Collective Decision-Making in Coupled Echo State Networks

Polyphony Bruna (pbruna@ucmerced.edu)[1†], Rick Dale (rdale@ucla.edu)[2],
Michael Spivey (spivey@ucmerced.edu)[1] & Christopher Kello (ckello@ucmerced.edu)[1]

[1]Department of Cognitive and Information Sciences, University of California, Merced
[2]Department of Communication, University of California, Los Angeles
[†]Corresponding author

Abstract

Interaction can sometimes hinder human collective performance, while aggregating independent responses ("wisdom of crowds") is expected to improve performance through statistical facilitation. Yet some coordination studies show that interacting pairs can outperform their individual members and nominal groups (Bahrami et al., 2010; Szary & Dale, 2013, 2014). We present a reservoir computing model demonstrating how interaction can enhance joint performance. Two echo state networks were trained independently to identify Japanese speakers from vowel recordings and tested jointly with and without coupled feedback. The model reproduces the dyadic advantage reported by Bahrami et al. (2010) with interaction via pooled-feedback, but not with independent self-feedback. Model dynamics revealed that interaction was beneficial when uncoupled outputs were prone to diverging from runaway feedback over time. Our results offer a formal framework for explaining when and why interaction enhances collective intelligence.

Keywords: collective intelligence; interaction dynamics; perceptual decision-making; recurrent neural networks; reservoir computing

Introduction

Human interaction sometimes helps and sometimes hinders collective performance. There are open questions about the underlying mechanisms driving these mixed observations. In some cases, interaction may complicate coordination (Dreyer et al., 2025). Other times, interaction between two people can achieve greater outcomes than either of their best individual performances (Bahrami et al., 2010). Yet other work suggests that independent aggregation, without interaction at all, may yield the best statistical outcome (Surowiecki, 2005). One approach to this puzzle is to investigate the underlying dynamics that interaction introduces to these collective decisions.

When humans interact they can exhibit interpersonal synchronization across multiple biomechanical and behavioral dimensions, including gait, posture, physiology, gaze, head movement, and speech (Abney et al., 2014; Branigan, 2007; Brennan, 1996; Felsberg & Rhea, 2021; Pardo, 2006; Richardson & Dale, 2005; Shockley et al., 2003; Trujillo et al., 2023; Yi & Palmer, 2025). Although the functional role of synchrony is debated, it has been argued to support joint action, prosociality, mutual understanding, and the formation of linguistic conventions (Dale et al., 2013; daSilva & Wood, 2025; Hawkins et al., 2023; Hove & Risen, 2009; Louwerse et al., 2012; Stolk et al., 2016). Consistent with this view, interpersonal synchrony has been linked to successful group problem-solving across multiple tasks (Abney et al., 2021; Dideriksen et al., 2023). However, classic findings in group performance show that interaction can impair collective outcomes relative to equivalent non-interacting groups, such as by reducing the number of unique ideas generated during category recall (Szary et al., 2015; Weldon & Bellinger, 1997), suggesting that interaction may suppress the benefits of independent exploration (Janis, 2008; Rotter & Portugal, 1969).

By contrast, aggregating independent guesses has been shown to produce a collective response that is more accurate than any individual guess, what is sometimes called the "wisdom of crowds" (Surowiecki, 2005). This effect leverages a general benefit of statistical facilitation in certain problem spaces and was first demonstrated by pooling guesses of the weight of an ox (Galton, 1907). Empirical illustrations of this statistical effect in general knowledge domains show that just a few pooled guesses from independent individuals can improve precision (Ariely et al., 2000). These effects of statistical independence are sufficiently robust that some research has shown that they can hold across multiple guesses from a single person (Vul & Pashler, 2008). Nevertheless, the potential for bias within underlying distributions of individuals can raise concerns about how to achieve optimal outcomes (Becker et al., 2017) and the effect is still subject to factors such as relative level of knowledge across a population (Bang & Frith, 2017).

Although the general advantage of the wisdom of crowds is well-known, and cited as a canonical example of collective intelligence, it does not shed light on the cognitive role of interaction in human coordination. In the present study, we investigate the finding that interacting dyads demonstrate a collective advantage in a joint perceptual decision-making task (Bahrami et al., 2010; Bang et al., 2014). We aim to shed light on how interaction can improve performance beyond the baseline benefit of statistical facilitation.

Bahrami et al. (2010) showed pairs of participants sets of Gabor patches (sinusoidal grating patterns with a Gaussian blur) in two consecutive intervals. Participants were tasked with identifying which interval, the first or the second, contained a contrast oddball. Individual responses were collected on each trial. When responses differed, participants conferred with each other to reach a joint decision. Individual and joint decisions across a range of contrast differences were fit with psychometric curves representing the perceptual sensitivity of

In G. Boleda, I. Dautriche, M. Maldonado, C. Saldana and T. Torrent (Eds.), *Proceedings of the 48th Annual Conference of the Cognitive Science Society*

each member and the dyad.

The study showed that dyad discriminability exceeded the perceptual sensitivity of its best member. Furthermore, the authors showed that this dyadic advantage was modulated by differences in individual member discriminability: Dyads composed of members with approximately similar perceptual sensitivities exhibited the greatest dyadic gain, and this gain diminished as the disparity between individual ability increased. The authors argued that this inverse linear relationship is consistent with the hypothesis that dyads communicated their level of confidence in their guesses to produce joint decisions.

