Group entrainment experiment: "Rhythm Battle"

Introduction

The purpose of this task is to simultaneously test entrainment *and* resisting entrainment in a large group of participants. The large group is divided into two teams who are then trying to maximise entrainment and tempo stability **within** the team, while minimising entrainment and influence **between** them.

This experiment was inspired by a study on Brazilian *Congado*, a tradition where marching bands parade around the city, playing and dancing. When they meet each other, they do a greeting ritual, where they attempt to maintain their own tempo and try not to be influenced by the other group. This tradition has been studied and these meetings have been analysed by Lucas, Clayton & Leante (2011). They showed that entrainment is such a strong, automatic tendency that it is very difficult to resist. Groups would often entrain fully, and at least adjust to each other. Even in cases where the groups seemed to manage to resist entrainment, they could have just entrained in a more complex way, such as 2/3 (one group has 2 beats in the time the other has 3).

We have been piloting a protocol that can be run e.g. as part of a movement improvisation workshop. Here, we will explain how the protocol works and show an example of how the data we collect could be analysed.

Protocol

Participants are divided into two teams. In each team, a facilitator or group leader has a metronome. The two metronomes are set to different (non-harmonic) tempi, e.g. 85 and 100 BPM. The game has four stages (figure 1).
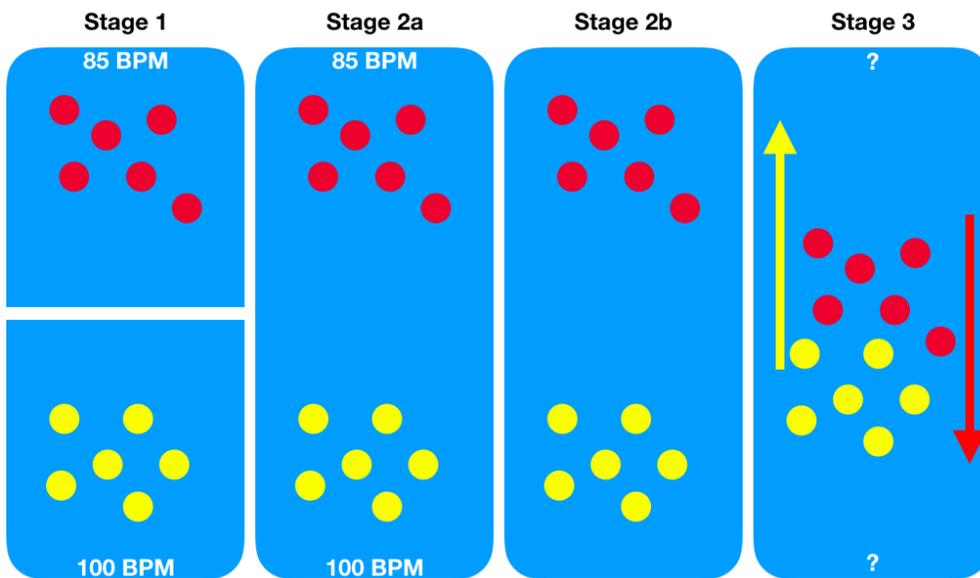

Figure 1. Overview of the protocol.

Stage 1:

- the two teams are in different rooms, or in our case, the room is divided into two halves using a movable wall.
- facilitators have metronomes (we used a mobile app Metronome Beats by Stonekick), and their tempi are set to 85 or 100 BPM.
- Teams are tasked to create a group rhythm pattern that they can move to. They can clap their hands, stomp feet, snap fingers, tap their chest etc. The team effectively becomes a moving, walking drum kit, performing their rhythm pattern. Participants are therefore not limited to just tapping /clapping etc. on the beat, but can do so also off-beat, half tempo, double tempo etc., as long as the whole team maintains their original tempo.
- In stage 1, there is no interaction between the teams.

Stage 2a:

- The teams are brought into the same space (or the divider is removed). This starts the *between teams* interaction.
- The metronomes are still on, allowing participants to use that as a reference and get used to performing their rhythm under the distraction of the other team.

Stage 2b:

- metronomes are turned off, and teams try to maintain their group cohesion and resist entraining with the other team.

