

Perpetual Crossers without Sensory Delay: Revisiting the Perceptual Crossing Simulation Studies

Eduardo J. Izquierdo, Gabriel J. Severino, Haily Merritt

Cognitive Science Program, Indiana University Bloomington
Luddy School of Informatics, Computing, and Engineering, Indiana University Bloomington
edizque@iu.edu

Abstract

We revisit the perceptual crossing simulation studies, which are aimed at challenging methodological individualism in the analysis of social cognition by studying multi-agent real-time interactions. To date, all of these simulation studies have reported that it is practically impossible to evolve artificially a robust behavioral strategy without introducing temporal delays into the simulation. Also, all of the studies report on a single strategy: a perpetually crossing agent pair. Here, we systematically report on the evolutionary success of neural circuits on the perceptual crossing task, with and without sensory delay. We also report on two different strategies in the ensemble of successful solutions, only one of which had been discussed in the literature previously.

Introduction

Research on social cognition has largely assumed that studying a single individual engaged in a social interaction is sufficient to understand the dynamics and behavior that constitute a social interaction. In the last couple of decades, there have been calls to take the social interaction itself, instead of the individuals in isolation, as the object of study (Schilbach et al., 2013). From this interactionist perspective, social interaction is more than simply the arena in which social cognition plays out; it enables or constitutes social cognition (De Jaegher et al., 2010; Froese and Di Paolo, 2011).

Making social interaction the object of study, instead of a social agent, need not entail more complex experiments. One minimalist example of an interactionist experimental paradigm is perceptual crossing (Auvray et al., 2009). In these experiments, participants are asked to identify when they think they are interacting with a partner participant in a simple one-dimensional virtual environment while unaware of what they are actually interacting with. The tasks are designed so that they cannot be solved by either participant independently; successful identification of the partner necessitates mutual interaction. The dynamics of behavior that result suggest that studies of social interaction should never be limited to analyzing a single individual's behavior. Recent work has expanded the paradigm to two dimensions (Rohde and Paolo, 2008; Rohde, 2010; Lenay et al., 2011), to

the domain of human computer interaction (Barone et al., 2020) and to different populations, from adults and adolescents (Hermans et al., 2020; Froese et al., 2020; Iizuka et al., 2015; Froese et al., 2014; Lenay and Stewart, 2012) to individuals with autism (Zapata-Fonseca et al., 2018).

There has also been a growing interest in using simulation studies in order to investigate the dynamics of social interaction (Di Paolo, 2000; Quinn, 2001; Iizuka and Ikegami, 2004; Ikegami and Iizuka, 2007; Iizuka and Di Paolo, 2007; Williams et al., 2008; Di Paolo et al., 2008; Froese and Di Paolo, 2008; Reséndiz-Benhumea and Froese, 2020; Reséndiz-Benhumea et al., 2021). Some of these models have been specifically designed to generate insights for mutually informing collaborations between the field of artificial life and the traditional empirical sciences (e.g. Ikegami and Iizuka, 2007; Di Paolo et al., 2008; Rohde and Paolo, 2008). This is particularly true for the perceptual crossing paradigm, where there have been a set of rich contributions from simulation studies that have managed to successfully replicate the experiment and contribute to hypotheses to be tested in further social experiments (Di Paolo et al., 2008; Froese and Di Paolo, 2010). In particular, the simulation studies have predicted challenges and patterns of interactions that would be faced by human participants (Di Paolo et al., 2008). In some cases, these predictions have then been supported by experimental evidence from humans (Auvray et al., 2009), facilitating model-experiment dialogue (for a review see Auvray and Rohde, 2012).

Despite the advances, important questions remain open. First, simulation studies have all relied on the introduction of a sensory delay for the agents to perform the perceptual crossing task successfully (Di Paolo et al., 2008; Froese and Di Paolo, 2010, 2009). Crucially, the practical need for delays in the models has been considered a potentially important component for the explanation of the adaptive performance of the task in human participants, and has motivated psychological studies. However, the necessity of a sensory delay in human participants is unlikely (Iizuka et al., 2015). Second, the dominant (or potentially the only) strategy that has been discussed in the simulation literature so far has

been a perpetually crossing strategy, where agents continue to cross back and forth perpetually. It is unclear whether any other strategies are feasible. However, most of the simulation studies report on only one of the solutions, not on the full ensemble of possible solutions.

In this paper, we revisit the original work on the evolution of perceptual crossing agents and we extend this work to answer the open questions above. The rest of this paper is organized as follows. In the next section, we describe the perceptual crossing task and the set up of the agents for all experiments. Then we present results from a series of three experiments which explore the evolution of perceptual crossers under various conditions. Finally, we conclude with a general discussion of the experimental results, and outline some directions for future work.

Methods

We set out to replicate the agent and task as described in previous simulation studies (Froese and Di Paolo, 2010; Di Paolo et al., 2008; Froese and Di Paolo, 2009). Two agents coexist on a ring (i.e., a one-dimensional environment that wraps around; Fig. 1A). Agents are able to move around the ring with a maximum velocity (2 units of space per unit of time) in either direction. There are three distinct types of objects that can be encountered by an agent (Fig. 1A): the other agent's avatar, the shadow of the other agent, and a static object. Each object occupies a total of 2 units of space. The shadows are 48 units of space away from the agent. The ring is 600 units in circumference, and the fixed objects are placed across from each other at 150 and 450. Agents can move past each other and their respective static objects unimpeded. The neural controller that governs movement (described below) is rotated from one agent to the other, so that left and right movement aligns with the orientation of the agent. The shadows of the agents are reflected about the ring, so that one agent's shadow is to its left and the other agent's shadow is to its right, as depicted in Fig. 1A. The sensory input of an agent is activated (set to 1) when its receptor field overlaps with another object; otherwise it remains off (set to 0).

