Learning multiple distinct patterns in a task of rhythmic interlimb coordination

Hendrik Reimann

August 31, 2007

1 Introduction

A basic model for learning new patterns in a task of performing rhythmic movement of two limbs with same frequency, but shifted in phase by a set amount, the relative phase ϕ, was described in the report *"A dynamic field model of learning new patterns in rhythmic interlimb coordination"*, from May 29th, 2007. The model consists of three separate, interconnected parts, command, coordination and feedback, forming a closed loop. The three layered command field develops self-sustained peaks, which change the coordination dynamic, creating an attractor at the activated site. The feedback field compares the produced to the desired results, and merges that information with the memory layer from the command field to create corrective input.

Learning takes place in the form of preactivation of field sites, or preshape. Whenever a field site supports an active peak, preshape is increased at that sites, thus forming preshape peaks at those sites that tend to produce the best results, which in turn support the quick formation of peaks at those sites in subsequent trials, and help stabilizing these peaks by providing additional input.

The model gave a good description of the process of learning to produce a single new relative phase. The predictions of the model were qualitatively consistent with the experiment results in terms of mean produced relative phase, and stability.

In the state it is described here, the model has been enhanced in two major ways. First, the dynamic feedback loop was closed. In the earlier version, feedback was only given in large intervals, between trials. Now feedback is given on a similar, but slower timeframe as the command and coordination dynamics run. While seemingly a cosmetic change, the new model shows far nicer behaviour in several aspects, as changes that were previously limited to the switch between trials can now occur at different times and with varying frequency.

Second, the model was extended to account for the learning and reproduction of multiple different relative phases. In tackling that problem, the first question was how those multiple relative phases should be encoded in memory. New relative phases are encoded in the form of peaks in a preshape field in the single pattern model. Now a new peak for a new to-be-learned relative phase can either build up in the same field, or live in a completely new field. Taken further, the first answer would produce a single field with a landscape of multiple different, pronounced preshape peaks at different sites. The second answer would lead to a bundle of different preshape fields, each with a single peak.

Obviously adding a new field for a new to-be-learned relative phase would need a sophisticated mechanism for deciding when something new is being learned. Similarly, when asked to reproduce a given relative phase, there must be a way to decide which, if any, of the preshape patterns fits the task. While a mechanism such as this is clearly within the scope of a biologically motivated, competitive attractor dynamic, designing such a mechanism seems like a much more general problem than the one at hand. Reaching a solution with the relatively simple means of storing all preshape in one field seemed like a cleaner, more to the point solution, and was chosen for the present model and simulator.

2 The Model

Only those parts of the model that were changed will be described here. As a reference I reproduce the model overview graph (Fig. 1), for a detailed description of the model, see the earlier report.

2.1 Feedback

In the old model, collection of knowlege of results *(kr)* information and its application to initiate change were implemented in a very makeshift, straightforward way by creating a histogram of the error in the performed relative phase during trial and merging it with the command field working memory once after each trial, to create input for the next trial. This in a way reflected the specific setup of the fingertapping experiment (described in e.g. [SZK92]), but otherwise restricted the dynamics of the model. Also from a biological viewpoint, it makes limited sense to have error correction only in intervals between trials. When trying to produce rhythmic movements in more natural settings as playing a music instrument, one receives and applies feedback directly, not in large intervals.

The feedback cycle was thus modified to a full closed loop on each timestep. The error of the performed relative phase was measured, and this measurement was used to create an excitative input to the kr-field. Now, when an error value persists for an extended period of time, the continued