Stage 3:

- the teams start moving, as groups. They can be given a movement objective, or the movement pattern in the room can be improvised. The movement objective in our case was to reach the other end of the room, so the place where the other team starts from.
- From analysis point of view, this is the most interesting stage; mainly, what tempi the teams end up with.

Based on our pilots, the first stage should be given as much time as the teams need to get ready. This makes the rest of the game more fun, as participants feel more comfortable and confident with their rhythms. Stage 2 can be quite short. Stage 3 (and the game) can be stopped when both teams have achieved their movement goals, or when one team has clearly "lost" and the two rhythms have assimilated into one.

While the movement tempo in stages 1 and 2a should be the as the team metronome, what the tempo develops to be in stages 2b and 3 is the main question for the analysis.

Pilot data

Here we demonstrate some of the analyses that can be useful in analysing group synchrony. First, we show some raw data from one performance, to give a taste of the type of data that the game produces. In this game, there were 10 participants that had five accelerometers each; on their left and right wrists and ankles, and on their chests. In the second analysis, we look at three other games. In that analysis, we look to analyse movement tempi of the two teams and how they change from one stage to another. Finally, we demonstrate how a network analysis of these games could progress.

Data processing

The acceleration data from the 50 sensors is synchronised, and re-sampled to have a constant sample rate of 100 Hz. The original data consists of accelerations along three axes, with positive and negative values depending on the direction of the acceleration. However, as the accelerometers change their orientation during the recording (e.g. a wrist sensor's X-axis can point to any direction depending on the posture of the participant), these directions are not relevant. Therefore we compute the absolute acceleration, that combines the three axes and only contains positive values. Thus, the data represents acceleration to any direction, or acceleration along the trajectory of movement. This data is then low-pass filtered so that we can get rid of the high-frequency components that are mostly noise, e.g. generated because the accelerometers or their straps can move a bit etc. As in this analysis, we want to focus on gross movements and their synchrony, we set the threshold of the filter quite low, to 2 Hz, this should emphasise the movements that occur at the beat-level at the expense of faster movements and jitters, that might be interesting for other analyses.

Amount of movement

A plot of the raw accelerometer data from each phase shows that the amount of movement increases in every step as is expected. In the second phase, participants started performing their rhythm louder, which required larger accelerations. In the third phase, participants were moving around the room, which of course resulted in yet more movement (larger accelerations). (Click on the figures for full-soze image).
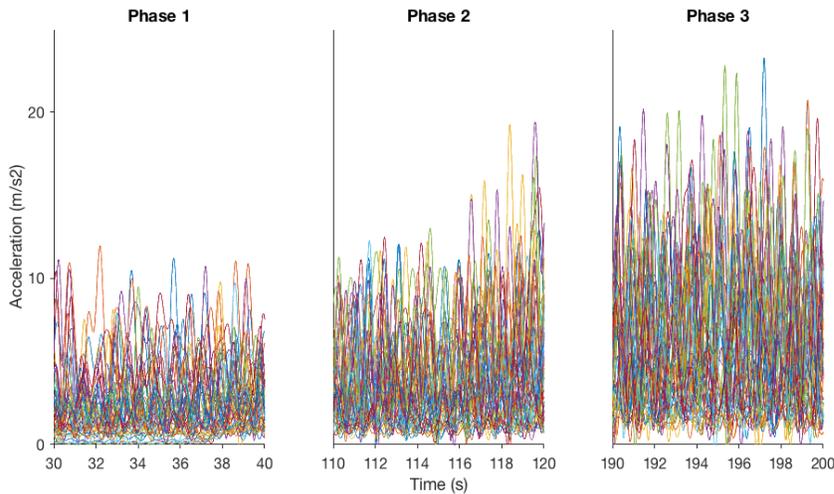

Figure 2. Acceleration data from all sensors, in the three sample windows.