There is an open question about the process by which dyads shared their confidences during joint decision-making. On the one hand, communication may only serve to integrate guesses after individual decision-making has terminated. In this case, the dyadic advantage would reflect the same baseline statistical facilitation also leveraged by wisdom of crowds. On the other hand, individual decision-making may continue during communication, and interaction between the decision processes themselves may produce an emergent benefit beyond what is attained by statistical facilitation. These theoretical alternatives are illustrated in Figure 1b.

An approach to answering open questions of this kind is to build computational models that implement well defined and parameterized processes to test potential underlying mechanisms systematically (McClelland, 2009). This is the approach we take here. We devised a computational architecture that offers a window into the role of collective information processes. The architecture uses a recurrent neural network model with two key benefits. First, these networks are flexible, allowing definition of various kinds of information pooling. Second, a recurrent neural network model permits visualization of the dynamics of unfolding decisions, helping to uncover the role of pooled or independent information. In our simulations, we reproduce the dyadic advantage observed in Bahrami et al. (2010) using a pair of coupled echo state networks (ESNs), a prominent recurrent neural network framework well suited to studying dynamic processes. In doing so, we shed light on the mechanisms underlying the original effect, in particular, and examine the functional role of interaction in collective intelligence, in general.

We first present a novel dyadic ESN architecture, which may be tested either as an aggregate of individual processes (*qua* wisdom of crowds) or as a coupled information-processing system (*qua* dyadic interaction). We then test this model on a canonical signal classification task to evaluate collective performance with and without interaction. We show that interaction via pooled-feedback over the course of processing produces a dyadic advantage similar to the original study. This effect is diminished with independent (self-) feedback, showing that interaction yielded a benefit beyond statistical facilitation. Lastly, we analyze the network decision trajectories and show that pooled-feedback stabilizes network dynamics over time in unstable agents, allowing coupled networks to exhibit a collective benefit when they process the task as a dyad.

Model

ESNs belong to a family of recurrent neural networks called *reservoir computers* (Jaeger, 2001). Reservoir computers leverage the dynamics of a fixed, recurrent network of artificial neurons to perform general-purpose computations. When the dynamical memory of the reservoir balances signal retention with signal separation the resulting nonlinear embeddings may be used to "read-out" properties of input signals, analogous to how biological intelligences may flexibly exploit naturally-occurring intrinsic dynamics for cognition without requiring task-specific wiring (Maass et al., 2002).

Crucially only the parameters of the readout layer, which projects reservoir states onto task-relevant targets, are trained, classically using a least mean squares algorithm in one pass. As a result, reservoir computers offer simple, dynamical models of computation that have been applied to modeling the brain (Enel et al., 2016), chaotic time series forecasting (Jaeger & Haas, 2004; Pathak et al., 2018), and speech recognition (Verstraeten et al., 2006). We employ ESNs as a dynamical model of decision-making that can be readily extended to collective processing among two or more networks.

Speaker Classification Task

We trained two ESNs to classify speech signals from the Japanese Vowels dataset (Kudo et al., 1999), consisting of 640 time series (270 training; 370 testing) of nine native, male Japanese speakers pronouncing the vowels /ae/ repeatedly. Speech recordings were cepstrum-coded to produce 12 linear prediction coefficients (LPCs) that encode the auditory signal on each time step. Targets were coded as nine-bit one-hot vectors representing classification decisions for each of the nine possible speakers.

Signals were truncated to 10 time steps, and signals with fewer than 10 time steps were dropped. Signals were then selected at random such that there were an equal number of signals per class in both the training (28 per class) and testing (24 per class) data. Elements in the input and target vectors were standardized with respect to the training data such that the total activation perturbing the reservoir on a given time step had an expected value of zero.

ESN Architecture

Each ESN was composed of three layers: An input layer, a recurrent network, called the reservoir, and an output layer, called the readout. The input layer was represented as a set of 12 nodes that deliver input signals, $u(t)$, to the reservoir. Connections from the input layer into the reservoir were defined by an input weight matrix, W^{in}, and input nodes were randomly connected to reservoir nodes with a probability of p_{inlink}. Input weights took on a fixed value of either 1 or -1 and were drawn with equal odds from a Bernoulli distribution ($p = 0.5$).