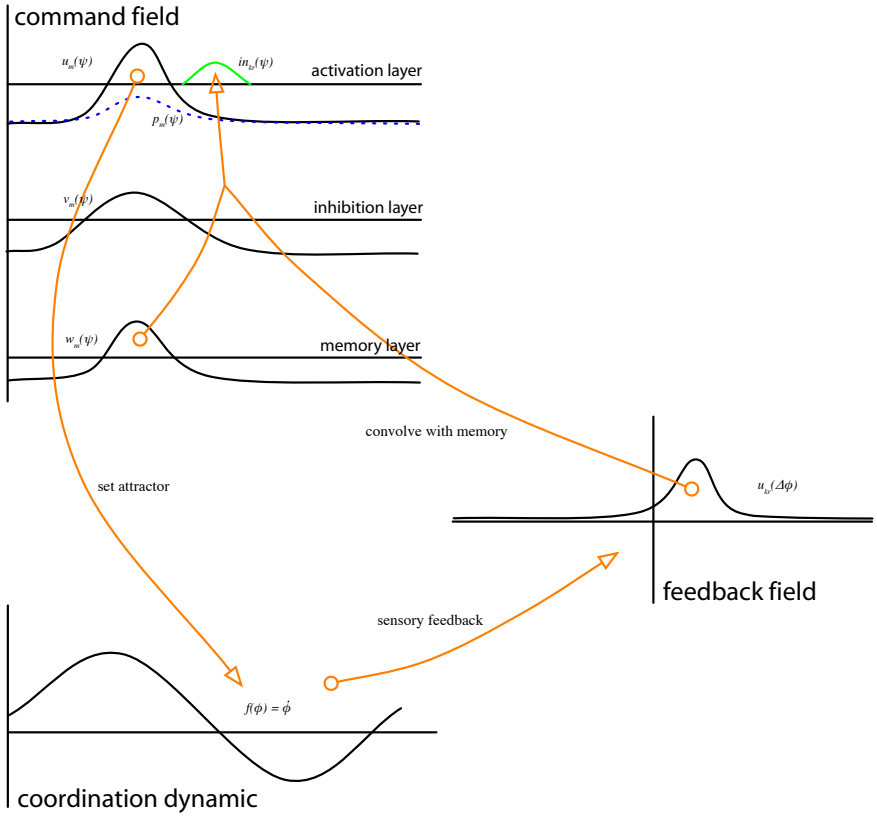

Figure 1: Overview of the closed loop model. A peak in the activation layer of the command field induces a sinoid attractor in the coordination dynamic. The systematic error of the attractor solution is shown in the feedback curve, which together with the memory layer of the command field produces new input (green). Any super-threshold activation in the command field leads to preshape at the activated sites (blue).

excitation will create a peak in the *kr*-field. As before, this peak is merged with the working memory layer of the command field via a convolution, to create a corrective input to the command field's activation layer.

To reflect the fact that the feedback part is still somewhat detached from the command-coordination side of the model, relying on sensory information about the effects of the initiated actions, the time scale of the *kr*-field was chosen to be much slower than that of the command field. This created the problem of "old" peaks in the *kr*-field not disappearing for an extended period of time, up to a point where in addition to the active, full blown peak, the *kr*-field would hold the remnants of 3-4 dying peaks, slowly disappearing, at sites that were active before. When a new peak had already formed in the

activation and then memory layer of the command field, the leftover peaks in the kr-field would be merged with the command memory field just as the legitimate peak, thus creating totally unwarranted inputs to the command activation field.

In order to solve this problem, an output nonlinearity was added to the kr-field, in the usual sigmoid form. This allowed for a strongly negative resting level in the kr-field, which meant the peaks not supported by input anymore would decay faster. Once under the sigmoid threshold, an old peak's influence quickly becomes insignificant. The negative resting level also made sure that short, transient inputs do not reach over into the command field. Only activities with a minimum of stability in the kr-field can arouse a corrective response in the command activation. The sigmoid thus acts as a filter towards noise as well. Finally, this new version of the feedback part of the model is biologically much more plausible than the old histogram realization.

The mathematical form was chosen as

$$\tau_{kr}\dot{u}_{kr} = -u_{kr} + h_{kr} + c_{fb}in_{fb},$$

with in_{fb} being the input from the sensory feedback, in the form of a gaussian peaked at the measured error in relative phase, $\psi_{tbl} - \phi$.

2.2 Command

The command field was subject to some changes and retuning. The new version of the feedback field made it possible to get rid of a workaround solution to problems that arose from the suboptimal implementation of the kr-field.

Activity in the memory layer w_m was more closely linked to peaks in the activation layer, excitation from u_m to w_m was increased, while lateral excitation in w_m was reduced. The memory layer is mirroring the activation layer, with a delay, due to different time scales. Peaks in w_m now depend on input from u_m, in the absence of which they collapse. As in the other two layers, activity in the memory layer fades between trials, while in the old version a sustained memory peak was needed in w_m, to be merged with the kr-field results in a convolution, creating the new input to u_m. This is not necessary with the new, dynamical feedback loop.