The behavior of each agent is controlled by a continuous-time recurrent neural network (Beer, 1995) with the following state equation:

$$\tau_i \dot{y}_i = -y_i + \sum_{j=1}^N w_{ji} \sigma(y_j + \theta_j) + g_i s + I_i \quad (1)$$

where y_i is the state of each neuron, τ is the time constant, w_{ji} is the strength of the connection from the j^{th} to the i^{th} neuron, θ is a bias term, $\sigma(x) = 1/(1 + e^{-x})$ is the standard logistic activation function, g_i is the sensory weight from the binary sensor s to neuron i , and I_i represents an external input to each neuron. The output of a neuron is

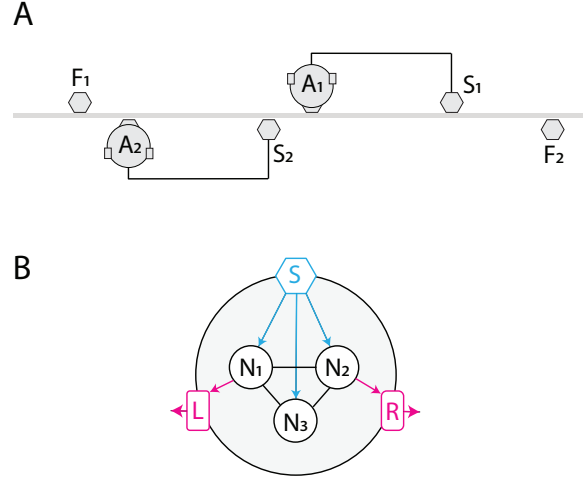


Figure 1: Task and agent setup. **(A)** The task takes place in a 1-dimensional ring where two agents face each other. Each agent can sense the other's avatar (A), a shadow of the other's avatar (S), and a fixed object (F). **(B)** Each agent has a sensor (cyan) that can send information to all N neurons (black). The neurons in the circuit are fully inter-connected, including self-connections (not depicted). The output from one neuron drives the left motor and another neuron drives the right motor (magenta). The neural circuits in the two agents are identical (i.e., they have the same parameters).

$o_i = \sigma(y_j + \theta_i)$. Following the original simulation studies (Fig. 1B), the sensor, s , is fully connected to all neurons in the circuit; the neurons are fully interconnected (including self-connections); and two of the neurons are chosen to drive the left and right motors, respectively. The velocity of an agent is proportional to the difference between the outputs of the two motor neurons: $v = \gamma(o_1 - o_2)$, where o_1 and o_2 represent the outputs of the neuron controlling the left and right motors, respectively, and γ is a constant that determines the agent's maximum possible velocity. In all of our simulations, the maximum velocity was set to $\gamma = 2$.

As with the original simulation studies, the two agents have identical neural controllers. The neural parameters of the controller are evolved using a real-valued genetic algorithm. Given that both agents are clones of each other in terms of their neural controller, each genome encodes the parameters for only one neural controller. The following neural parameters, with corresponding ranges, are evolved: time-constants $\tau \in [1, 10]$, biases $\theta \in [8, 8]$, and all connection weights (from sensors to neurons, g , and between neurons, w) $\in [8, 8]$. We used a generational algorithm with rank-based selection and a population size of 96 genotypes. Successive generations are formed by first applying random Gaussian mutations to each parent genome with a mutation variance of 0.05 (see Beer, 1996 for details). In addition, uniform crossover is applied with 50% probability. A child

replaces its parent if its performance is greater than or equal to that of the parent; otherwise the parent is retained.

The fitness evaluation is intended as a replication of the original simulation studies, such that neural controllers are evolved so that the two agents successfully find each other. We evaluate the performance of a pair of agents by systematically varying the starting location of the two agents. Specifically, the starting location for the first agent in a pair is chosen between 0 and 600 in steps of 50; the starting location for the second agent in the pair is between 0 and the first agent's starting location, for a total of 78 trials. Each trial lasts 800 time units and proceeds as follows: (1) the neural states of both agents are initialized to 0. (2) During the first 400 units of time, the agents interact without evaluation. We treat this as a transient period because it allows for agents initialized at the maximum starting distance moving at their maximum velocity enough time to traverse the ring environment and find each other. (3) For the remainder of the simulation after the transient period, we record and normalize the distance between the two agents. For a given trial, the score that a given pair of agents with a given neural controller receive is:

$$f = 1 - \frac{\bar{d} - 2}{298} \quad (2)$$

where \bar{d} is the average separation between the two agents during a trial (excluding the initial transient period), 298 is the maximum spatial distance between the two agents. Since the 1-D environment wraps around between 0 and 600 units, 300 is the maximum spatial distance between points on the ring; and because the agents are 2-units wide and the sensors are binary, the agents cannot detect proximity beyond 2 units of space away from each other. The final fitness of the evaluation is the average fitness across all trials.

While we try to maintain as close a replication to the original study as possible, we summarize the key differences between the original fitness evaluation and ours. First, the distance during the initial transient period is not taken into consideration for the fitness here. Second, the fitness here is normalized to run between 0 and 1 based on the minimum distance at which an agent can sense the other agent. Finally, this evaluation is deterministic: the starting positions of the agents are deterministic, the position of the fixed objects does not change, and the relative position of both shadows to the agent is fixed. Additionally, there are two minor differences between our agent/task setup and the original study: the objects (the agent's avatars, the agents' shadows, and the static objects) occupy 2 units of space instead of 4; and the maximum velocity is 2 instead of 1. Also, the stochastic search algorithm used in our simulation is different from those used previously. As far as we can tell, each simulation study that has replicated the work has used different stochastic search algorithms, and we have no strong reason to believe that the results depend on the particularities of it. Crucially, as with the original studies, since the avatars,

shadows, and fixed objects are *indistinguishable* to either agent, success in this task requires that the agents evolve a system for accurately detecting *mutual* interactions. As it has been demonstrated by previous implementations of the model, including only distance in the fitness function makes for a non-trivial task (Froese and Di Paolo, 2010, 2009).