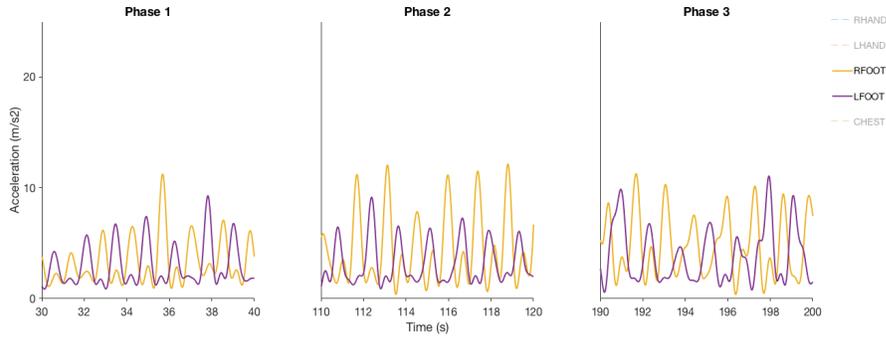Let us then take a look at an individual participant, and how their five sensors are related.

Figure 3. Individual participant; feet

In phases one and two, participants remain still and huddled in a circle. However, many participants were using their feet for the rhythms they were creating. Figure 3 shows data from one participant. They do a stepping motion, with acceleration peaks alternating between the two sensors. In the third phase, they are walking around the room. This seems to disturb the pattern somewhat, as especially the left foot is less consistent than in the second phase. Banking the basic beat to large body parts and large muscles is a good idea. Here we can also easily estimate the tempo from the figure. In the first and second phase, there are 7 peaks per foot, in the ten-second window. This indicates a tempo of 14 * 6 = 84 beats per minute. In the third phase, the tempo has increased somewhat, as there are now clearly 15 peaks, indicating a tempo of 90 bpm. This is of course a very rough measure, kin to estimating your HR by counting heartbeats for ten seconds and multiplying this by 6. Just as we have accurate HR-monitors, we can get a more accurate tempo estimates by using autocorrelation-based periodicity analysis.
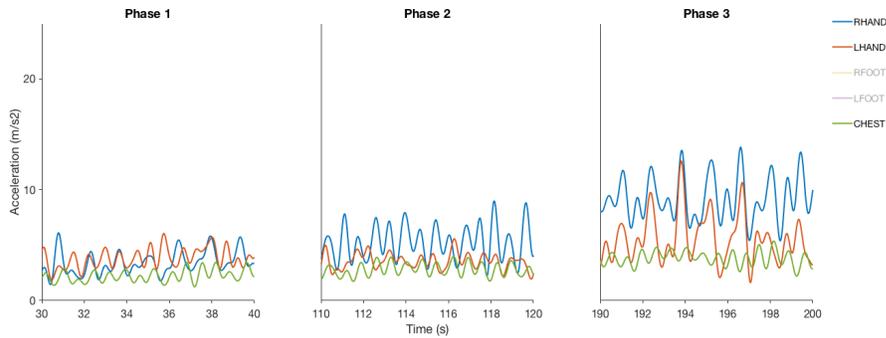

Figure 4: Hands and chest

Looking at the movements of hands and chests, we can see that much of the increased movement in phase three comes actually from increased hand movements, at least for this participant. While different hands and feet tend to alternate in their movements, the chest seems to be the most consistent, and clearly periodic signal from the beginning to the end. It seems to capture the main beat and ignore the "frills". This is very clear for this participant but tends to be the case more generally. Therefore for a simple synchronisation analysis, using data from chest sensors is probably a good idea.
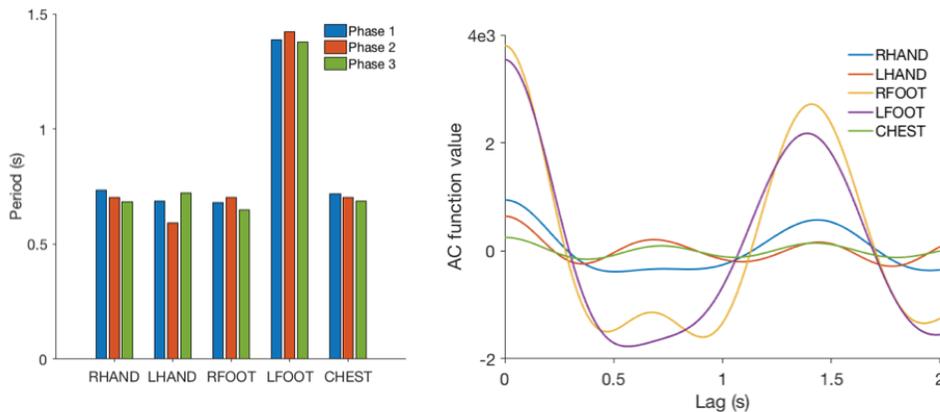
Periodicity of movement