The mexican hat shape of the kernel for lateral interaction in w_m has thus become unnecessary as well, as it was introduced to prevent the explosion of peaks in the memory layer in absence of inhibition from activity in the other layers. It was replaced with a simple gaussian kernel, the same shape as all other interaction kernels in the command field.

2.3 Preshape

The preshape mechanism was subject to substantial change, to allow for the parallel coexistence of multiple peaks at different sites in the same field. In the former version of the simulator, whenever there was activity at a given field site, preshape was increased at the corresponding site in the preshape field. At the same time, preshape was reduced at all sites currently inactive in the command field, although the reduction was set to be slower than the increase. This enabled initially unstable systems to form a clean peak after a row of trials, by slowly flattening out the preshape traces from erraneous activity, while building up the successful solution.

This is a very mechanistic approach, and obviously not useful for creating a landscape of multiple distinct peaks for different relative phases in the preshape field. To create a pattern for preshape accumulation that was both fitting the problem at hand, and biologically more plausible than the old one, the leading idea was to localize the influence of the field activity.

The preshape field p_m is seen as a separate, but isomorphic, dynamic field parallel to the command activation field u_m. Connections between the two fields run in both directions. Excitation from the preshape to the command field activation layer is straightforward between field sites, a high preshape level leads to a high input into the command field at the corresponding site. The connection in the other direction was localized, from a mix between very narrow excitative, and very broad inhibitive influence in the earlier model. Activity in a command field site is now understood to increase the level of the preshape neurons in corresponding and adjacent field sites, with the influence waning as distance increases. Preshape at sites further away than a certain threshold is inhibited, while field sites still further away are not affected at all.

For an activated command field site, the sites in the first region are those that will most likely lead to a peak supporting a successful solution. The field sites in the second region are still close enough that command field activity in them would influence the original field site, but in an unsupportive way, leading to a too broad or unclear peak, so they are inhibited. Field sites still further away barely affect the original site, so preshape at these sites is left unchanged.

In the model, this is realized by using a mexican hat style kernel for the influence from u_m to p_m. To prevent the preshape field from exploding, a coefficient function γ is added, which starts low, then quickly reaches a maximum of 1, and declines to 0 again as a threshold is approached. This leads to the preshape rate of change being reduced for low values, so previously inactive sites will take a bit longer to develop significant activation. When the preshape reaches its saturation threshold, the rate of change will again be reduced to zero. The equations are as follows:

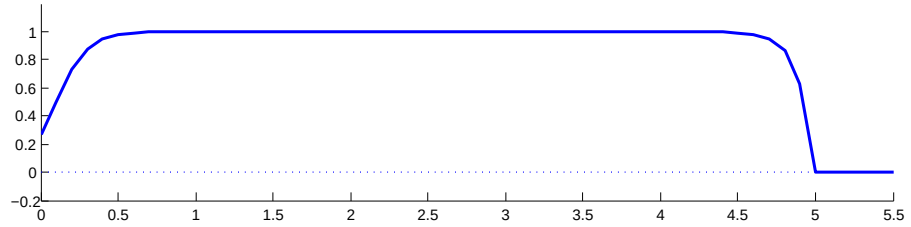

Figure 2: The coefficient function γ for the connection from the command to the preshape field.

$$\tau_{pre}\dot{p}_m(\psi) = \gamma(p_m(\psi)) \int k_{pre}(\psi - \psi')\sigma(u_m(\psi))d\psi',$$

with k_{pre} being the mexican hat style kernel function, and

$$g = \frac{1}{1 + e^{(-\beta(x-x_1))}} - \frac{1}{e^{(-\beta(x-x_2)),}}$$

$$\gamma = \begin{cases} g & : g \geq 0 \\ 0 & : g < 0. \end{cases}$$

While the command field can be interpreted as a biological neural field, the preshape field should be seen as a more abstract entity. As it is implemented here, the name preshape is not used in the traditional way. In the present model, p_m is a separate field, technically independent of the command field, while *preshape* usually refers to field sites being pre-activated as a result of repeated activation in the immediate past. The field p_m in this model should best be understood as an abstract collection of memory neurons, innervating the command field. The value of p_m at a given field site represents the strength of the synapses between the memory neurons and that site in the command activation field, so when the memory is active, the relevant sites become pre-activated. We thus keep using *preshape* as a legitimate name to refer to p_m.