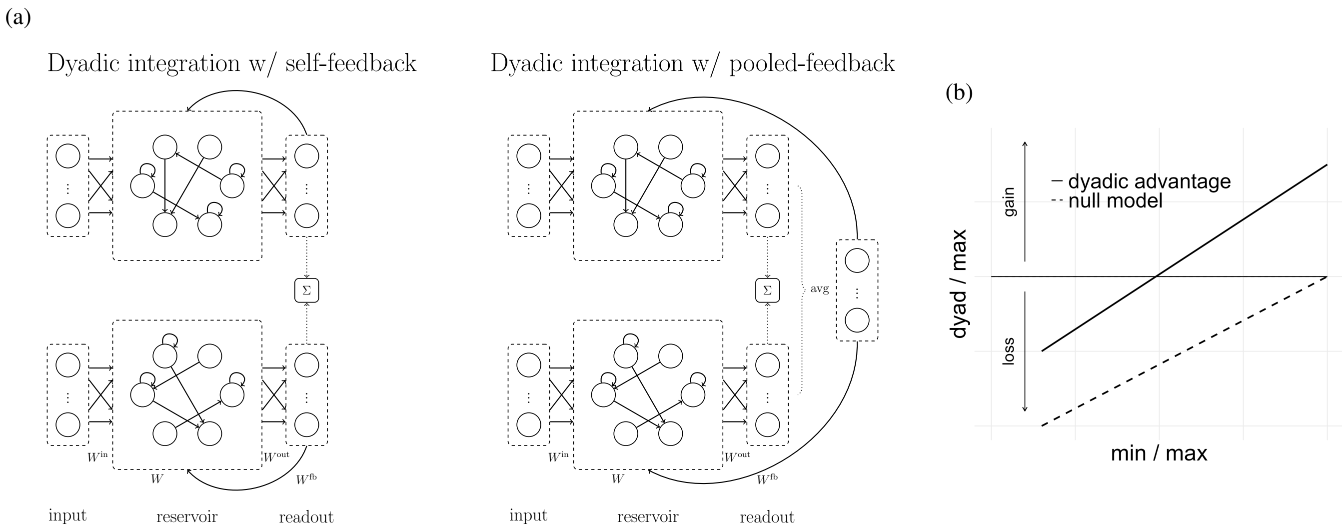

Figure 1: **(a)** Schematic of the dyadic ESN model in each condition for one time step. In the self-feedback condition, readout from independent ESNs is integrated node-wise across networks (shown by the summation box) as well as over time (not shown) in order to produce a collective response. In the pooled-feedback condition, ESNs are coupled through their feedback channels and each ESN receives the average of the two readouts as feedback. **(b)** Theoretical predictions based on Bahrami et al. (2010). The null model is defined by dyadic performance equal to the average of its members. We aim to determine how much of the dyadic advantage is accounted for by statistical facilitation via self-feedback versus coupling via pooled-feedback.

The reservoir was represented as a recurrent network of N_{res} nodes such that each node was connected to other nodes in the reservoir, including themselves, with a probability of $p_{reslink}$. Internal reservoir connections, W^{res}, took on weight values drawn from a Gaussian distribution ($\mu = 0, \quad \sigma = 1$). Reservoir weights were fixed at initialization, and the reservoir matrix was rescaled to obey a spectral radius of $sr = 0.9$. Reservoir dynamics, $x(t)$, updated according to Equation 1.

$$x(t+1) = \tanh(W^{in}u(t) + W^{res}x(t) + W^{fb}y(t)) \quad (1)$$

Lastly, the readout was represented as a set of nine nodes with each node representing a classification decision by the network. Readout nodes were fully connected to the reservoir, and weights were determined using the least mean squares algorithm with Tikhonov regularization (Tikhonov & Arsenin, 1977). Feedback connections from the readout to the reservoir sent prior network outputs, $y(t)$, back into the reservoir on each time step. Feedback connections were defined by the weight matrix, W^{fb}, and readout nodes were randomly connected to reservoir nodes with a probability of p_{inlink}. Like the input weights, feedback weights took on a fixed value of either 1 or -1 and were drawn with equal odds from a Bernoulli distribution ($p = 0.5$).

During training, the reservoir received teacher forcing as substitute feedback, where feedback is estimated using the target readout signal on each time step (Lukoševičius, 2007). Gaussian noise with a standard deviation of $\sigma_{teacher}$ was added to each unit of the teacher forced signal, which controls how the readout layer learns to interpret reservoir dynamics impacted by noisy feedback.

Dyadic Architecture

The dyadic ESN architecture is shown in Figure 1a. A joint decision was calculated from two trained ESNs by integrating the activation of their readout layers over time and across networks. At the end of each signal, the integrated readout node with the highest summed activation was taken as the decision of the dyadic network (Equation 2).

$$\text{Joint Decision} = \arg\max\left(\sum_{r=1}^{2}\sum_{t=1}^{10} y_{r,t}\right) \quad (2)$$

Networks were tested according to two conditions: Dyadic integration with self-feedback and dyadic integration with pooled-feedback. In the self-feedback condition, networks only received feedback from their own readout layer, essentially completing the task independently. In the pooled-feedback condition, networks received the average of the two readout layers as feedback on each time step, allowing reservoir dynamics to influence each other during the processing of each signal. In both conditions, a joint decision was calculated on each test signal. This allowed us to collect two different measures of collective performance from the model.

Parameter Space

Individual ESN performance on the speaker classification task varied along two key parameters defining the model space: Size of the reservoir (N_{res}) and noise added to teacher forcing ($\sigma_{teacher}$). Performance within this model space averaged over 20 simulations is shown in Figure 2a. Performance varied nonlinearly, with accuracy on the task reaching ceiling with moderate N_{res} and $\sigma_{teacher}$ (unreported simulations showed a comparable trend without feedback). This pattern did not change substantially by varying the internal connectivity of the reservoir nor the input connectivity into the reservoir, which, consequently, remained fixed ($p_{reslink} = 0.1$, $p_{inlink} = 0.1$).

