

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Bayesian rational memory model simulates temporal binding effect

Permalink

<https://escholarship.org/uc/item/8sr9m069>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 44(44)

Authors

Saad, Laura

Musolino, Julien

Hemmer, Pernille

Publication Date

2022

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Bayesian rational memory model simulates temporal binding effect

Laura Saad, Julien Musolino, and Pernille Hemmer

Department of Psychology, Rutgers University
152 Frelinghuysen Rd.
Piscataway, NJ 08854

Abstract

Temporal Binding (TB) is standardly regarded as an implicit measure of the sense of agency (Haggard, 2017). Though the TB effect is robust, an underlying mechanism has not been agreed upon (Hoerl et al., 2020). Here we propose a memory process as an explanation for the observed error in two publicly available datasets. We first replotted the data and found that on average, across both experiments, participants overestimate the length of the shortest timing interval and underestimate the longest interval, a classic regression to the mean pattern. Summary statistics extracted from the data from each experiment were then used as parameters in a simple Bayesian model of memory. Model simulations reproduced the behavioral data for almost all timing intervals and experimental trial-types across two experiments. Adjusting one of the parameters in the model (prior mean for actions) resulted in an improved qualitative fit. We suggest that other more likely sources of error, apart from experienced agency, may account for this result.

Keywords: temporal binding; Bayesian models of cognition; sense of agency; memory

Introduction

The sense of agency is defined as the feeling of control over our actions and their associated consequences (Haggard, 2017). Though this experience can be assessed explicitly by asking individuals to report their feelings of control over outcomes in their environment (Daprati et al., 1997; Metcalfe et al., 2013; Metcalfe & Greene, 2007; Spengler et al., 2009), it has also been assessed implicitly via measures like temporal (also sometimes referred to as “intentional”) binding. Temporal binding is the perceived subjective compression of the timing interval between a voluntary action and its associated outcome.

Temporal binding is often (though not exclusively – see Haggard et al. (2002)) measured using a timing estimation method known as interval estimation (e.g., Caspar et al., 2016; Engbert et al., 2008; Fereday et al., 2019; Obhi et al., 2013; Pfister et al., 2014; Seghezzi & Zapparoli, 2020; Zhao et al., 2016). In the prototypical voluntary action trial, participants are asked to perform an intentional voluntary **action**, e.g., a button press, which triggers a **tone** 250ms later. At the end of each trial, participants are asked to freely recall their best estimation (in milliseconds) of the length of the interval that elapsed between the two events. This trial-type is typically compared to a baseline, observational (a.k.a. passive) condition where participants (most commonly) hear two computer-generated tones separated by 250ms and are

again asked to recall the interval that elapsed between the two events. Participants, on average, show increased compression of the interval between events in the voluntary action condition compared to the baseline, which is referred to as the temporal binding effect.

Since the original report (Haggard et al., 2002), the effect has more than 200 reported replications across multiple distinct timing estimation methods (e.g., Barlas & Kopp, 2018; Berberian et al., 2012; Cavazzana et al., 2014; Takahata et al., 2012; see Tanaka et al., 2019 for a review). Several mechanisms have been proposed to explain this effect, though no consensus has been reached (Hoerl et al., 2020). Initially, binding was thought to be specific to voluntary actions and so a pre-reflective motor mechanism was proposed (Haggard et al., 2002). This account posits that an implicit “feeling of control” yields temporal binding, a process that is thought to be driven by a motor control system that does not require conscious reflection.

However, there have been reported instances where binding is present in passive/observational (Graham-Schmidt et al., 2016, other condition; Humphreys & Buehner, 2010; Poonian et al., 2015; Suzuki et al., 2019) as well as involuntary actions (Borhani et al., 2017; Buehner, 2015; Graham-Schmidt et al., 2016, passive condition; Kirsch et al., 2019). As a result, another account proposes that the perceived causality between two events can by itself elicit binding. Importantly, this account predicts that binding can occur with cause/effect pairs that do not involve intentionality or voluntary motor action (Hoerl et al., 2020).