Part I: Replicating original results

The goal of our first set of simulations was to replicate the experiments in the original studies. The agents must solve the perceptual crossing task with the original sensory delay of 2.5 units of time (25 timesteps using a step size of integration of 0.1) and with a fitness function that selects for close average proximity between the two agents. We report on the evolutionary performance across different circuit sizes and on two different strategies observed in the ensemble of successful perceptual crossers.

What general trends are observed in the evolution of perceptual crossers with sensory delay? One hundred evolutionary runs were performed for two-, three-, and four-node circuits (see Fig. 2). There are two main groups of solutions. The first and most dominant ones can be seen in the peak around 0.92 in the histograms, comprising 93% of all evolutionary runs with two-neuron circuits, 70% of three-neuron circuits, and 62% of four-neuron circuits. These solutions entirely fail to find the other agent from a small subset of the starting conditions. Because this group fails to solve the task from all possible starting configurations, we do not study them in any more detail in this paper. There is a second group of solutions that solve the problem nearly perfectly (>0.99). The size of this second group increased with the number of neurons: 1% of all two-neuron circuits, 10% of three-neuron circuits, and 16% of four-neuron circuits. We analyze the behavior of this group of solutions in more detail in the remainder of this section.

What are the overall tendencies observed in the behavior of successful perceptual crossers? We analyzed all solutions with a fitness greater than 0.99 across a wider range of conditions than were examined during evolution. Specifically, we analyzed the performance of this high-performing ensemble over 100×100 starting conditions across the full spatial range $[0, 600]$ and over a range of shadow distances $[48, 52]$. We kept track of the performance of these individuals in three different ways: (a) the performance as measured in the original simulations studies (i.e., with the transients and without normalization based on sensory-range; x -axis, Fig. 3A); (b) the performance measured without the initial transient and with normalization (y -axis, Fig. 3A); and (c) the number of times the two agents crossed each other in a trial (Fig. 3B). The results of the behavioral robustness analysis is shown for all 27 high-performing solutions in the ensemble in Fig. 3.

We highlight three key observations from this analysis. First, most of the circuits are behaviorally robust across a

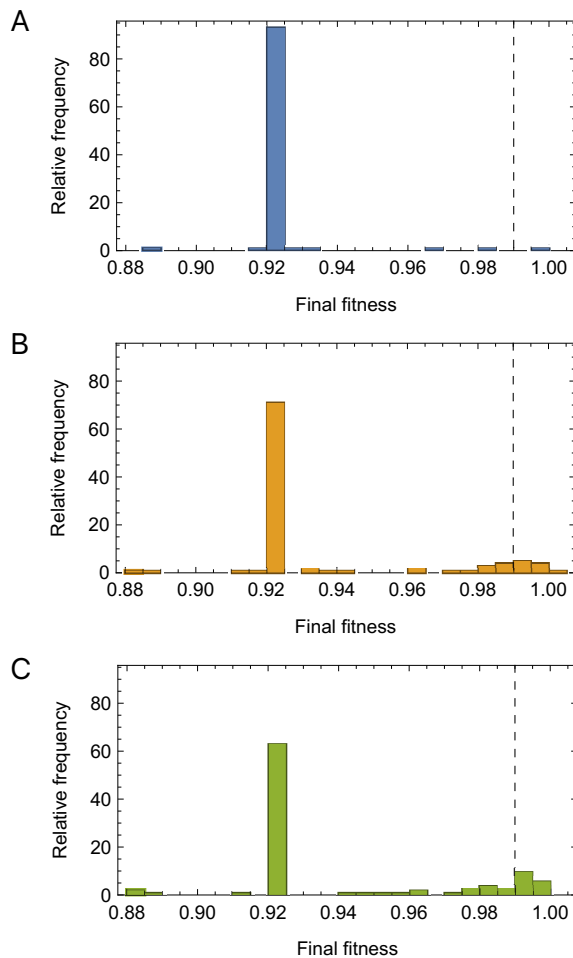


Figure 2: Evolutionary performance statistics. Performance histograms for two-neuron (A), three-neuron (B), and four-neuron (C) circuits. Relative frequency of binned proximity as a percentage of total trials is plotted. The maximum proximity is 1.0. The dashed line depicts the cutoff of 0.99 for agents analyzed in more detail.

wider set of conditions: 81.48% of the solutions maintained a performance greater than 0.95. This suggests that the conditions presented during evolution, including particularly a relatively small set of starting positions and a fixed shadow distance, were sufficient for agents to generalize. Second, our fitness function offers two improvements on the original fitness function: (a) Equally fit solutions no longer obtain a wide range of fitness; (b) The fitness of some great solutions is no longer indistinguishable from much worse solutions. The fitness function is one potential explanation for the difficulties evolving mentioned in the original simulation studies. Finally, not all successful solutions perpetually cross. This last observation was unexpected. As far as we gathered from the literature, there were no reports of minimally crossing solutions; only of perpetual crossing solutions.