Figure 5. Periodicity as measured from individual sensors of one participant

Using the autocorrelation method to estimate periodicities, we see that different limbs move in slightly different tempi. In the left panel on figure 5, the average periodicities for the three phases, for the sensors of one participant are plotted. These are the period lengths associated with the first peaks in the autocorrelation function, or the fastest rate for each limb. We can see that e.g. for the right hand, the average tempo gets slightly faster in each phase, same goes for chest. One option for From the right panel, where the autocorrelation functions themselves are plotted, we can see that the "clearest" periodicities (highest peaks) for this participant are mostly at double the metronome tempo, or at the 1.4 second area. The feet give the clearest periodical signals overall, and as they move in an alternating pattern (see figure 4), together they form a rhythm with one step / one beat.

Synchronisation

When analysing "togetherness" or "matching", or "adaptation", we can simply look how close to each other's tempo people are (e.g. standard deviation of the tempi of the team). However, if the accelerometers are synchronised well enough, we can use a more complete definition of synchronisation, according to which, two systems converge in phase and/or period when they sync. In other words, we can look at phase synchrony.

First of all, as we concluded above, the chest markers seem to capture the different limb patterns and sum them up into a periodic signal that is quite stable and at the frequency of the metronome. Figure 6 shows the chest accelerations of all participants, divided to the three phases, and this time also the participants are divided into the two teams or groups.
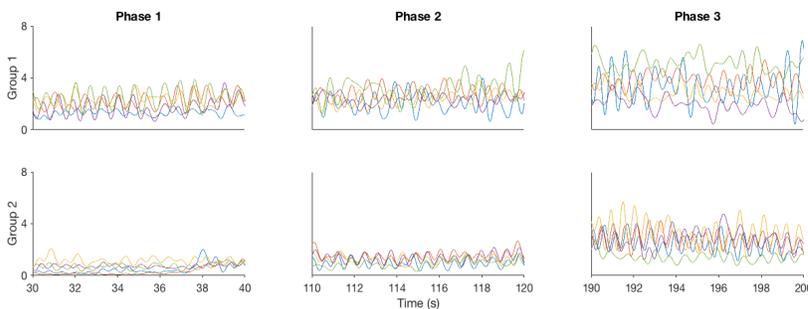

Figure 6. Chest sensor accelerations for teams 1 and 2 (the "competing teams"), samples from the three phases of the game

If we only have two systems (two dancers in this case), we could take the acceleration of their chests, for example, and use Hilbert transform to obtain their instantaneous phases, and then get their phase difference through subtraction and then see whether it stays stable or not. If the two are synchronised, they should maintain a stable phase difference (which does not have to be zero). But, as we have ten players in two teams, we need a different method. We can use the Kuramoto model, which describes collective synchronisation of a large number of independent oscillators, spontaneously synchronising their period and phase. The model explains how the oscillators sync, and we can also use it to quantify the level at which a group of oscillators have synchronised. This can be summarised in an order parameter, an index that ranges from 0 to 1, with 0 indicating disarray and 1 perfect sync. We calculated the order parameter for team 1, team 2 (in figure legend "group 1" and "group 2"), and everyone together. Figure 7 shows the average OP's for each of the three phases.
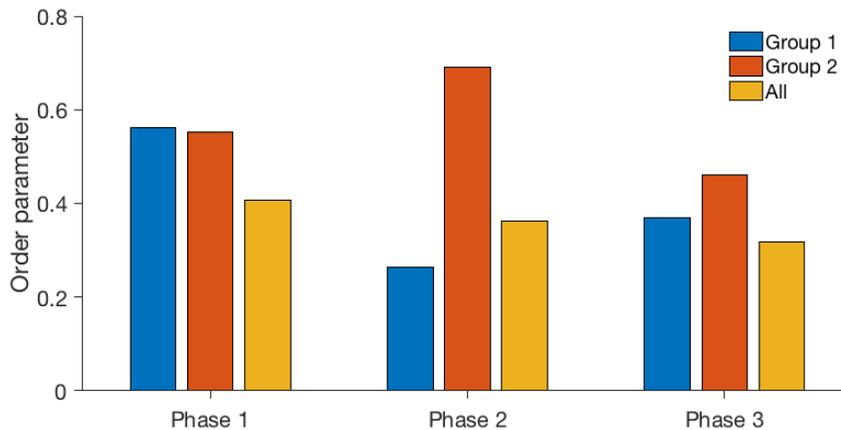