More recently, cue integration, borrowed from perception research (Lush et al., 2019; Moore & Fletcher, 2012), has been suggested as a potential mechanism underlying the binding effect. Cue integration is the process by which an observer combines information associated with multiple cues from different domains to decrease error in perception. The cues are weighted differentially based on the perceptual certainty associated with each. For example, in the voluntary action case, the timing of the button press has more perceptual certainty compared to the timing of the first tone in the baseline trials (i.e., participants can control the timing of their voluntary actions which increases certainty compared to the occurrence of the first tone in the baseline condition which they do not control). Therefore, the estimation of the outcome is pulled toward the event with higher perceptual certainty which can explain the observed compression seen in binding. A Bayesian process has been suggested as a potential mechanism for optimally combining the




Figure 1: Reproduced results from Weller et al. (2020) experiment 3A (left) and 3B (right) publicly available datasets. Bars depict average raw estimations by trial-type across intervals. Horizontal dotted black lines indicate average baseline error estimates for each timing interval.

information from these varying cues (Moore & Fletcher, 2012).

Within this literature, a common assumption is that temporal binding is an implicit marker of agency. Therefore, the proposed accounts attempt to explain the effect by looking at the differences in potential agency between experimental conditions. However, these differences may not be related to agency.

To see this, consider the fact that at its core, the interval estimation task is a free recall task in which, regardless of trial-type, participants are asked to encode a time interval between two events and then recall their estimation of that interval at the end of each trial. Approaching the problem from the perspective of memory opens the possibility that the effect may be explained, for example, by differences in how participants encode information across trials, regardless of agency. Therefore, the contribution of memory to the observed binding patterns should be assessed.

Here we suggest that a simple memory mechanism can entirely account for the results reported in two publicly available temporal binding experiments (Weller et al., 2020, experiment 3A and 3B). We first demonstrate that when replotting the data as a function of the error (or bias from the objective timing) in the participants' interval estimates, a classic regression to the mean pattern appears. This is important, as regression to the mean is a well-known pattern in time perception and memory. For example, Vierordt's law is a well-established regression effect in time perception where participants tend to overestimate short durations and underestimate long durations (Lejeune & Wearden, 2009).

What's more, Huttenlocher et al. (2000) reported results from memory experiments where participants learned to represent a distribution over a particular stimulus feature (e.g., sizes of fish) and then regressed to the mean of that distribution when recalling. Here, we propose that participants might be doing the same thing in the temporal estimation task. To find out, we simulated observed results using the same Bayesian rational memory model that has

been successfully implemented in other studies (Huttenlocher et al., 2000; Hemmer et al., 2015; Persaud & Hemmer, 2014). These results show, for the first time in the temporal binding literature, that we can qualitatively account for the pattern of results in two temporal binding studies with a memory mechanism and, importantly, without appealing to agency or causal inference as additional underlying mechanisms.

Data Sets

The data described here come from a publicly available paper (Weller et al., 2020, experiments 3A and 3B). Code for all analysis, figures, and supplementary material included in this paper are also publicly available (<https://osf.io/juh5y/>). A detailed description of the experimental method and procedure can be found in the supplementary material as well as in the original Weller et al. (2020) paper.

Here we briefly describe the procedure for both experiments 3A and 3B. Both experiments included three trial-types: action, non-action, and baseline. At the beginning of each trial, participants were asked to choose between an action and a non-action which would each produce distinct outcomes. In the non-action trials, participants chose not to act and a default outcome would occur, whereas in the action trials, participants acted (i.e., pressed a button at a timing of their choosing) to change the default outcome. Note that a participant could in principle choose not to act on any trials. However, only participants with a minimum of five or more observations per cell were included in the final analysis.

At the end of each trial, regardless of which type, participants were asked to recall and report their estimate of the interval between two events in milliseconds using a visual analog scale presented on-screen. There was also a baseline trial-type which was completely passive (i.e., no initial decision to act or not) and was compared to the other two trial-types. Three different time intervals were used between events (100ms, 400ms, and 700ms in 3A; 300ms, 500ms, and 700ms in 3B). The presentation of these intervals was randomized across the different blocks of trials. Data




Figure 2: Regression lines by experimental trial-type for Weller et al. (2020) experiment 3A (left) and 3B (right). Plots indicate average error (objective value subtracted from the average subjective estimation) of each interval across experimental trial-types. In experiment 3A, the slope for actions is significantly steeper compared to the other trial-types driven by larger average underestimation at the longest timing interval. For experiment 3B, the general pattern across trial-types is similar to results from experiment 3A though there are qualitatively larger differences in slopes across conditions indicating more noise in the memory process.

presented here are from 27 participants from experiment 3A and 40 participants from experiment 3B.