Figure 2a also illustrates dyadic performance (averaged over 10 simulations) over the best member for both conditions in this model space. These results show a small dyadic advantage in most of the model space for both independent and coupled networks and a substantial advantage in coupled networks localized to the bottom-right corner, corresponding to a region where individual performance drops off. To investigate this difference in collective performance, we sampled ESNs randomly from $N_{res} \in [800, 1600]$ and $\sigma_{teacher} \in [0.2, 0.4]$ to generate heterogeneous dyadic networks and used those dyads to simulate results from Bahrami et al. (2010) shown in Figure 2b. In line with the original study, Gaussian noise was added to each input node ($\sigma_{input} = 2$), either for both networks ("equal noise") or one network ("unequal noise"). Ambient noise was simulated in the "no noise" condition ($\sigma_{input} = 0.5$).

Results

Model simulations were produced using ReservoirPy (Trouvain et al., 2020), and analyses were performed in the R environment. Results were averaged over 10 simulations for dyads, or 20 simulations for individual networks. Performance was measured as proportion correct on test signals.

Coupled ESNs Reproduce Dyadic Advantage

We examined the relationship between member disparity and collective performance for 60 simulated dyads (20 per noise condition). The ratio between the performance of each member of a dyad measures their difference in individual performance. The ratio between dyadic performance and the performance of the better member measures collective benefit.

Figure 2b shows the relationship between member disparity and collective benefit under the self- and pooled-feedback conditions. Both conditions display an inverse relationship between collective benefit and member disparity, indicating that dyads with members that were matched on individual performance tended to perform better collectively. A mixed-effects linear regression with dyad as a random effect revealed a significant interaction between feedback condition and member disparity, indicating that the effect of member disparity on collective benefit was larger with pooled-feedback than self-feedback ($\beta = 0.29, p = 0.016$). Estimated marginal means (EMMs) from this regression were used to test collective benefit at various levels of member disparity in each condition.

Consistent with prior findings from Bahrami et al. (2010), pooled-feedback showed a significant collective benefit for dyads with matched members (EMM = 0.57, $p < 0.0001$) that persisted until the worse member performed roughly half as well as their partner, after which dyads began to underperform their best member. In contrast, the magnitude of this effect was reduced with self-feedback. In matched dyads, self-feedback showed a significant, but smaller, collective benefit (EMM = 0.21, $p < 0.0001$). In closely matched dyads ($P_{min}/P_{max} = 0.75$), pooled-feedback still showed a significant improvement over the best member (EMM = 0.25, $p < 0.0001$), but self-feedback did not (EMM = −0.037, $p = 0.22$).

In the original study, dyadic performance approximated that of the best member when the ratio of the worse to better member was roughly 0.42, which is more closely fit by pooled-feedback (0.55) compared to self-feedback (0.79). Most dyads with self-feedback exhibited a collective detriment relative to their best member (72%), compared to pooled-feedback (53%) and the original study (32%). Together, these findings indicate that the pooled-feedback condition offers a better qualitative fit to Bahrami et al. (2010), suggesting that the dyadic advantage depends on interaction between members during decision-making.

Interaction Stabilizes Network Dynamics

Next we compared the dynamics in the readout layer over the time course of signal processing with self- and pooled-feedback to understand how interaction alters the decision process of networks in the dyadic model. Figure 2c shows time series for each of the nine decisions averaged over signals. Time series are compiled such that the first trajectory corresponds to the correct response and the following trajectories indicate competing decisions ordered by rank.

With self-feedback, networks failed to retain the correct classification decisions until the end of a signal, demonstrated by trajectories that increasingly diverge from their target activation as signals persist. However, with pooled-feedback this divergence was mitigated. Networks maintained more stable decision trajectories over time, shown by a delay in how quickly, and to what extent, each trajectory diverges in the coupled networks. The first-order temporal difference (Equation 3) averaged over test signals measures the dispersion of the trajectories over time in each condition (Figure 2d). A linear mixed-effects regression model with condition as a predictor of dispersion and random effects of node, signal, and time step confirmed that trajectories with pooled-feedback dispersed significantly less over time compared to self-feedback ($\beta = -0.13, p < 0.0001$).

$$\Delta(t)_{act} = \frac{1}{9}\sum_{i=1}^{9} |y_i(t) - y_i(t-1)| \tag{3}$$

These findings suggest that integrating the feedback of multiple networks over time helps coupled networks maintain more stable dynamics compared to independent systems. Comparing these dynamics to networks that are more robust to perturbation in their feedback signal (modulated by increasing $\sigma_{teacher}$) revealed that the effect of pooled-feedback vanishes ($\beta = -0.0015, p = 0.15$), suggesting that the dyadic advantage arises when agents are prone to runaway positive feedback in their decision-making process. Interaction introduces a negative feedback dynamic that stabilizes individual