3 Simulation and Results

The matlab model used for the last simulation was slightly retuned to account for the changes described in the last chapter. Most notably, both lateral excitation in u_m and global inhibition from v_m into u_m were increased. This made peaks in u_m more self sustained and less dependent on preshape input to be stable, increasing the chance of a switch to another field site that is receiving correctional input from the kr-field.

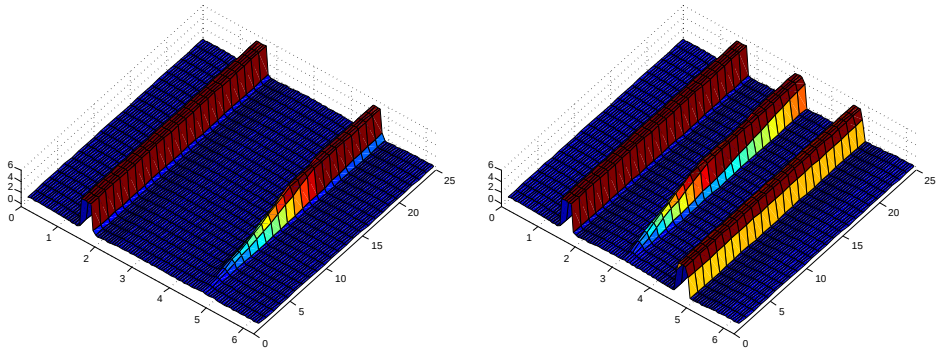

Figure 3: Learning to produce multiple different relative phases, as represented by preshape patterns. On the left the new relative phase $\psi_{tbl} = \frac{3}{2}\pi$ is being learned, in addition to the previously learned $\psi = \frac{\pi}{2}$. On the right, a third task of $\psi_{tbl} = \pi$ was added. Each session consisted of 25 trials and was started with a preshape field shaped from earlier learning trials.

With the new system for storing preshape memory, the preshape landscape changed appearance. A single preshape peak will now be surrounded with a small valley of negative preshape, resulting from the mexican hat shaped kernel. This gives a lower bound for how close two preshape peaks can be to each other. That lower bound is dependent on the width of the interaction kernels, mostly the preshape kernel k_{pre}. Given a larger distance between each other, however, the preshape peaks do not directly interact with each other, in the sense that a peak initialized by one peak does not contribute substantial input to the field sites of the other peak. This makes it possible for multiple different relative phases to be learned by the model. As can be seen in Fig. 3, the development of the preshape fields nicely shows the learning process. For each different relative phase ψ that was learned, there is a clear and distinct peak in the preshape field.

While the results of the preshape fields look remarkably clear cut, other measures of the process show different effects. Fig. 4 indicates a much stronger fluctuation in relative phase than was known from sessions with only one RP to be learned in the earlier model (called single-RP model from now on). The mean error shows relatively large shifts to both directions in the earlier trials, then stabilizes very close to 0 at about the 15th trial of each session. In both the single-RP model and the experiments described in [SZK92], the mean error decreased very quickly, while the stability increased at a much slower rate.

In the present model, stability seems to be bound to the mean error. In the first half of each session, the large mean error is accompanied by a low stability level. Stability seems to take a quick rise around the middle of each session, where as the mean error declines to almost zero, stability increases

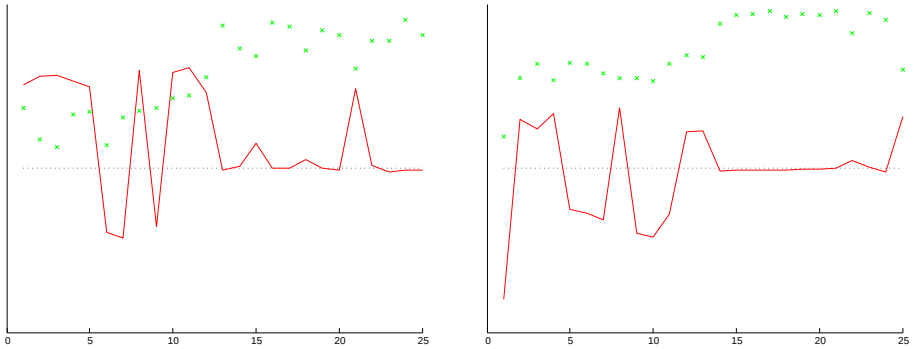

Figure 4: Mean attractor position (red line) and stability (green cross) during each trial, for the two learning sessions for $\psi_{tbl} = \frac{3}{2}\pi$ and $\psi_{tbl} = \pi$ with different prior learning corresponding to the two sessions described in Fig. 3. Values were directly measured from the coordination dynamic phase plot, stability taken as the mean slope at the attractor during that trial. The analytical measures show a clearer picture than the empirical ones, as these are overly obscured by noise, so they were chosen for this display.