In the original publication (Weller et al., 2020), the authors compared trial-type and delay, and reported two significant results for experiment 3A: actions compared to baseline at the 700ms interval and non-actions compared to baseline at the 400ms interval. There were also two significant results for experiment 3B: actions compared to baseline at the 700ms interval and non-actions compared to baseline at the 700ms interval. No other comparisons were significant. From these results, the authors concluded that “temporal binding ha[d] [also] emerged for non-actions”.

Figure 1 depicts the average raw estimation across the three trial-types (actions, non-actions, and baseline) and across the three different timing intervals used. These plots show a qualitative reproduction of the published results.

Reanalysis

We re-analyzed, replotted, and simulated the original results after using exclusion criteria reported in the published paper. While we sought to reproduce and present the data exactly as Weller et al. (2020) published them, our reanalysis did produce some minor differences in the total number of trials removed due to exclusionary criteria. Our reanalysis removed a total of 2.63% of trials due to error or SD criteria from the data for experiment 3A and a total of 2.46% of trials from the data for experiment 3B which are comparable to those reported in the paper (2.1% and 2.2% of trials in 3A and 3B, respectively).

To evaluate the regression effect, we first replotted the raw estimates as error, or the difference between the average estimates and the objective timing at that interval. The regression pattern can clearly be seen in Figure 2 for experiments 3A and 3B. There is a consistent overestimation of the shortest interval (i.e., the value below the overall mean across intervals) and underestimation of the longest timing

interval (i.e., the value above the overall mean) across both experiments. All regression lines have a negative slope across the three intervals and are significantly different from zero ($p < 0.05$).

Note that sequential dependencies (i.e., the influence of a characteristic of the stimulus immediately prior to the current stimulus on the recall of the current stimulus) have been suggested as a potential cause of the regression pattern that is typically seen in memory experiments (Sailor & Antoine, 2005). An important step was to ensure that this pattern of results cannot be explained by the effect of sequential dependencies.

To that end, we first separated the trials into three groups: cases where the previous trial’s interval was longer than the current trial, cases where the previous trial’s interval was shorter than the current trial, and cases where the previous trial’s interval was the same as the current trial. In cases where the interval in the previous trial was shorter, we expect to see underestimation since the response will be drawn toward the shorter interval in the previous trial. We expect to see the opposite case when the previous trial’s interval was longer. The most informative case is when the previous trial’s interval was the same as the current trial. If the error disappears on these trials, this would suggest that sequential dependencies on the other trial types are driving the aggregate effect. However, if the error remains, a different explanation is in order.

We completed this analysis by plotting the frequency of all error across participants for cases where the previous trial’s interval was the same. We further split this data into three subsets (over length of intervals) and calculated the mean at each interval. We also completed Bayesian paired samples t-tests evaluating whether these means differed significantly from the overall means for each interval length. These analyses provided evidence in favor of the null hypothesis and confirmed that sequential dependencies do not account




Figure 3: Results of the Bayesian memory model simulation (including 95% CIs) for experiment 3A (top row) and experiment 3B (bottom row) across experimental trial types. Model parameters extracted from the data using overall averages for prior mean and variance as well as memory noise value based on overall standard deviation from each experimental condition. Model qualitatively fits well for all intervals and trial-types across both experiments except for action trials in experiment 3B.

for this pattern of results (see Tables S1 and S2 in supplementary material for more detailed information). That is, the regression pattern remained for the “same as current” trials.

A Bayesian Model of Memory

Next, we used a Bayesian rational memory model to simulate the results from experiments 3A and 3B. As previously mentioned, our goal for implementing this model in this context was to provide a plausible alternative explanation for interval timing in this context. We planned to evaluate whether a memory mechanism could account for this pattern of results without appealing to the influence of experienced agency. The model we use here is based on the Bayesian models reported in Huttenlocher (1991, 2000) as well as Hemmer and Steyvers (2009). This simple Bayesian memory model assumes that recall is a combination of noisy memory traces and prior expectations of interval lengths learned across the trials.