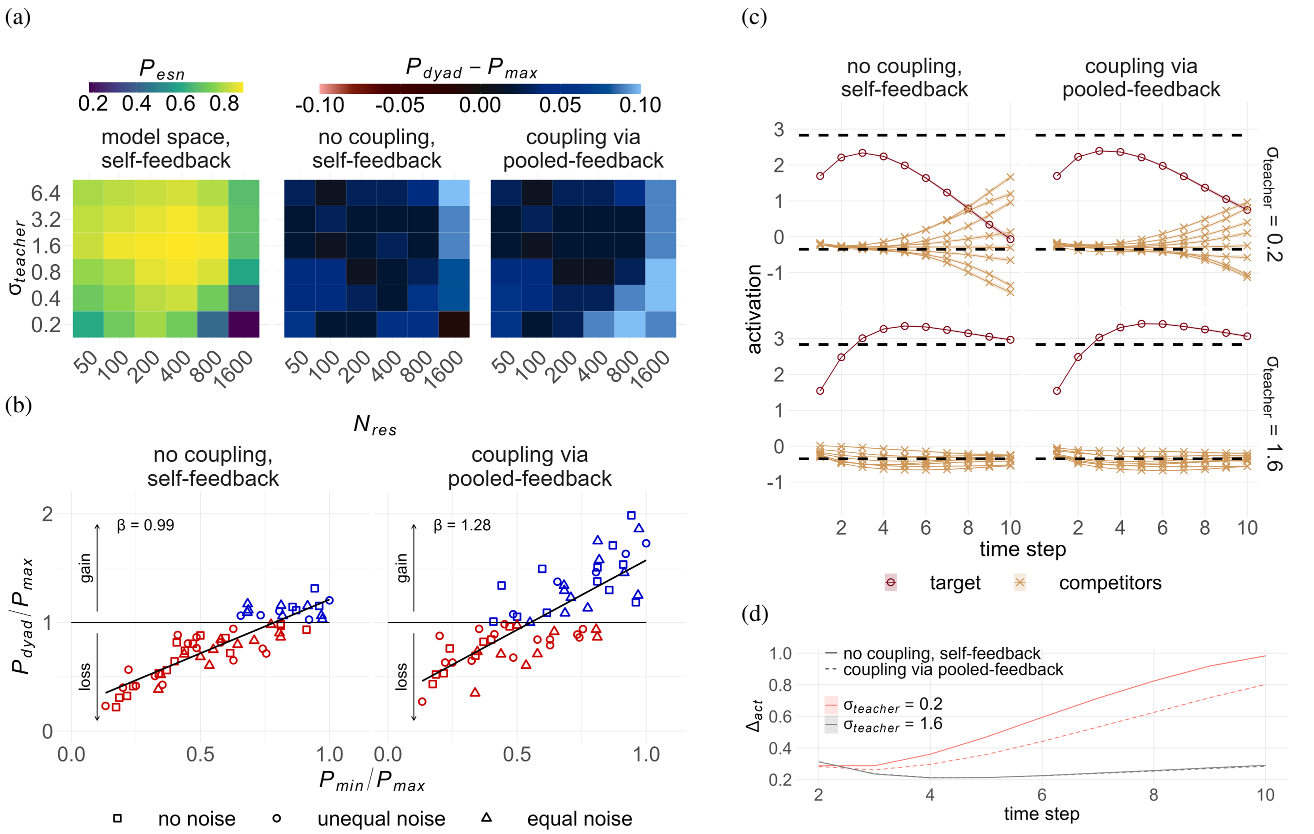

Figure 2: **(a)** Performance as a function of reservoir nodes (N_{res}) and noise on teacher forcing ($\sigma_{teacher}$). **(b)** Collective benefit compared to member disparity across both conditions. Values above P_{dyad}/P_{max}=1 indicate the dyad outperforms its best member and values below P_{dyad}/P_{max}=1 indicate the dyad underperforms its best member. **(c)** Time series of readout node activation across both conditions for two levels of $\sigma_{teacher}$. Incorrect trajectories are averaged to preserve the rank-order of the decision. Dashed lines indicate the target activation value for the correct and incorrect class. **(d)** Absolute first-order temporal difference averaged over readout nodes for each condition and two levels of $\sigma_{teacher}$. Ribbons indicate standard error.

network behavior, enabling collectives to outperform even their best individuals.

General Discussion

In this paper, we ask what mechanisms enable humans to perform better when deliberating together versus independently. We offer a novel model of joint decision-making using coupled ESNs that reproduces empirical findings that dyads can outperform their best individual member. Importantly, this finding was not replicated in collective decisions from independent networks, revealing that the advantage is an effect of interaction between the two networks. Our results provide a computational foundation for understanding when interaction benefits collective intelligence and why.

Interaction and Crowd Wisdom

We show that interaction mitigates the destabilization of network dynamics by reducing endogenous noise introduced through feedback. Networks that pool feedback outperform their independent counterparts, producing a dyadic advantage.

In our model, pooled feedback offers a simple mechanism to fulfill the function of weighted confidence-sharing, proposed by Bahrami et al. (2010). Each network's readout vector encodes both its decision (most active node) and its graded confidence across alternatives (magnitude of node activations) on each time step. This interpretation is straightforward because networks learn readouts with zero-mean activation. When readout vectors are combined, low-magnitude activations contribute little, while high-magnitude activations strongly bias the pooled signal towards or against certain outcomes. Networks with higher confidence therefore exert greater influence on the joint feedback.

Weighted confidence-sharing does not specify how interaction alters decision-processing over time. Individual guesses could be combined after independent decision-making has terminated, which would predict that integrating decisions across independent networks would show the same dyadic advantage as coupled networks, or better. Instead, we find that pooling feedback during signal processing better reproduces prior findings, suggesting that the collective benefit emerges from interaction shaping the dynamics of decision-making itself.