by about 50%. Even very late in the sessions, there are still substantial mavericks in both curves, e.g. on the 21st trial in the first session, or the last trial in the second session. Those mavericks, too, simultaneously affect mean error and stability.

The reason for this can be easily seen in Fig. 5, which shows the details of the performed phase's development for example trials in both sessions. In the normal trials, the model finds a command that yields the correct relative phase at the beginning, resulting in the low error and high stability seen in Fig. 4. In the maverick trials, on the other hand, the relative Phase is far off for a considerable length of time at the beginning of the trials. In the $\psi_{tbl} = \frac{3}{2}\pi$ trial, the attractor lingers at $\phi = \frac{\pi}{2}$, in the $\psi_{tbl} = \pi$ around $\phi = \frac{3}{2}\pi$. This means that the initial peaks stabilized by the command field were at $\frac{\pi}{2}$ and $\frac{3}{2}\pi$ respectively. The system fell into the 'trap' of the wrong preshape peak, until the error was noted by the feedback field and the corrective input became strong enough to stabilize a new peak at the correct site and let the old, misplaced peak decay.

Scanning trials

An important tool for understanding the development of the coordination dynamics during the learning process are so called scanning trials. This is a session of trials during which the task variable, the to-be-produced relative phase ϕ_{tar}, is varied over the whole range of available values. The mean performed error $\phi_{mean} - \phi_{tar}$ of these scanning trials gives information about the state of the coordination dynamic at the time the session is taken.

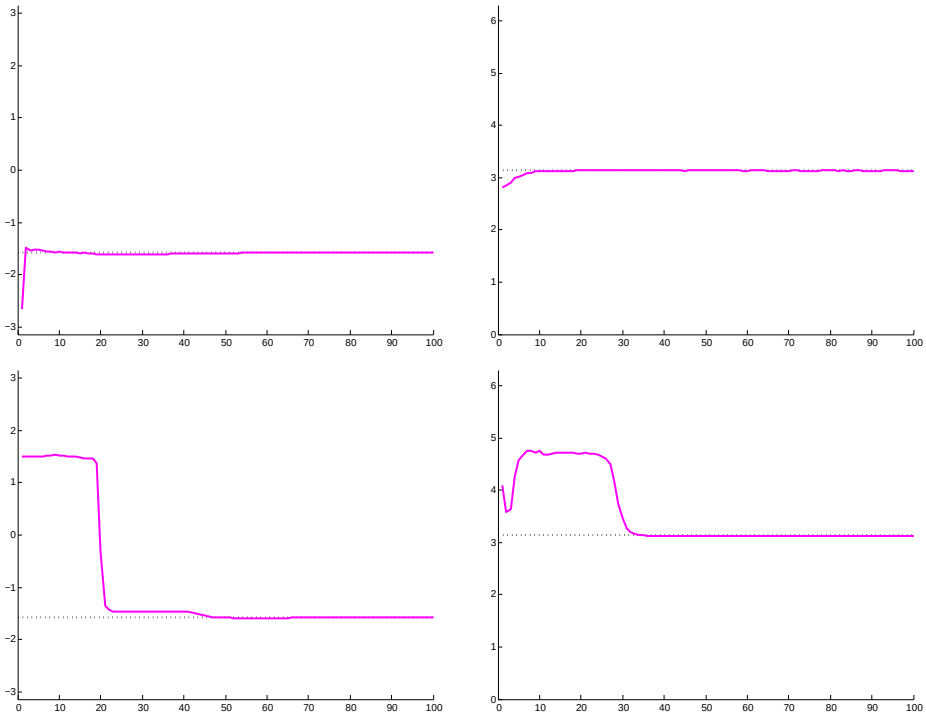

Figure 5: Performed relative phase over single trials. Examples for normal trials (top) and mavericks (bottom) for the two sessions described in Fig. 3. Displayed are trial 20 (top) and 21 (bottom) for the first session with $\psi_{tbl} = \frac{3}{2}\pi$, and trial 24 (top) and 25 (bottom) for the session with $\psi_{tbl} = \pi$. For better display the phase has been shifted by -2π in the first session.