In the Weller et al. (2020) experiments, an observer is faced with the task of recalling features of a study stimulus (e.g., estimating a time interval). The observer’s goal is to reconstruct the original study stimulus feature θ using noisy samples y retrieved from memory. Bayes’ rule gives a principled method for combining prior expectations and evidence from memory to calculate the posterior probability,

$$p(\theta|y) \propto p(y|\theta)p(\theta) \quad \text{Eq (1)}$$

The posterior probability $p(\theta|y)$ describes how likely attribute values θ are given the noisy memory contents y and

prior expectations about the attributes. Suppose the feature values of the stimulus are Gaussian distributed, $\theta \sim N(\mu^*, \tau^*)$, where μ^* and τ^* are the prior mean and precision of the feature values. Further suppose that the samples y being drawn from memory have a Gaussian noise distribution centered on the original studied value, $y \sim N(\theta, \psi)$. The variance of the noise process, ψ , determines the resemblance of the stored representations to the original feature of the study stimulus. Suppose the observer also has some expectations about the general distribution of attributes, $\theta \sim N(\mu, \tau)$. This distribution corresponds to the prior in the observer’s memory model and assumes that the observer has learned the environmental statistics which can be used as a proxy for the prior in the observer’s model. That is, $\mu = \mu^*$ and $\tau = \tau^*$. Bayesian techniques can now be used to calculate the posterior distribution:

$$\theta|y, \psi, \mu, \tau \sim N\left(\frac{\psi y + \mu \tau}{\psi + \tau}, \psi + \tau\right) \quad \text{Eq (2)}$$

The mean of the recalled stimulus values (e.g., time intervals) is a weighted linear combination of the prior mean μ and the mean of memory content y . The prior mean μ is weighted more heavily when the prior has a higher precision (τ) or when the memory noise increases. This corresponds to the intuition that if the prior is strong, it will exert a strong influence. Similarly, if memory contents are very noisy, the prior will also exert a strong influence. This linear combination explicitly predicts the regression pattern that is well-known in memory.

Simulation Results

We applied this simple Bayesian rational model of memory described above to the Weller et al. (2020) data. The goal of this analysis was to compare the predictions of the Bayesian model of memory and the empirical data at a qualitative level. In the model, the priors were based directly on the environmental statistics, i.e., the average over subjective estimates. However, rather than hand fitting the parameters ad hoc, we used the summary statistics from the experimental data as parameter input. For the prior in Eq. 2, we used distinct mean μ and precision τ values corresponding to the overall error values from each experiment ($\mu = 454.78\text{ms}$ [overall μ] and $\tau = 91.11\text{ms}$ [overall SD] for experiment 3A and $\mu = 448.72\text{ms}$ [overall μ] and $\tau = 97.95\text{ms}$ [overall SD] for experiment 3B). Furthermore, we assumed the same prior and memory noise setting for all participants (i.e., no individual variation). We chose to use prior parameter values extracted from aggregate participant responses as we assumed that participants learned the mean of the intervals over the course of the experiment. We felt this choice was appropriate considering there is evidence to suggest that priors are learned early in an experiment (Berniker et al., 2010). We also completed simulations using the mean of the stimulus feature ($\mu = 400\text{ms}$ for EXP 3A and $\mu = 500\text{ms}$ for EXP 3B) and did not find a substantial difference in qualitative fit (See Supplemental Material for more details). For the memory precision ψ , we used a value of 90, which is based on the overall standard deviation of timing estimates in the experimental data (while it is technically unitless, this value could be interpreted in milliseconds). We obtained 1000 samples for each distribution and simulated the same number of participants, trials, and time intervals that were used in each experiment. The model simulated raw participant estimations for each trial which were then used to calculate the bias on each trial. Figure 3 depicts the simulated responses compared to those obtained from participants in the Weller et al. (2020) study. The triangles represent the maximum a posteriori (MAP) estimates from the posterior predictive recall distribution simulated from the model with a 95% CI around each simulated mean.

For both simulated and observed responses, the results show effects of the prior (i.e., the overall mean across intervals). For experiment 3A, Figure 3 (top row) shows that the model qualitatively simulates the overall effect. The time intervals that are shorter than the mean of time intervals (i.e., 400ms) are overestimated while the time interval that is longer than the mean is underestimated. The model qualitatively fits best for actions and non-actions while under- and over-estimating the error in the 400 and 700ms timing interval baseline trials, respectively. For experiment 3B, Figure 3 (bottom row) shows that the model provides an excellent fit for baseline and non-action trials though interestingly, it does not fit action trials as well.