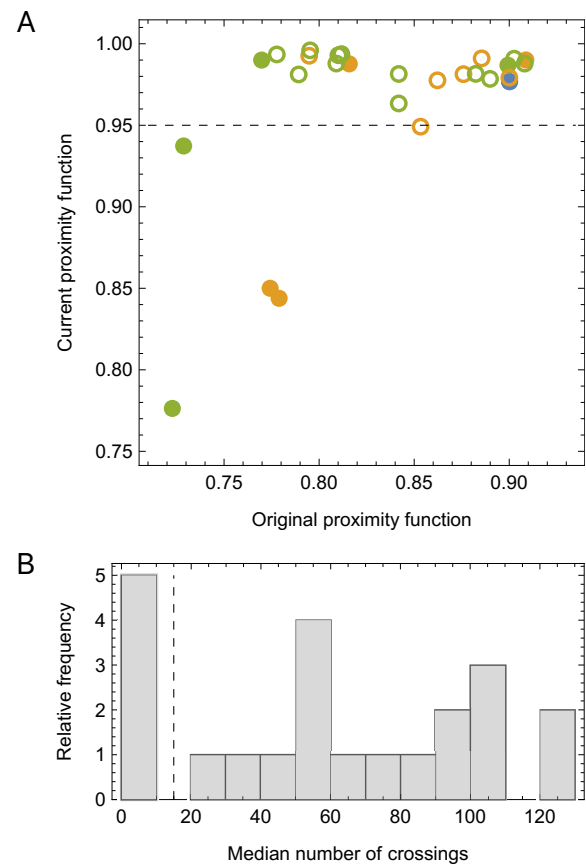


Figure 3: Behavioral robustness statistics. (A) Relationship between original and current measure of performance. Each point represents one solution from the top ensemble of solutions. The current measure of proximity excludes transients and normalizes the distance based on the sensory range. The color of the point represents the size of the circuit: two-neuron (blue), three-neuron (orange), and four-neuron (green). Filled disks represent solutions with minimal number of crossings. Open circles represent solutions that cross perpetually. (B) Median number of crossings across the solutions that achieved a robustness performance greater than 0.95 (above the dashed line in panel A). Solutions with fewer than 15 crossings (dashed line) per trial were labeled as minimal crossers. Solutions with greater than 15 crossings were labeled perpetual crossers.

Part II: Agents without sensory delay

Having replicated the original results, the second major goal of ours was to attempt to evolve perceptual crossers without a sensory delay. Given that previous reports always included a sensory delay (Di Paolo et al., 2008; Froese and Di Paolo, 2010, 2009), we did not expect these evolutionary runs to succeed. Nevertheless, it would be important and informative to understand how and why the no-sensory-delay

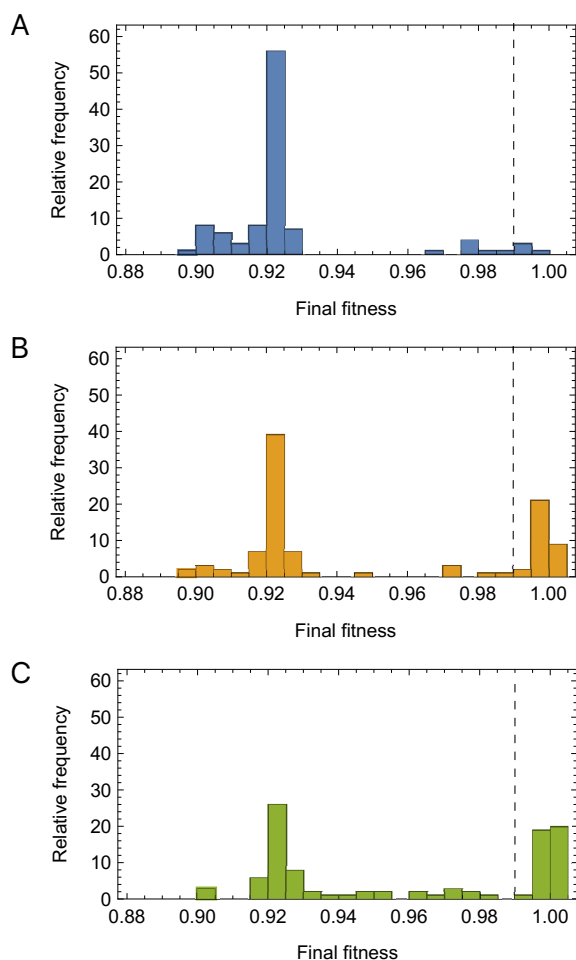


Figure 4: Evolutionary performance statistics for agents without sensory delay. Performance histograms for two-neuron (A), three-neuron (B), and four-neuron (C) circuits. Dashed line depicts 0.99 cutoff for selecting agents to analyze further.

condition failed.

We again performed one-hundred evolutionary runs for two-, three-, and four-node circuits to solve the perceptual crossing task, this time without a sensory delay. We maintained the same modified version of the fitness function that selects for proximity between the two agents. As in the previous experiment, there was a peak of evolutionary runs that became stuck around a fitness of 0.92 (Fig. 4). Crucially, and contrary to what had been reported until now, we observed that a substantial number of evolutionary runs succeeded (fitness > 0.99). Moreover, the proportion of successful runs was substantially larger without a sensory delay than with it: 4% of all two-neuron circuits, 32% of three-neuron circuits, and 40% of four-neuron circuits.

Do solutions without a sensory delay generalize well such that agents find each other across a broad range of condi-

tions? As with the first set of experiments, we performed a behavioral analysis of all 76 solutions with fitness > 0.99 (Fig. 5). We highlight the key insights from this analysis. First, solutions without the delay were largely robust to the wider range of starting conditions and to the different distances between an agent's shadow and avatar. All but one of the 76 solutions had performance above 0.95 using our updated proximity function (Fig. 5A). Second, only one of the 75 robust solutions crossed perpetually; all others crossed fewer than 15 times on average (Fig. 5B). Based on these results, we define minimal crossers as pairs of agents that cross each other fewer than fifteen times on average and perpetual crossers as pairs of agents that cross each other continuously, or more than fifteen times. This last finding prompted the next set of experiments.