Figure 7. Average order parameters for the different phases of the game

It seems that both teams reached a similar sync level in the first phase, but team 1 somehow lost it when teams started communicating with each other. In the third stage, the teams get closer in terms of sync, but both are weaker than in the first stage. But, let's look at how the OP evolves in time.
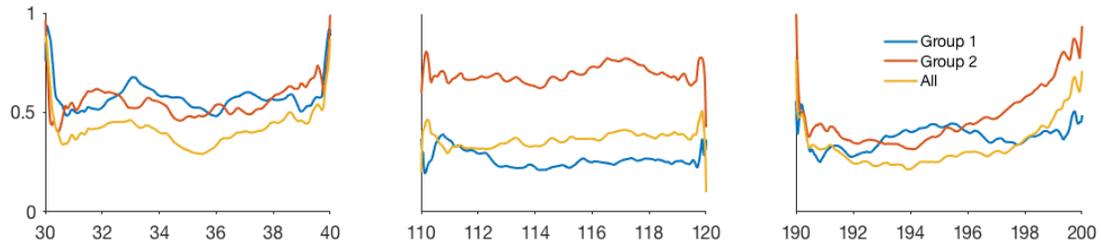
Figure 8. Order parameter evolution

The data is smoothed with a long averaging window, as the actual parameter is quite noisy, as the participants' slightly different tempi make the phase relationships somewhat unstable. The peak values at the ends of each graph are a result of the smoothing. But, these smoothed curves show some interesting features. In the first phase, the fact that participants are moving very little means that the chest signals are small, this might contribute to the jittery order parameters. In the second phase, both teams seem to have a more steady state, although as we already saw in the bar graph, team 1 seems to have a difficult time, while team 2 has a very good run. In the third phase, things evolve more gradually. While team 1 improves in the middle, it seems that team 2 then gets better and better, overtaking team 1. And as the overall order also grows towards the end, following the path of team 2, it might be that team 2 have managed to pull at least some in team 1 to sync with them. Definitely some team 1 members seem to correlate highly with team 2 especially in phase 3.

Analysing the battling teams as a network

For this analysis example, we use data from another pilot. Another way to look at what goes on in the game is to visualise the teams as networks. Here, we visualise the average tempi of each participant, and the tempo differences between each pair of participants (both within and between teams) to form our graph. First, lets take a look at the overview of how this game progressed. In Figure 9 we have plotted the mean tempi in two teams for the different phases of the game. Based on these, we can say that the fast tempo team "won" the game–they kept their average tempo close to where they started, whereas the slow tempo team shifted to the fast tempo already after the metronomes were switched off, even before the teams started moving (figure 9).
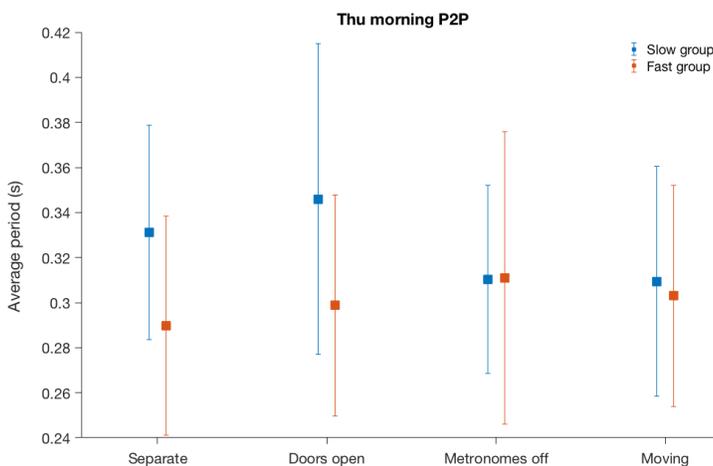

Figure 9. Average tempi in the different phases of the game. These means are calculated from 20-second segments from each phase. Error bars represent the means of individual standard deviations of inter-beat intervals in these segments.