This last result appears to be due to a substantially different y-intercept value in the regression line which suggests that the parameter values used in the model for action trials in this experiment are not likely the ones participants used. It is




Figure 4: Bayesian memory model simulation results (including 95% CIs) for Weller et al. (2020) experiment 3B action trials. Decreasing the value of the prior mean parameter for action trials only improved the overall qualitative fit of the model to the trial data. This indicates that participants may have used a different prior mean when making estimations during action trials in this task.

plausible that there may be more memory noise in the memory process for these trials compared to the baseline trials. It could also be that participants used a different prior mean or variance to make their estimates for these trials. To assess this, we tested larger and smaller memory noise values without changing the prior parameter values which did not improve the qualitative fit (see Supplemental Material Figure S2). We then varied the prior mean parameter value in the model and found that a value of 350 (slightly smaller than the 448.72ms value used for the simulations in Figure 3) improved the fit substantially (see Figure 4).

Discussion

We investigated whether a memory process could explain the pattern of results reported in two publicly available temporal binding datasets. After replotting the raw estimations of timing intervals as error (or the difference between the estimations and the objective timing), a clear regression pattern emerged. We implemented a well-established Bayesian rational memory model and found that the model qualitatively simulates the experimental data across three different trial types. Using parameter values extracted from the data resulted in a good qualitative fit across all intervals and trial-types in experiment 3A. The model also produced a good qualitative fit for non-action and baseline trials but did not produce as good a qualitative fit for action trials in experiment 3B.

The improved fit using a lower value for the prior mean parameter suggests that participants may be regressing to a shorter interval length when recalling estimations after action trials as compared to other trials in the experiment. This could be because participants perceive the time between their actions and the outcomes as shorter than in the other trials which may be due to an increase in perceived agency in these trials.

A simpler and more likely explanation may be that encoding is different in the action trials compared to the nonaction and baseline trials. The presence of the action may be more distracting which would increase the error in the initialization of the timing process. There may also be increased noise in the recall process which was reflected in an increased standard deviation for actions ($SD = 115.57ms$) compared to nonaction ($SD = 100.29ms$) and baseline ($SD = 78ms$) trials. However, these explanations are not entirely satisfactory, as the model fit is qualitatively better for experiment 3A compared to experiment 3B and there is no perceptual difference between the conditions in terms of agency or encoding across the experiments. This may represent a superficial and idiosyncratic feature of the experiment rather than the reflection of something deeper. We acknowledge that this is also not a satisfactory explanation and recognize it as an area for future research.

Relatedly, the aggregated results across the studies also suggest that the range of intervals used in the experiment may have some effect on the pattern of errors. In experiment 3A, the range of possible intervals was larger than in experiment 3B. Though the overall regression pattern was still the same, the directionality of the average error was switched at the middle timing interval: overestimation in 3A and underestimation in 3B. At first blush, this seems like a perplexing result. However, one should note that the overall average estimation value was close to the midpoint of the visual response scale: approximately 450ms for both studies (454ms for 3A and 448ms for 3B). This suggests that there may be a more general, environmental prior that participants are using to make estimations in these experiments or that they are being influenced by the range of responses presented to them. This could explain the reversal in directionality of the error for the middle timing interval across the two experiments. It is also possible that the length of the general prior differs depending on the type of action or the outcome modality. Finding out requires further investigation.

The iteration of the model reported here assumed that all participants used the same parameter values to make their estimations. This is obviously an oversimplification and represents a limitation in our implementation as it is likely that the prior mean, variance, and memory noise values differ from person to person. Future iterations of the model will allow us to uncover individual differences in these parameters which may elucidate these results.

Our model fits well with other models of time perception. Specifically, the idea that memory may be a significant influence for interval timing (and recall) is not new (Addyman et al., 2011; Addyman & Mareschal, 2014; Fountas & Zakharov, 2022; Jazayeri & Shadlen, 2010). The novel contribution of our work lies in the application of this simple Bayesian memory model in the context of temporal binding as well as our suggestion that memory can account for an effect commonly attributed to agency. We recognize that this is not a novel concept in the time perception literature more broadly and as such more complex models could also be incorporated in future work.

Importantly, we also acknowledge that this model does not necessarily exclude agency as a potential influence on the TB effect. It could be that the mediating factor for a smaller prior mean for actions is an increase in perceived agency, though this seems unlikely since this result did not replicate across the experiments. Though the relationship between agency and binding remains unclear (Klaffehn et al., 2021; Suzuki et al., 2019), one could speculate that when participants perform voluntary actions, this may be preceded by a prior expectation based on their experience that they can use to make predictions or explicitly recall time intervals. Such intervals may have developed as a result of their experienced agency (i.e., “I expect an immediate interval in cases where I feel in control”), or it may be that the intervals determine their experienced agency (i.e., “if something happened that quickly after my action then I must have caused it”). Determining the answers to these intriguing questions will require further research.