Fig. 6 gives plots of the mean performed error for two different model states. The change in the coordination dynamic during the trials is clearly visible in the graphs. The first session, taken before learning, shows bigger fluctuations with a pattern influenced by the intrinsic part of the coordination dynamic. Around $\psi_{tar} = 0$ and to a lesser extent $\psi_{tar} = \pi$ the mean error is near zero, with high stability, corresponding to attractors in the intrinsic dynamic, between these sites, the mean error increases and stability declines. The stability curve displays the same characteristic as the 'seagull shape curve', only inverted, because the stability measure is taken to be the slope at the attractor, instead of standard deviation of the performed relative phase.

The second session shows a much more consistent behaviour, with less fluctuations. The mean error ist lowest around $\psi_{tar} = \frac{\pi}{2}$, the phase that was practiced in the learning session that preceded the scanning trials. To both sides, the mean error increases in absolute value, but in different directions, as the performed relative phase is drawn towards the learned pattern at $\frac{\pi}{2}$.

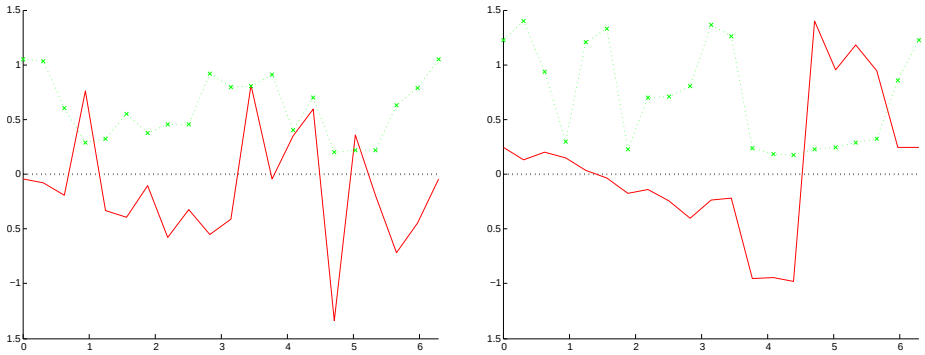

Figure 6: Mean performed error (red) and stability (green) from scanning trials for two different conditions. The first session was done with a fresh system, without any structure in the preshape field. The second session was taken after the system had learned to produce a relative phase of $\psi_{tbl} = \frac{\pi}{2}$. Scans were done in 20 trials, increasing ψ_{tar} by equal steps from 0 to 2π each trial. The trial for $0 = 2\pi$ was duplicated in the plot to align the edges of the graph. Analytical measures were chosen, so mean performed relative error is measured in attractor position, stability in attractor slope. Stability was rescaled for display reasons, y-axis applies to mean error.

4 Discussion

The scanning results conform very nicely with the examples of scanning trials for similar situations, before and after a practice session, given in [SZK92]. The reasons those curves look how they do, however, differ.

In the model given in [SZK92], the coordination dynamic is the result of several influences, some of which work in opposition to each other. We have

$$\dot{\phi} = f_{intr}(\phi) - c_{env} f_{info}(\phi, \psi_{env}) - c_{mem} f_{info}(\phi, \psi_{mem}),$$

with $f_{info}(\phi, \psi) = -\sin(\psi - \phi)$. The attraction of the performed relative phase towards memorized patterns is a result of the $c_{mem} f_{info}(\phi, \psi_{mem})$ term diverting the attractor towards previously memorized relative phase ψ_{mem}. During each trial, the attractor is fixed.

In the model presented here, the divertions of the relative phase result from the pattern described above. At the beginning of a trial, the command field develops a peak at the site supported by the memory preshape. After a while, the corrective input from the *kr*-field leads to the formation of a new peak, destabilizing the old one. So during the trial, the attractor itself moves towards ψ_{tbl}, which diverts the average relative phase of that trial towards the site of the first command, initiated by the preshape peak.