Part III: Promoting perpetual crossing

When agents were evolved with sensory delay, we replicated the successful findings of the original simulation studies. Notably, in addition to the perpetual crossing strategy that had been reported originally, we observed a second strategy: minimal crossers (i.e., pairs of agents that crossed each other fewer than 15 times on average within a given trial). When agents were evolved without sensory delay, we were surprised to find that they still succeeded at finding each other, but primarily using the minimally crossing strategy. Given that the previous simulation studies primarily focused on analyzing solutions that cross perpetually, the natural follow up question was: Can we reliably generate agents without sensory delay that solve the problem with a perpetual crossing strategy?

In this third and final set of experiments, we set out to evolve agents without sensory delay to find each other and cross perpetually. Our goal was to accomplish this by only redesigning the fitness function to more closely match the desired behavioral goal. Specifically, we introduced a term that promoted perpetual crossing, in addition to proximity. Including the additional term allows us to ask whether the perpetual crossing strategy only arises in the presence of sensory delay. Thus, here our motivation was more about refining and further understanding the modeling approaches to perceptual crossing than directly replicating perceptual crossing work in humans.

The additional perceptual crossing term simply counted the number of times two agents crossed and averaged this across all the different starting conditions. We explored three different ways to introduce this additional term and we only succeeded with one. In one batch of experiments, we multiplied or added the proximity term and the crossing-count term together. In a second batch, we used an incremental approach in two stages: in the first stage, only proximity was evaluated; in the second stage, the proximity term and the crossing-count term were again multiplied or added together. Neither of these two strategies successful results.

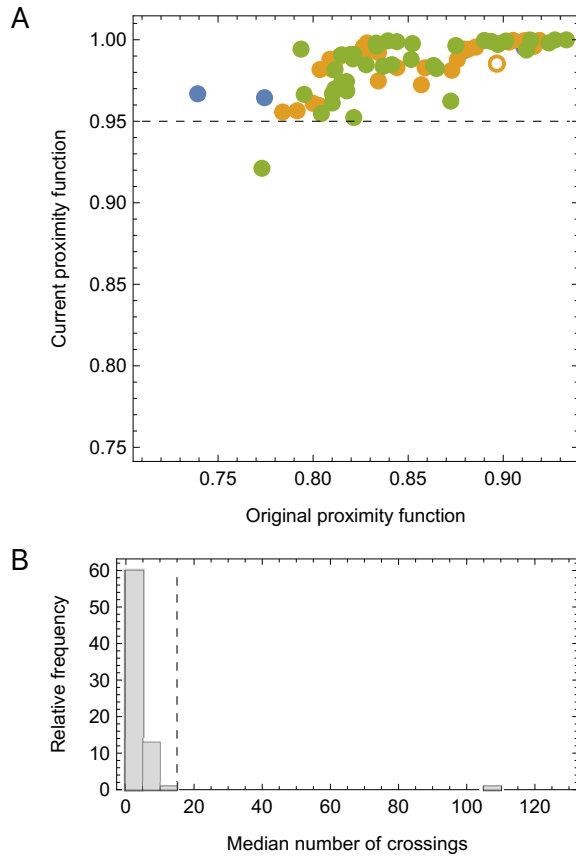


Figure 5: Behavioral robustness statistics of agents with no sensory delay. (A) Relationship between original and current measure of performance. Each point represents one solution from the top-performing ensemble. Color represents circuit size: two-neuron (blue), three-neuron (orange), and four-neuron (green). Filled disks represent solutions with minimal number of crossings. One open circle represents a perpetual crossing solution. (B) Median number of crossings across the solutions that achieved a robustness performance greater than 0.95 (above the dashed line in panel A).

In the final batch of experiments, we employed a conditional fitness evaluation: if the proximity term was lower than 0.99, only it counted towards fitness; if the proximity term was higher than 0.99, then the fitness involved the sum of the proximity and the crossing-count term. Using this formulation, agent pairs can achieve a fitness greater than 1. Only with this approach did we obtain successful perpetual crossers.

Is it possible to evolve perpetual crossers without a sensory delay? We performed the final set of one-hundred evolutionary runs for two-, three-, and four-node circuits. As with all previous experiments, there was again a large number of runs that became stuck around a fitness of 0.92 (Fig. 6). In this run, a smaller batch of runs got stuck also at

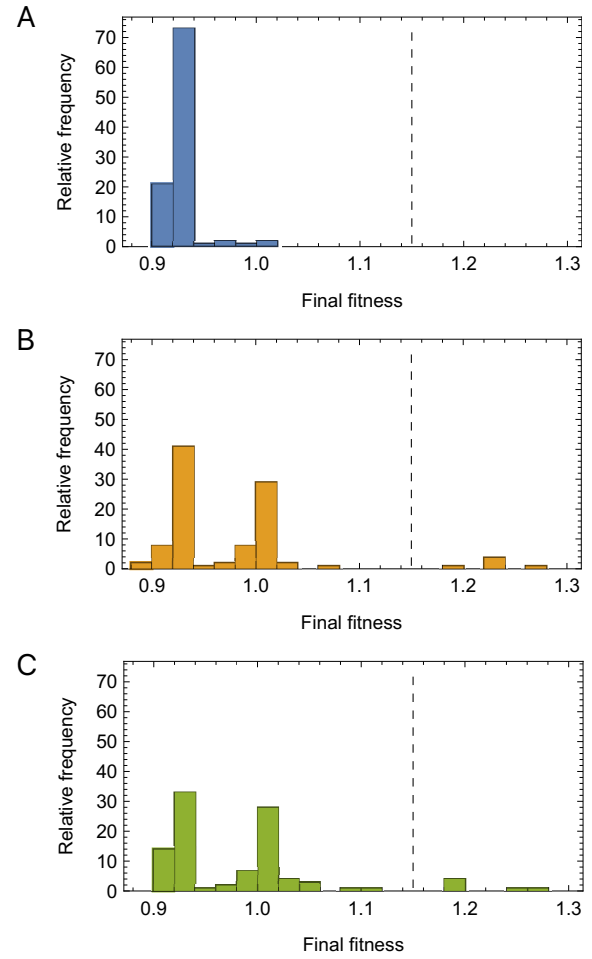


Figure 6: Evolutionary performance statistics for agents without sensory delay and with a fitness function that encourages crossing. Performance histograms for two-neuron (A), three-neuron (B), and four-neuron (C) circuits. The dashed line depicts the cutoff of 1.15 for selecting agents to analyze in more detail.