Additionally, it is important to note that our model is compatible with approaches based on cue integration. These approaches do not specify a memory process though one is necessary to explain the results in a temporal binding task. Cue integration attempts to explain the process by which an individual optimally encodes and combines information from multiple sources. Furthermore, it has been suggested that a Bayesian process may explain how this is accomplished thereby allowing us to experience agency (Legaspi & Toyoizumi, 2019; Moore & Fletcher, 2012). Though it remains unclear how cue integration can account for all aspects of the binding effect (Wolpe et al., 2013), our model could act as an extension to existing cue integration approaches. In this combined model, memory content would be updated using the cue integration process and then later combined with the prior mean and variance and memory noise parameter values in our model (using the linear weighted combination) to arrive at an estimate of the interval on each trial. This provides another interesting avenue for future research.

Here we have provided evidence for a novel explanation of the temporal binding effect. Approaching binding from the perspective of memory allows us to account for the effect as (partially) epiphenomenal. Although the role of memory does not preclude that agency may also be involved, it needs to be fully accounted for before claims about agentic processes can be made. This perspective opens new avenues of research regarding the role of memory in temporal binding tasks and in our everyday experiences of control.

Acknowledgments

This research was supported by the National Science Foundation CAREER grant #1453276 to PH. We would like to thank Lisa Weller and her collaborators for making their data publicly available. Their contribution to open science made this work possible and will help improve research in this domain.