In our network, each participant is a node. The edges (connections between the nodes) come from We calculated next the tempo differences for each pair of participants, and then converted these into edge weights for the graph - the closer two participants are in tempo, the stronger their connection. Also, if the tempo difference gets too large, the two are not connected at all. This way we get a more interesting graph than the fully connected one that we would get if everyone was connected to everyone else by default. Figure 10 shows the networks for this group, in phase 1 and phase 3. (We did have two teams of equal size in the game, however due to technical reasons we only have data from 5 participants in one team and 3 in the other. We used chest accelerometer data for this analysis.) We visualised the networks in Gephi, using a force-algorithm that pushes weakly connected nodes away from each other and pulls strongly connected nodes closer.
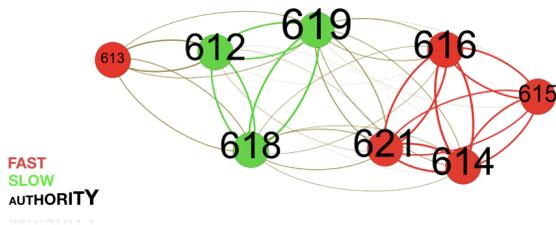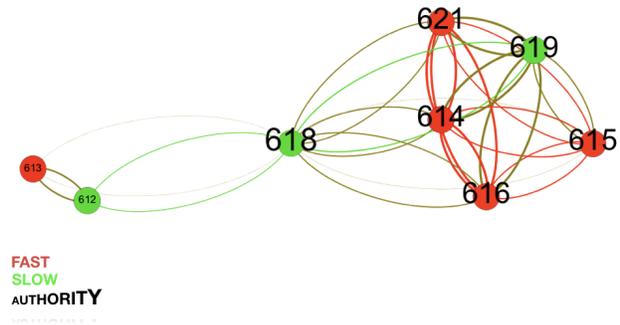
Figure 10. Networks of the game in the beginning (left) and the end (right) of the game. Node colours indicate team membership. The thickness of the edges reflect how close the two are in tempo. The font size of the participant number reflects the "authority" of the participant, or how influential they are in the network (font sizes are not compareable across the left and the right graphs, only within each).

In the beginning (figure 10, left) the participants are in two groups that are relatively weakly interconnected, apart from participant 613 who seems to be an outlier here. Perhaps their pattern was such that the chest accelerometer did not pick it up very clearly, or they were just not in the tempo the rest of their team was. Apart from this one outlier, the graph looks exactly as we'd expect for this first phase: two groups, each highly coherent within, and with a clear difference to the other group.

Figure 10 right shows the network at the end of the game, when they have already been moving for a while. Our original outlier is still an outlier, but now has gained another, from the slow team. These two are somewhat close to each other, but almost not connected to anyone else. The rest of the fast team is still together, and now crucially they've managed to attract one of the other slow team participants (619) to be with them. 619 was also the most authoritative member of the slow team in the beginning, so that had a big effect on the slow team. 618 has also drifted closer to the fast group. Also the visual analysis of the graph supports the view from looking at the mean tempi: the fast team was stronger.

Rhythm battle group and subjective assessment of closeness

In this pilot, we also asked participants to rate how close they felt to the other participants. These ratings were collected at the very end of the day, and the ratings were not specifically focused on the group synchronisation task, but generally to the whole day, that consisted of a number of experiments, an improv glass and various rhythm games alone and in groups.

We used a version of the Inclusion of Others in the Self -measure (Aron, Aron & Smollan, 1992). Pooling together ratings from three groups of participants (three games), we could compare whether on average, members in the same team were rated closer than members in the opposing team. Previous research points out to the social benefits of synchronisation (e.g. Hove & Risen, 2009; Rabinowitch, et al. 2013), and thus our hypothesis would be that team membership in the game would have an effect here although the participants experience together a longer program of various activities.