a fitness of around 1. This corresponds to solutions that can find each other perfectly, but do not cross perpetually. Only a small fraction of the evolutionary runs surpassed both challenges. Based on observations of the behaviors, we counted the number of successful solutions as those that surpassed a fitness of 1.15 on the combined task: none of the two-neuron circuits, 6% of three-neuron circuits, and 6% of four-neuron circuits. All 12 of the successful perpetual crossing solutions had a performance above 0.95 on the robustness test (Fig. 7A); and all of them had a median number of crossings above 80 (Fig. 7B).

What can we learn from the behavior of a perceptual crossing agent without sensory delay that crosses perpetually? Although a detailed analysis of the dynamics of one of these circuits is outside the scope of this contribution, we

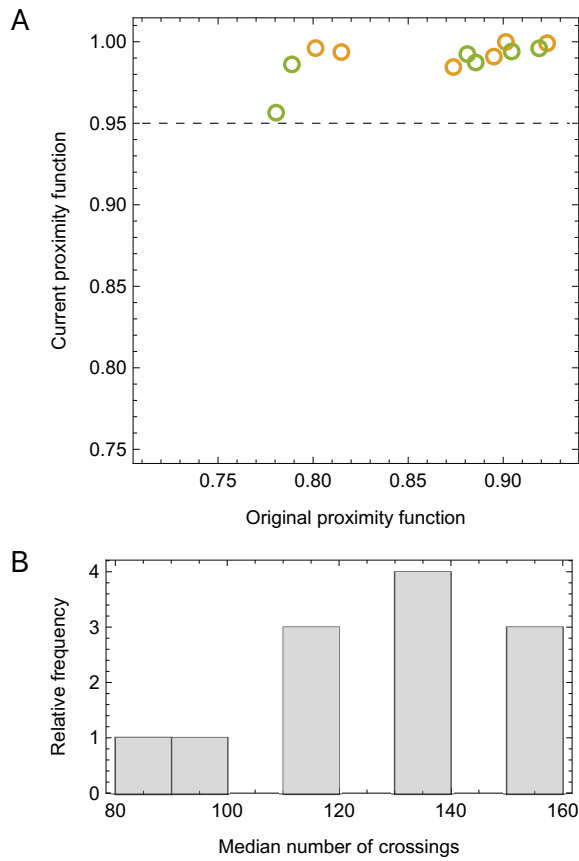


Figure 7: Behavioral robustness statistics for agents without sensory delay evolved to cross perpetually. (A) Performance of the best solutions (each point) using the original and the current measure of proximity. Color represents circuit size: two-neuron (blue), three-neuron (orange), and four-neuron (green). All solutions are perpetual crossers, as depicted by the open circles. (B) Median number of crossings across the solutions that achieved a robustness performance greater than 0.95 (region above dashed line in panel A).

can learn something about the operation of these circuits by looking at examples of their behavior. In this final section, we visualize the behavior of one of the top three-neuron circuits without sensory delay in three stages of detail (Fig. 8). First, we visualize the average proximity performance of the two agents throughout the full duration of a trial as a function of a wide range of starting conditions (100×100) (Fig. 8A). It is important to keep in mind that in this solution, like in most of the successfully evolved solutions, the agents always find each other and remain close to each other thereafter. This map of performance, then, does not reflect their ability to find each other, but rather how much exploration the agents exhibit as a function of the starting positions. A performance of 1.0 represents starting configurations where the agents find each other early in the evaluation trial; while