Our findings align with work showing that structured interaction can enhance the benefit of crowd wisdom. One study found that small-group consensus can outperform simple aggregation of individual estimates when participants deliberate independently before interacting (Navajas et al., 2018). Likewise, we observe that while aggregating independent network decisions yields a modest collective benefit, this benefit increases in both magnitude and scope when pooled feedback is introduced.

Recent research showing that greater variance among individual judgments improves crowd wisdom (Barrera-Lemarchand et al., 2024) suggests that independent processing prior to interaction reduces stifling effects. In our model, ESNs are trained independently before testing, analogous to the individual deliberation phases used by Bahrami et al. (2010) and Navajas et al. (2018).

Finally, studies demonstrate that network structure and response weighting can amplify or suppress crowd wisdom (Baten & Mahjabin, 2025; Becker et al., 2017). Our model offers a framework for examining this phenomena at the mechanistic level. Extending our model to larger social networks under various coupling rules may help bridge this work in computational social science with underlying cognitive mechanisms.

Collective Stabilization

Networks that are trained to be more robust to perturbations in their feedback signals do not exhibit an effect of interaction, revealing that the dyadic advantage depends on collective decision-making in agents with unstable dynamics. Such instability arises from runaway positive feedback. In decision-making, this can manifest as "overthinking," such as when a student revises a correct multiple-choice answer because deliberation continues after sufficient evidence has been reached. When unstable dynamics persist, incorrect information can be amplified over time.

Positive feedback is common in natural and biological systems, enabling exploration and adaptation but risking instability when unchecked (Brandman & Meyer, 2008; May, 1972). In humans, for instance, attempts to maintain a constant vocal volume often lead to unintentional amplification (Zollinger & Brumm, 2011), while in the brain excessive positive feedback can produce seizures (Huguenard & McCormick, 2007; Jirsa et al., 2014). Negative feedback stabilizes neural dynamics, supporting homeostasis, and disruptions in the balance of positive and negative feedback contribute to psychiatric disorders like schizophrenia (Molina et al., 2020). Neural computation is thought to emerge near this balance at the "edge of chaos" (Hengen & Shew, 2025; Kello, 2013), a regime also exploited by reservoir computing to maximize both stability and flexibility (Bertschinger & Natschläger, 2004).

Our work demonstrates that balancing positive and negative feedback can emerge from the interplay of individual and collective dynamics. Simple agents that are individually susceptible to positive feedback can achieve stable performance by leveraging negative feedback loops introduced by dyadic interaction, highlighting one computational solution available to social agents. The possibility that collectives of relatively "cheap" agents, with respect to computational costs, can equal or outperform an expensive individual suggests that leveraging collective dynamics may be favored on the grounds of computational efficiency from an evolutionary perspective. Further work examining the tradeoff between efficiency and performance in collectives versus individuals is needed to explore this potentiality.

Conclusion

We present a minimal model that reproduces findings from Bahrami et al. (2010), in particular, and explores the cognitive mechanisms underlying productive interaction in collectives, in general. One limitation of this work is the use of a nine-choice signal classification task, which differs from the task used in the original study. Consequently, we measure performance as proportion correct instead of estimating discriminability from psychometric curves. Although the role of interaction discussed in Bahrami et al. (2010), and explored in the present work, is mathematically general, future work should examine how task variation influences these results.

Our work also contributes to multi-agent reservoir computing, which has attracted modest attention under the name "ensemble reservoirs" (Casanova et al., 2023). For example, Hu et al. (2022) implement a similar coupling mechanism in reservoirs trained on chaotic attractors to model synchronization between oscillators. Our findings suggest that such coupling mechanisms may yield performance gains in time series forecasting on these chaotic signals, a topic to be explored in further research.

Furthermore, negative correlation learning (Liu & Yao, 1999), a technique in which neural networks are trained to produce negatively correlated outputs, has been shown to yield better performance in ensemble reservoirs (Rodan & Tino, 2011), mirroring the finding that crowd wisdom improves in maximally divergent crowds (Barrera-Lemarchand et al., 2024). Alongside the work presented here, this research points to an intersection between multi-agent machine learning and computational social science as well as motivates reservoir computing as a promising approach for modeling interpersonal cognitive dynamics.

Acknowledgments

Scripts, data, and other materials are available at: https://github.com/pjbruna/coupled-esns. Part of this research was performed while Polyphony Bruna was visiting the Institute for Pure and Applied Mathematics, which is supported by the National Science Foundation (Grant No. DMS-1925919), and the Santa Fe Institute. This project also received support from the Institute for Humane Studies (Grant Nos. IHS019149 and IHS018920) and from a UC Merced Cognitive and Information Sciences Fellowship to Polyphony Bruna. AI was not used in the preparation of this project and manuscript.

References

Abney, D. H., Paxton, A., Dale, R., & Kello, C. T. (2014). Complexity matching in dyadic conversation. *Journal of Experimental Psychology: General*, *143*(6), 2304.

Abney, D. H., Paxton, A., Dale, R., & Kello, C. T. (2021). Cooperation in sound and motion: Complexity matching in collaborative interaction. *Journal of Experimental Psychology: General*, *150*(9), 1760.