Indeed, our analysis indicates that participants rated those participants closer that they shared a team with in the rhythm battle. This difference was highly significant (figure 11).
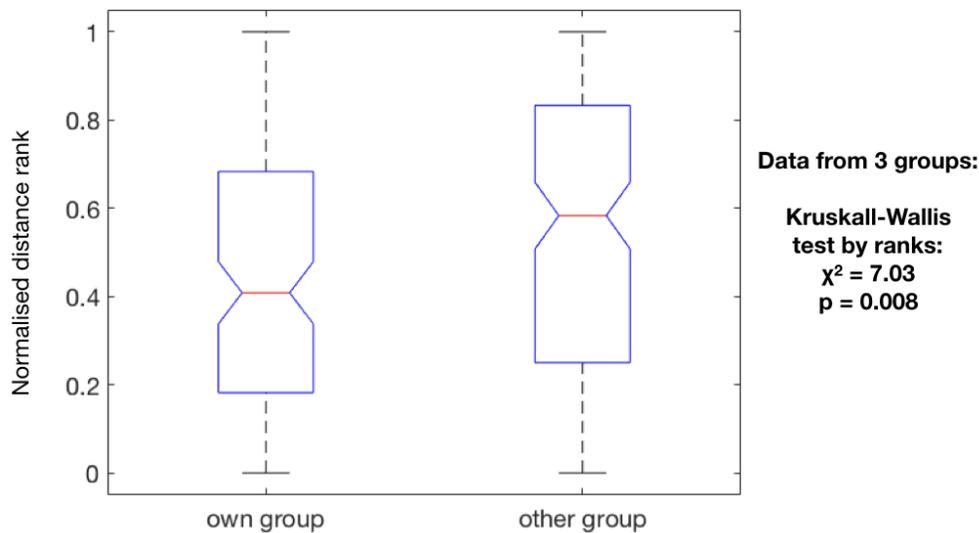
Figure 11. Comparison of IOS ratings between those in the same team in the rhythm battle vs. those not in the same team.

Conclusions

In this example, we have walked through a possible group synchronisation analysis with some accelerometer data recorded from two groups playing a sync - resist sync game. Using the Kuramoto model and the order parameter as a group synchronisation metric yields interesting information about the evolution of group synchrony. The example analysis looks at just three short windows of the original data, and is just from one group, one game. However, the whole performance could be analysed using similar methods, and for example the average OP's could be compared across different games or different groups, to compare performance within (as in pre- and post-intervention) or between participant groups.

We also presented a network-based analysis of the game. Constructing a graph with participants as nodes and their tempo differences as edges produces a visualisation that allows for analysis of the dynamics during the game and how the "winning" team won.

The team membership in the rhythm battle also was found to have an effect on how close participants felt at the end of the day. To see that the IOS measure seems sensitive to synchronisation is very interesting, it would be in line with previous studies that have used different measures, and usually dyadic or at least shorter interventions.

We are currently pursuing the network analyses further, looking e.g. dynamic models where we could visualise the evolving tempo differences more continuously. Also, the protocol is constantly being refined. Furthermore, we can use the other rhythm games (solo and triad finger tapping) as background measures, to e.g. identify if outliers in the game are ones that have lowest coherence also in solo tapping or whether their unsynchrony is just related to the group context. Also, we might be able to group people based on their tempo preferences (those with faster spontaneous tempi make up the faster team), instead of the grouping being (pseudo)random as it is currently.

References:

Acebrón JA, Bonilla LL, Vicente CJP, Ritort F, Spigler R. (2005). The Kuramoto model: A simple paradigm for synchronization phenomena. Reviews of modern physics. 77(1):137–185.

Aron, A., Aron, E. N., & Smollan, D. (1992). Inclusion of Other in the Self Scale and the structure of interpersonal closeness. Journal of Personality and Social Psychology, 63(4), 596.

Lucas, G., Clayton, M., & Leante, L. (2011). Inter-group entrainment in Afro-Brazilian Congado ritual. Empirical Musicology Review, 6(2), 75–102.

Oullier O, de Guzman GC, Jantzen KJ, Lagarde J, Kelso JAS. (2008). Social coordination dynamics: Measuring human bonding. Social neuroscience. 3(2): 178–192.

Spiro, N. & Himberg, T. (2012). Empathy and Resisting Entrainment: Mapping the dimensions of pairwise rhythmic interaction. 16th Annual Symposium for Music Scholars in Finland, Jyväskylä, Finland.