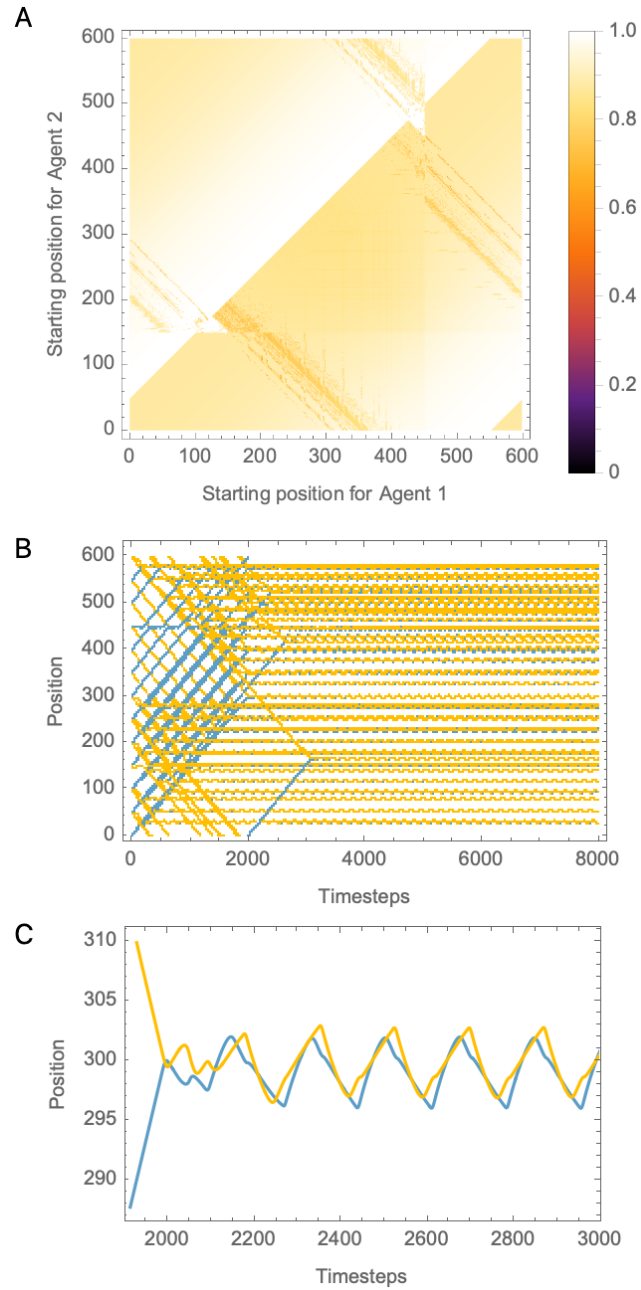


Figure 8: Behavior of one of the top three-neuron circuits without sensory delay. (A) Proximity performance of the two agents as a function of their starting positions. Color indicates proximity performance. (B) A sample of 78 trials from the full range of different starting conditions as examined during the fitness evaluation. (C) Detailed look at the interaction between two agents for one trial.

a performance of 0.8 represents agents that find each other after the first 200 units of time. As a second step, we can observe the agents' movement in the one-dimensional ring

over time given a smaller set of starting conditions (78 total) across all of time (Fig. 8B). One thing that is interesting to note here is that these agents find each other somewhat uniformly along the length of the environment; other solutions in the ensemble exhibited different patterns, and not always uniform. Finally, we can visualize a single trial over the small window of time where the agents first interact and then maintain a mutual crossing (Fig. 8C). One thing to note is that the pattern of crossing was quite different across all 12 of the different top three- and four-node circuits. For some of the solutions, the absolute position of the cycle of crossing stayed constant; for other solutions the pair drifted slowly in time while continuing to cross around each other. Overall, the main take-home message from looking at examples of the behavioral trajectories of some of these agents is that there is a wide variety of patterns of behavior according to which they could be grouped.

Discussion

In this paper, we set out to replicate the perceptual crossing simulation studies (Froese and Di Paolo, 2010; Di Paolo et al., 2008; Froese and Di Paolo, 2009) and refine the approaches used. First, we observed that evolving agents with a sensory delay resulted in two clearly distinct behavioral strategies: perpetual crossers (agents that find each other and continuously cross) and minimal crossers (agents that stop moving after crossing each other a limited number of times). As far as we are aware, only the former had been reported in the literature (Froese and Di Paolo, 2010; Di Paolo et al., 2008; Froese and Di Paolo, 2009). Presumably, perpetual crossers are preferred because they continuously interact. Second, we succeeded at artificially evolving agents without the sensory delay, contrary to what has been previously reported (Di Paolo et al., 2008; Froese and Di Paolo, 2010). However, an analysis of the successful solutions revealed that nearly all of them adopted a minimally crossing strategy. Finally, by modifying the fitness function to select for both proximity and crossings, we were able to generate agents without sensory delay that adopted the perpetually crossing strategy.

There are two factors that are likely to have played an important role in the success of the evolutionary runs in our simulation studies. First, by including the initial transient of the behavioral trajectories of the agents in each trial, the fitness function in the original simulation studies blurred the performance of otherwise successful circuits (c.f. Fig. 3A). By eliminating the transients, we could measure with more precision the percentage of trials where agents found each other. Second, the original simulation studies involved a stochastic fitness evaluation, varying the starting position of the agents and the distance of the shadows in each trial. The purpose of this was to make sure the agents learned the task robustly. However, for this task, a deterministic fitness evaluation was sufficient to produce equally robust solutions.

Finally, the positioning of the shadow in these experiments led to some initial exploratory and counter-intuitive results. In every one of the formulations of the task that we examined (Froese et al., 2014; Froese and Di Paolo, 2010), there was no explicit statement on whether the shadow was positioned to the left or the right of the agent/participant. Most crucially, it was not stated whether the shadow of one agent was reflected or rotated with respect to the other agent. However, in all the schematics of the task (Froese et al., 2014; Froese and Di Paolo, 2010), except for one of them (Auvray and Rohde, 2012), the shadow appeared to be reflected. Importantly, the original paper (Auvray et al., 2009) has a schematic that represents the shadows as reflected. In preliminary experiments, we examined both conditions (although we only report here on the reflected condition). The rotated-shadow condition provided counter intuitive results. Although we might predict that the task would be impossible because agents would end up mutually oscillating around each other's shadows, thinking they have found the other agent, evolution reliably found a clever hack that relied on the symmetry of the nervous systems. Because the two agents start moving in the same direction (left or right), they always encounter the other's avatar and shadow in the same sequence: one first and then other (depending on the direction). This allows them to "hardcode" which stimulus to center on without the requirement to mutually interact. This is, of course, not a possibility for the condition where the shadows are reflected.

Future Work

The next step for future work is to perform detailed analyses of the evolved perceptual crossers without sensory delay. Using the mathematical tools of dynamical systems theory, one can identify how the underlying dynamical structure of the agents supports their joint interaction. A psychophysical analysis might illuminate additional differences between the perpetual and minimal crossers that we identify. Such an analysis might also be useful for connecting our results to empirical work on humans, further strengthening the model-experiment loop that the perceptual crossing paradigm has established. Of particular interest would be to run new empirical experiments with humans to test whether participants can be prompted to use perpetual versus minimal strategies depending on either the sensory-delay condition and/or different variations of the task prompt.

Acknowledgements

We are thankful for the reviewers' comments. This material is based upon work supported by the National Science Foundation under Grant No. 1845322 and NSF Research Traineeship Grant No. 1735095, "Interdisciplinary Training in Complex Networks and Systems."

