulatory constraints on interpersonal postural coordination. *Journal of Experimental Psychology: Human Perception and Performance*, 33(1), 201–208.
- Shockley, K., Santana, M. V., & Fowler, C. A. (2003). Mutual interpersonal postural constraints are involved in cooperative conversation. *Journal of Experimental Psychology*, 29(2), 326–332.
- Skuban, E. M., Shaw, D. S., Gardner, F., Supplee, L. H., & Nichols, S. R. (2006). The correlates of dyadic synchrony in high-risk, low-income toddler boys. *Infant Behavior and Development*, 29(3), 423–434.
- Valdesolo, P., & DeSteno, D. (2011). Synchrony and the social tuning of compassion. *Emotion*, 11(2), 262–266.
- Welch, G., & Foxlin, E. (2002). Motion tracking: No silver bullet, but a respectable arsenal. *Computer Graphics and Applications, IEEE*, 22(6), 24–38.