Ariely, D., Tung Au, W., Bender, R. H., Budescu, D. V., Dietz, C. B., Gu, H., Wallsten, T. S., & Zauberman, G. (2000). The effects of averaging subjective probability estimates between and within judges. *Journal of Experimental Psychology: Applied*, *6*(2), 130.

Bahrami, B., Olsen, K., Latham, P. E., Roepstorff, A., Rees, G., & Frith, C. D. (2010). Optimally interacting minds. *Science*, *329*(5995), 1081–1085.

Bang, D., & Frith, C. D. (2017). Making better decisions in groups. *Royal Society open science*, *4*(8), 170193.

Bang, D., Fusaroli, R., Tylén, K., Olsen, K., Latham, P. E., Lau, J. Y., Roepstorff, A., Rees, G., Frith, C. D., & Bahrami, B. (2014). Does interaction matter? testing whether a confidence heuristic can replace interaction in collective decision-making. *Consciousness and cognition*, *26*, 13–23.

Barrera-Lemarchand, F., Balenzuela, P., Bahrami, B., Deroy, O., & Navajas, J. (2024). Promoting erroneous divergent opinions increases the wisdom of crowds. *Psychological Science*, *35*(8), 872–886.

Baten, R. A., & Mahjabin, M. R. (2025). The wisdom of intellectually humble networks. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *47*.

Becker, J., Brackbill, D., & Centola, D. (2017). Network dynamics of social influence in the wisdom of crowds. *Proceedings of the National Academy of Sciences*, *114*(26), E5070–E5076.

Bertschinger, N., & Natschläger, T. (2004). Real-time computation at the edge of chaos in recurrent neural networks. *Neural Computation*, *16*(7), 1413–1436. https://doi.org/10.1162/089976604323057443

Brandman, O., & Meyer, T. (2008). Feedback loops shape cellular signals in space and time. *Science*, *322*(5900), 390–395.

Branigan, H. (2007). Syntactic priming. *Language and Linguistics Compass*, *1*(1-2), 1–16. https://doi.org/https://doi.org/10.1111/j.1749-818X.2006.00001.x

Brennan, S. E. (1996). Lexical entrainment in spontaneous dialog. *Proceedings of ISSD*, *96*, 41–44.

Casanova, M., Dalena, B., Bonaventura, L., & Giovannozzi, M. (2023). Ensemble reservoir computing for dynamical systems: Prediction of phase-space stable region for hadron storage rings. *The European Physical Journal Plus*, *138*(6), 559.

Dale, R., Fusaroli, R., Duran, N. D., & Richardson, D. C. (2013). The self-organization of human interaction. In B. H. Ross (Ed.), *Psychology of learning and motivation* (pp. 43–95, Vol. 59). Academic Press. https://doi.org/10.1016/B978-0-12-407187-2.00002-2

daSilva, E. B., & Wood, A. (2025). How and why people synchronize: An integrated perspective. *Personality and Social Psychology Review*, *29*(2), 159–187.

Dideriksen, C., Christiansen, M. H., Tylén, K., Dingemanse, M., & Fusaroli, R. (2023). Quantifying the interplay of conversational devices in building mutual understanding. *Journal of Experimental Psychology: General*, *152*(3), 864.

Dreyer, T., Haluts, A., Korman, A., Gov, N., Fonio, E., & Feinerman, O. (2025). Comparing cooperative geometric puzzle solving in ants versus humans. *Proceedings of the National Academy of Sciences*, *122*(1), e2414274121.

Enel, P., Procyk, E., Quilodran, R., & Dominey, P. F. (2016). Reservoir computing properties of neural dynamics in prefrontal cortex. *PLoS computational biology*, *12*(6), e1004967.

Felsberg, D. T., & Rhea, C. K. (2021). Spontaneous interpersonal synchronization of gait: A systematic review. *Archives of Rehabilitation Research and Clinical Translation*, *3*(1), 100097. https://doi.org/https://doi.org/10.1016/j.arrct.2020.100097

Galton, F. (1907). Vox populi. *Nature*, *75*, 450–451.

Hawkins, R. D., Franke, M., Frank, M. C., Goldberg, A. E., Smith, K., Griffiths, T. L., & Goodman, N. D. (2023). From partners to populations: A hierarchical bayesian account of coordination and convention. *Psychological Review*, *130*(4), 977.

Hengen, K. B., & Shew, W. L. (2025). Is criticality a unified setpoint of brain function? *Neuron*, *113*(16), 2582–2598.

Hove, M. J., & Risen, J. L. (2009). It's all in the timing: Interpersonal synchrony increases affiliation. *Social cognition*, *27*(6), 949–960.

Hu, W., Zhang, Y., Ma, R., Dai, Q., & Yang, J. (2022). Synchronization between two linearly coupled reservoir computers. *Chaos, Solitons & Fractals*, *157*, 111882.

Huguenard, J. R., & McCormick, D. A. (2007). Thalamic synchrony and dynamic regulation of global forebrain oscillations. *Trends in neurosciences*, *30*(7), 350–356.

Jaeger, H. (2001). *The "echo state" approach to analysing and training recurrent neural networks* (tech. rep. No. 148). German National Research Center for Information Technology (GMD).

Jaeger, H., & Haas, H. (2004). Harnessing nonlinearity: Predicting chaotic systems and saving energy in wireless communication. *science*, *304*(5667), 78–80.