References

- Auvray, M., Lenay, C., and Stewart, J. (2009). Perceptual interactions in a minimalist virtual environment. *New Ideas in Psychology*, 27(1):32–47.
- Auvray, M. and Rohde, M. (2012). Perceptual crossing: the simplest online paradigm. *Frontiers in Human Neuroscience*, 6:181.
- Barone, P., Bedia, M. G., and Gomila, A. (2020). A minimal turning test: reciprocal sensorimotor contingencies for interaction detection. *Frontiers in Human Neuroscience*, 14:102.
- Beer, R. D. (1995). On the dynamics of small continuous-time recurrent neural networks. *Adaptive Behavior*, 3(4):469–509.
- Beer, R. D. (1996). Toward the evolution of dynamical neural networks for minimally cognitive behavior. *From Animals to Animats*, 4:421–429.
- De Jaegher, H., Di Paolo, E., and Gallagher, S. (2010). Can social interaction constitute social cognition? *Trends in Cognitive Sciences*, 14(10):441–447.
- Di Paolo, E. A. (2000). Behavioral coordination, structural congruence and entrainment in a simulation of acoustically coupled agents. *Adaptive Behavior*, 8(1):27–48.
- Di Paolo, E. A., Rohde, M., and Iizuka, H. (2008). Sensitivity to social contingency or stability of interaction? modelling the dynamics of perceptual crossing. *New Ideas in Psychology*, 26(2):278–294.
- Froese, T. and Di Paolo, E. A. (2008). Stability of coordination requires mutuality of interaction in a model of embodied agents. In *International Conference on Simulation of Adaptive Behavior*, pages 52–61. Springer.
- Froese, T. and Di Paolo, E. A. (2009). Toward minimally social behavior: social psychology meets evolutionary robotics. In *European Conference on Artificial Life*, pages 426–433. Springer.
- Froese, T. and Di Paolo, E. A. (2010). Modelling social interaction as perceptual crossing: an investigation into the dynamics of the interaction process. *Connection Science*, 22(1):43–68.
- Froese, T. and Di Paolo, E. A. (2011). The enactive approach: Theoretical sketches from cell to society. *Pragmatics & Cognition*, 19(1):1–36.
- Froese, T., Iizuka, H., and Ikegami, T. (2014). Embodied social interaction constitutes social cognition in pairs of humans: a minimalist virtual reality experiment. *Scientific Reports*, 4(1):1–10.
- Froese, T., Zapata-Fonseca, L., Leenen, I., and Fossion, R. (2020). The feeling is mutual: clarity of haptics-mediated social perception is not associated with the recognition of the other, only with recognition of each other. *Frontiers in Human Neuroscience*, 14:373.
- Hermans, K. S., Kasanova, Z., Zapata-Fonseca, L., Lafit, G., Fossion, R., Froese, T., and Myin-Germeyns, I. (2020). Investigating real-time social interaction in pairs of adolescents with the perceptual crossing experiment. *Behavior Research Methods*, 52(5):1929–1938.
- Iizuka, H. and Di Paolo, E. A. (2007). Minimal agency detection of embodied agents. In *European Conference on Artificial Life*, pages 485–494. Springer.
- Iizuka, H. and Ikegami, T. (2004). Adaptability and diversity in simulated turn-taking behavior. *Artificial Life*, 10(4):361–378.
- Iizuka, H., Saitoh, S., Marocco, D., and Yamamoto, M. (2015). Time delay effect on social interaction dynamics. In *Proceedings of the 3rd International Conference on Human-Agent Interaction*, pages 217–219.
- Ikegami, T. and Iizuka, H. (2007). Turn-taking interaction as a co-operative and co-creative process. *Infant Behavior and Development*, 30(2):278–288.
- Lenay, C. and Stewart, J. (2012). Minimalist approach to perceptual interactions. *Frontiers in Human Neuroscience*, 6:98.
- Lenay, C., Stewart, J., Rohde, M., and Amar, A. A. (2011). “you never fail to surprise me”: The hallmark of the other: Experimental study and simulations of perceptual crossing. *Interaction Studies*, 12(3):373–396.
- Quinn, M. (2001). Evolving communication without dedicated communication channels. In *European Conference on Artificial Life*, pages 357–366. Springer.
- Reséndiz-Benhumea, G. M. and Froese, T. (2020). Enhanced neural complexity is achieved by mutually coordinated embodied social interaction: A statespace analysis. In *Second International Workshop on Agent Based Modelling of Human Behaviour*.
- Reséndiz-Benhumea, G. M., Sangati, E., Sangati, F., Keshmiri, S., and Froese, T. (2021). Shrunken social brains? a minimal model of the role of social interaction in neural complexity. *Frontiers in Neurorobotics*, 15:72.
- Rohde, M. (2010). *Enaction, embodiment, evolutionary robotics: simulation models for a post-cognitivist science of mind*, volume 1. Springer Science & Business Media.
- Rohde, M. and Paolo, E. D. (2008). Embodiment and perceptual crossing in 2d. In *International Conference on Simulation of Adaptive Behavior*, pages 83–92. Springer.
- Schilbach, L., Timmermans, B., Reddy, V., Costall, A., Bente, G., Schlicht, T., and Vogeley, K. (2013). Toward a second-person neuroscience I. *Behavioral and Brain Sciences*, 36(4):393–414.
- Williams, P. L., Beer, R. D., and Gasser, M. (2008). Evolving referential communication in embodied dynamical agents. In *Eleventh International Conference on the Simulation and Synthesis of Living Systems*, pages 702–709.
- Zapata-Fonseca, L., Froese, T., Schilbach, L., Vogeley, K., and Timmermans, B. (2018). Sensitivity to social contingency in adults with high-functioning autism during computer-mediated embodied interaction. *Behavioral Sciences*, 8(2):22.