References

- Addyman, C., French, R. M., Mareschal, D., & Thomas, E. (2011). Learning to perceive time: A connectionist, memory-decay model of the development of interval timing in infants. *Proceedings of the 33rd Annual Conference of the Cognitive Science Society*, 7.
- Addyman, C., & Mareschal, D. (2014). GAMIT-Net: Retrospective and prospective interval timing in a single neural. *Proceedings of the 36th Annual Meeting of the Cognitive Science Society*, 7.
- Barlas, Z., & Kopp, S. (2018). Action Choice and Outcome Congruency Independently Affect Intentional Binding and Feeling of Control Judgments. *Frontiers in Human Neuroscience*, 12, 137. <https://doi.org/10.3389/fnhum.2018.00137>
- Berberian, B., Sarrazin, J.-C., Le Blaye, P., & Haggard, P. (2012). Automation Technology and Sense of Control: A Window on Human Agency. *PLoS ONE*, 7(3), e34075. <https://doi.org/10.1371/journal.pone.0034075>
- Berniker, M., Voss, M., & Kording, K. (2010). Learning Priors for Bayesian Computations in the Nervous System. *PLoS ONE*, 5(9), e12686. <https://doi.org/10.1371/journal.pone.0012686>
- Borhani, K., Beck, B., & Haggard, P. (2017). Choosing, Doing, and Controlling: Implicit Sense of Agency Over Somatosensory Events. *Psychological Science*, 28(7), 882–893. <https://doi.org/10.1177/0956797617697693>
- Buehner, M. J. (2015). Awareness of voluntary and involuntary causal actions and their outcomes. *Psychology of Consciousness: Theory, Research, and Practice*, 2(3), 237–252. <https://doi.org/10.1037/cns0000068>
- Caspar, E. A., Christensen, J. F., Cleeremans, A., & Haggard, P. (2016). Coercion Changes the Sense of Agency in the Human Brain. *Current Biology*, 26(5), 585–592. <https://doi.org/10.1016/j.cub.2015.12.067>
- Cavazzana, A., Begliomini, C., & Bisiacchi, P. S. (2014). Intentional binding effect in children: Insights from a new paradigm. *Frontiers in Human Neuroscience*, 8. <https://doi.org/10.3389/fnhum.2014.00651>
- Daprati, E., Franck, N., Georgieff, N., Proust, J., Pacherie, E., Dalery, J., & Jeannerod, M. (1997). Looking for the agent: An investigation into consciousness of action and self-consciousness in schizophrenic patients. *Cognition*, 65(1), 71–86. [https://doi.org/10.1016/S0010-0277\(97\)00039-5](https://doi.org/10.1016/S0010-0277(97)00039-5)
- Engbert, K., Wohlschläger, A., & Haggard, P. (2008). Who is causing what? The sense of agency is relational and efferent-triggered. *Cognition*, 107(2), 693–704. <https://doi.org/10.1016/j.cognition.2007.07.021>
- Fereday, R., Buehner, M. J., & Rushton, S. K. (2019). The role of time perception in temporal binding: Impaired temporal resolution in causal sequences. *Cognition*, 193, 104005. <https://doi.org/10.1016/j.cognition.2019.06.017>
- Fountas, Z., & Zakharov, A. (2022). Bayesian sense of time in biological and artificial brains. *ArXiv:2201.05464 [Cs, q-Bio, Stat]*. <http://arxiv.org/abs/2201.05464>
- Graham-Schmidt, K. T., Martin-Iverson, M. T., Holmes, N. P., & Waters, F. A. V. (2016). When one's sense of agency goes wrong: Absent modulation of time perception by voluntary actions and reduction of perceived length of intervals in passivity symptoms in schizophrenia. *Consciousness and Cognition*, 45, 9–23. <https://doi.org/10.1016/j.concog.2016.08.006>
- Haggard, P. (2017). Sense of agency in the human brain. *Nature Reviews Neuroscience*, 18(4), 196–207. <https://doi.org/10.1038/nrn.2017.14>
- Haggard, P., Clark, S., & Kalogeras, J. (2002). Voluntary action and conscious awareness. *Nature Neuroscience*, 5(4), 382–385. <https://doi.org/10.1038/nn827>
- Hemmer, P., & Steyvers, M. (2009). A Bayesian Account of Reconstructive Memory. *Topics in Cognitive Science*, 1(1), 189–202. <https://doi.org/10.1111/j.1756-8765.2008.01010.x>
- Hemmer, P., Tauber, S., & Steyvers, M. (2015). Moving beyond qualitative evaluations of Bayesian models of cognition. *Psychonomic Bulletin & Review*, 22(3), 614–628. <https://doi.org/10.3758/s13423-014-0725-z>
- Hoerl, C., Lorimer, S., McCormack, T., Lagnado, D. A., Blakey, E., Tecwyn, E. C., & Buehner, M. J. (2020). Temporal Binding, Causation, and Agency: Developing a New Theoretical Framework. *Cognitive Science*, 44(5). <https://doi.org/10.1111/cogs.12843>
- Humphreys, G. R., & Buehner, M. J. (2010). Temporal binding of action and effect in interval reproduction. *Experimental Brain Research*, 203(2), 465–470. <https://doi.org/10.1007/s00221-010-2199-1>
- Huttenlocher, J., Hedges, L. V., & Duncan, S. (1991). Categories and Particulars: Prototype Effects in Estimating Spatial Location. *Psychological Review*, 98(3), 352–376.
- Huttenlocher, J., Hedges, L. V., & Vevea, J. L. (2000). Why Do Categories Affect Stimulus Judgment? *Journal of Experimental Psychology: General*, 129(2), 220–241.
- Jazayeri, M., & Shadlen, M. N. (2010). Temporal context calibrates interval timing. *Nature Neuroscience*, 13(8), 1020–1026. <https://doi.org/10.1038/nn.2590>
- Kirsch, W., Kunde, W., & Herbol, O. (2019). Intentional binding is unrelated to action intention. *Journal of Experimental Psychology: Human Perception and Performance*, 45(3), 378–385. <https://doi.org/10.1037/xhp0000612>
- Klaffehn, A. L., Sellmann, F. B., Kirsch, W., Kunde, W., & Pfister, R. (2021). Temporal binding as multisensory integration: Manipulating perceptual certainty of actions and their effects. *Attention, Perception, & Psychophysics*, 83(8), 3135–3145. <https://doi.org/10.3758/s13414-021-02314-0>
- Legaspi, R., & Toyoizumi, T. (2019). A Bayesian psychophysics model of sense of agency. *Nature Communications*, 10(1), 4250. <https://doi.org/10.1038/s41467-019-12170-0>
- Lejeune, H., & Wearden, J. H. (2009). Vierordt's *The Experimental Study of the Time Sense* (1868) and its