Janis, I. L. (2008). Groupthink. *IEEE Engineering Management Review*, *36*(1), 36–36. https://doi.org/10.1109/EMR.2008.4490137

Jirsa, V. K., Stacey, W. C., Quilichini, P. P., Ivanov, A. I., & Bernard, C. (2014). On the nature of seizure dynamics. *Brain*, *137*(8), 2210–2230.

Kello, C. T. (2013). Critical branching neural networks. *Psychological review*, *120*(1).

Kudo, M., Toyama, J., & Shimbo, M. (1999). Japanese vowels. *UCI Machine Learning Repository*. https://doi.org/10.24432/C5NS47

Liu, Y., & Yao, X. (1999). Ensemble learning via negative correlation. *Neural Networks*, *12*(10), 1399–1404. https://doi.org/https://doi.org/10.1016/S0893-6080(99)00073-8

Louwerse, M. M., Dale, R., Bard, E. G., & Jeuniaux, P. (2012). Behavior matching in multimodal communication is synchronized. *Cognitive science*, *36*(8), 1404–1426.

Lukoševičius, M. (2007). *Echo state networks with trained feedbacks* (Technical Report No. 4). Jacobs University Bremen, School of Engineering and Science. https://www.ai.rug.nl/minds/uploads/tfbesn_iubtechreport.pdf

Maass, W., Natschläger, T., & Markram, H. (2002). Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural Computation*, *14*(11), 2531–2560.

May, R. M. (1972). Will a large complex system be stable? *Nature*, *238*(5364), 413–414.

McClelland, J. L. (2009). The place of modeling in cognitive science. *Topics in Cognitive Science*, *1*(1), 11–38.

Molina, J. L., Voytek, B., Thomas, M. L., Joshi, Y. B., Bhakta, S. G., Talledo, J. A., Swerdlow, N. R., & Light, G. A. (2020). Memantine effects on electroencephalographic measures of putative excitatory/inhibitory balance in schizophrenia. *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging*, *5*(6), 562–568.

Navajas, J., Niella, T., Garbulsky, G., Bahrami, B., & Sigman, M. (2018). Aggregated knowledge from a small number of debates outperforms the wisdom of large crowds. *Nature Human Behaviour*, *2*(2), 126–132.

Pardo, J. S. (2006). On phonetic convergence during conversational interaction. *The Journal of the acoustical society of America*, *119*(4), 2382–2393.

Pathak, J., Hunt, B., Girvan, M., Lu, Z., & Ott, E. (2018). Model-free prediction of large spatiotemporally chaotic systems from data: A reservoir computing approach. *Physical review letters*, *120*(2), 024102.

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.

Rodan, A., & Tino, P. (2011). Negatively correlated echo state networks. *European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN)*.

Rotter, G. S., & Portugal, S. M. (1969). Group and individual effects in problem solving. *Journal of Applied Psychology*, *53*(4), 338.

Shockley, K., Santana, M.-V., & Fowler, C. A. (2003). Mutual interpersonal postural constraints are involved in cooperative conversation. *Journal of Experimental Psychology: Human Perception and Performance*, *29*(2), 326.

Stolk, A., Verhagen, L., & Toni, I. (2016). Conceptual alignment: How brains achieve mutual understanding. *Trends in cognitive sciences*, *20*(3), 180–191.

Surowiecki, J. (2005). *The wisdom of crowds*. Vintage.

Szary, J., & Dale, R. (2013). Dyadic cooperation enhances retrieval and recall of crossword solutions. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *35*(35).

Szary, J., & Dale, R. (2014). Familiarity modulates the dynamics of collaborative inhibition in a trivia game. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *36*(36).

Szary, J., Kello, C., & Dale, R. (2015). Memory foraging in a spatial domain. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *37*.

Tikhonov, A. N., & Arsenin, V. Y. (1977). *Solutions of ill-posed problems*. Winston & Sons.

Trouvain, N., Pedrelli, L., Dinh, T. T., & Hinaut, X. (2020). ReservoirPy: an Efficient and User-Friendly Library to Design Echo State Networks. *ICANN 2020 - 29th International Conference on Artificial Neural Networks*. https://inria.hal.science/hal-02595026

Trujillo, J. P., Dideriksen, C., Tylén, K., Christiansen, M. H., & Fusaroli, R. (2023). The dynamic interplay of kinetic and linguistic coordination in danish and norwegian conversation. *Cognitive Science*, *47*(6), e13298.

Verstraeten, D., Schrauwen, B., & Stroobandt, D. (2006). Reservoir-based techniques for speech recognition. *The 2006 IEEE International Joint Conference on Neural Network Proceedings*, 1050–1053.

Vul, E., & Pashler, H. (2008). Measuring the crowd within: Probabilistic representations within individuals. *Psychological Science*, *19*(7), 645–647.

Weldon, M. S., & Bellinger, K. D. (1997). Collective memory: Collaborative and individual processes in remembering. *Journal of experimental psychology: Learning, memory, and cognition*, *23*(5), 1160.

Yi, W., & Palmer, C. (2025). Respiratory synchrony and individual differences causally influence dyadic interpersonal coordination. *bioRxiv*, 2025–05.

Zollinger, S. A., & Brumm, H. (2011). The lombard effect. *Current Biology*, *21*(16), R614–R615.