This difference is inherent in the model created here to describe the learning process. Whether it is reconcilable with experimental data remains to be checked. It should be stated, though, that several values like the length

of a single trial and the amount of noise in the command field were not really tuned, so changing these might lead to different results, where the command peak might switch more often during a single trial, due to noise, and then be directed towards the preshape peak more often than not, producing a similar picture in the scanning trials, with a less awkward explanation for the divertion of the mean error.

Another thing that was left mostly untouched is the width of the interaction kernels. While the kernels used in this simulation were narrower than in the last simulation, supported by a higher resolution in the field, all results presented here have been reached with kernels of the same size. During the tuning of the kernel width, though, a phenomenon occurred where peaks in the command field would oscillate around the tasked relative phase ψ_{tbl} for a long time and over many trials, building two slim preshape peaks, one on each side of ψ_{tbl}. This turned out to be the result of the intrinsic part's influence on the coordination dynamic. When ψ_{tbl} would coincide with a repellor in the intrinsic dynamic, the system was never able to stabilize a satisfying pattern. Once a peak was stabilized, the feedback field would create corrective input. But this input would be inaccurate, as the error was excited by the intrinsic dynamic. The resulting solution would then overshoot, creating corrective input in the other direction, which would again be too large. The system then proceeded to oscillate between these two states.

There were two ways to prevent this phenomenon. The first was to simply increase the kernel range. Initially the same pattern established itself, of two separate peaks, one on each side of ψ_{tar}, and the command field oscillating between the two peak sites. Through the wider kernels, though, the preshape field sites between the peaks were sufficiently increased by activity at the two sites the command field oscillated between to stabilize peaks closer to ψ_{tbl}. These peaks increased preshape around ψ_{tbl} even more, so that after some time the two preshape peaks joined to form a single big peak centered at ψ_{tbl}.

The other solution to the oscillation problem was to simply decrease the influence of the intrinsic dynamic. Dividing the coefficients a and b of f_{intr} by two, the repellors of the intrinsic dynamic were not strong enough to create the problematic oscillations anymore.

In [HKB85], the involuntary shift between in-phase and anti-phase patterns that occurred as the frequency of the pattern was increased, was attributed to the attractor solution at π becoming unstable due to shifting coefficients of parts of the coordination dynamic. The present findings could be interpreted in a similar way, as the reduction of the coefficients a and b stabilizes the solution at $\frac{\pi}{2}$. If that reduction can be linked to a reduction in frequency, this might turn out to be a different instance of the same underlying mechanism as the one initially described by Kelso.

A Parameters

Command field

$\tau_u = 20, \quad \tau_v = 4, \quad \tau_w = 80$
$h_u = -10, \quad h_v = -7, \quad h_w = -4$
$\beta_u = 2, \quad \beta_v = 1, \quad \beta_w = 1$

$c_{uu} = 2.5, \quad \sigma_{uu} = 5$
$c_{uv} = 0.5, \quad \sigma_{uv} = 10, \quad g_{uv} = 1$
$c_{uw} = 0, \quad \sigma_{uw} = 5$

$c_{vu} = 1.4, \quad \sigma_{vu} = 5$
$c_{vw} = 0, \quad \sigma_{vw} = 5$

$c_{wu} = 2, \quad \sigma_{wu} = 5$
$c_{wv} = 0.3, \quad \sigma_{wv} = 10, \quad g_{wv} = 0.5$
$c_{ww} = 0.5, \quad \sigma_{ww} = 5$

Preshape

$\tau_{pre} = 6000, \quad x_1 = 0.1, x_2 = 5, \beta = 10$

Coordination

$a = b = 0.5, \quad \tau_{rel} = 0.5, c_{com} = 2, \quad \beta_{com} = 0.4, \quad \text{threshold}_{com} = 5$

Feedback

$\tau_{kr} = 160, \quad h_{kr} = -8, \quad c_{fb} = 1$

References

[HKB85] H. Haken, J. A. S. Kelso, and H. Bunz. A theoretical model of phase transitions in human hand movements. *Biol. Cybern.*, 51:347–356, 1985.

[SZK92] G. Schöner, P.G. Zanone, and J.A.S. Kelso. Learning as change of coordination dynamics: Theory and experiment. *J. Mot. Behav.*, 24(1):29–48, 1992.