- legacy. *European Journal of Cognitive Psychology*, 21(6), 941–960. <https://doi.org/10.1080/09541440802453006>
- Lush, P., Roseboom, W., Cleeremans, A., Scott, R. B., Seth, A. K., & Dienes, Z. (2019). Intentional binding as Bayesian cue combination: Testing predictions with trait individual differences. *Journal of Experimental Psychology: Human Perception and Performance*, 45(9), 1206–1217. <https://doi.org/10.1037/xhp0000661>
- Metcalfe, J., Eich, T. S., & Miele, D. B. (2013). Metacognition of agency: Proximal action and distal outcome. *Experimental Brain Research*, 229(3), 485–496. <https://doi.org/10.1007/s00221-012-3371-6>
- Metcalfe, J., & Greene, M. J. (2007). Metacognition of agency. *Journal of Experimental Psychology: General*, 136(2), 184–199. <https://doi.org/10.1037/0096-3445.136.2.184>
- Moore, J. W., & Fletcher, P. C. (2012). Sense of agency in health and disease: A review of cue integration approaches. *Consciousness and Cognition*, 21(1), 59–68. <https://doi.org/10.1016/j.concog.2011.08.010>
- Obhi, S. S., Swiderski, K. M., & Farquhar, R. (2013). Activating memories of depression alters the experience of voluntary action. *Experimental Brain Research*, 229(3), 497–506. <https://doi.org/10.1007/s00221-012-3372-5>
- Persaud, K., & Hemmer, P. (2014). The Influence of Knowledge and Expectations for Color on Episodic Memory. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 36(6), 7.
- Pfister, R., Obhi, S. S., Rieger, M., & Wenke, D. (2014). Action and perception in social contexts: Intentional binding for social action effects. *Frontiers in Human Neuroscience*, 8. <https://doi.org/10.3389/fnhum.2014.00667>
- Poonian, S. K., McFadyen, J., Ogden, J., & Cunningham, R. (2015). Implicit Agency in Observed Actions: Evidence for N1 Suppression of Tones Caused by Self-made and Observed Actions. *Journal of Cognitive Neuroscience*, 27(4), 752–764. https://doi.org/10.1162/jocn_a_00745
- Sailor, K. M., & Antoine, M. (2005). Is memory for stimulus magnitude Bayesian? *Memory & Cognition*, 33(5), 840–851. <https://doi.org/10.3758/BF03193079>
- Seghezzi, S., & Zapparoli, L. (2020). Predicting the Sensory Consequences of Self-Generated Actions: Pre-Supplementary Motor Area as Supra-Modal Hub in the Sense of Agency Experience. *Brain Sciences*, 10(11), 825. <https://doi.org/10.3390/brainsci10110825>
- Spengler, S., von Cramon, D. Y., & Brass, M. (2009). Was it me or was it you? How the sense of agency originates from ideomotor learning revealed by fMRI. *NeuroImage*, 46(1), 290–298. <https://doi.org/10.1016/j.neuroimage.2009.01.047>
- Suzuki, K., Lush, P., Seth, A. K., & Roseboom, W. (2019). Intentional Binding Without Intentional Action. *Psychological Science*, 30(6), 842–853. <https://doi.org/10.1177/0956797619842191>
- Takahata, K., Takahashi, H., Maeda, T., Umeda, S., Suhara, T., Mimura, M., & Kato, M. (2012). It's Not My Fault: Postdictive Modulation of Intentional Binding by Monetary Gains and Losses. *PLoS ONE*, 7(12), e53421. <https://doi.org/10.1371/journal.pone.0053421>
- Tanaka, T., Matsumoto, T., Hayashi, S., Takagi, S., & Kawabata, H. (2019). What Makes Action and Outcome Temporally Close to Each Other: A Systematic Review and Meta-Analysis of Temporal Binding. *Timing & Time Perception*, 7(3), 189–218. <https://doi.org/10.1163/22134468-20191150>
- Weller, L., Schwarz, K. A., Kunde, W., & Pfister, R. (2020). Something from nothing: Agency for deliberate nonactions. *Cognition*, 196, 104136. <https://doi.org/10.1016/j.cognition.2019.104136>
- Zhao, K., Hu, L., Qu, F., Cui, Q., Piao, Q., Xu, H., Li, Y., Wang, L., & Fu, X. (2016). Voluntary action and tactile sensory feedback in the intentional binding effect. *Experimental Brain Research*, 234(8), 2283–2292. <https://doi.org/10.1007/s00221-016-4633-5>