Behavioral signatures of temporal context retrieval during continuous recognition

Atharva Joshi (atharva.j@research.iiit.ac.in)
Cognitive Science Lab - IIIT-H,
Hyderabad, 500032

Kamalaker Dadi (kamalaker.dadi@ihub-data.iiit.ac.in)
iHub-data - IIIT-H,
Hyderabad, 500032

Vishnu Sreekumar (vishnu.sreekumar@iiit.ac.in)
Cognitive Science Lab - IIIT-H,
Hyderabad, 500032

Abstract

An influential mathematical model of memory, the temporal context model (TCM), posits that we encode items and their associations with *temporal context* (Howard & Kahana, 2002). Temporal context is conceived of as a recency-weighted average of past experiences. Critically, the model assumes that when an item is retrieved later, the associated temporal context is also obligatorily retrieved. Existing evidence for the idea of retrieved temporal context primarily comes from free-recall studies. However, free recall introduces some critical confounds that are difficult to resolve (Folkerts et al., 2018) and also encourages memory strategies that may mimic temporal context effects (Hintzman, 2011). To address these confounds, we investigate temporal context using an image recognition task. Schwartz et al. (2005) examined temporal context in an image recognition task using a short-list-based experimental design, and found that temporal context influenced recognition performance. Building on this, we use the Natural Scenes Dataset (NSD) to show that reinstating temporal context enhances recognition accuracy even across substantially longer timescales. We demonstrate that images that were temporally closer during encoding facilitated the recognition of each other. Critically, we show that this influence falls off with temporal distance at encoding only when the temporal context is successfully retrieved, as predicted by TCM. Furthermore, the slope of this temporal gradient increases as a function of the strength of the influence of the retrieved temporal context. These findings extend our understanding of temporal context effects in episodic memory by showing that temporal context is retrieved even in tasks that do not encourage linking between items as a memory strategy.

Keywords: Temporal context model; retrieved temporal context; continuous recognition; episodic memory

Introduction

One of the defining features of episodic memory is its ability to encode both the content of an event and the context in which it occurred (Tulving, 1972). When a stimulus is encoded, it becomes part of a network that also includes the surrounding context within which the stimulus was experienced. Critically, according to the temporal context model (TCM) (Howard & Kahana, 2002), retrieving the stimulus at a later time reactivates not just the features of the stimulus but also *obligatorily* retrieves its associated temporal context. The temporal context associated with a stimulus can be conceptualized as a recency-weighted average of ongoing events (Howard & Kahana, 2002), where events closer in time contribute more to the context than those further in the past. Evidence for retrieved temporal context has primarily come from free recall data. However, free recall encourages certain strategies (Hintzman, 2011) and introduces confounds (Folkerts et al., 2018) that make it difficult to distinguish retrieved temporal context from other mundane explanations. In the current study, we examine continuous image recognition data (Allen et al., 2022), which does not have these drawbacks of free recall, and show that temporal context is retrieved and influences memory performance as predicted by TCM.

Slow drift in internal representations can be thought of as the neural substrate underlying temporal context. Unlike the transient nature of external experiences, internal representations exhibit inertia, changing more gradually than the environment around us (DuBrow et al., 2017). Thus, two events close in time share a similar temporal context, which is encoded and retrieved alongside the event. Note that unlike associative chaining, which directly links two temporally close events (Caplan et al., 2022), temporal context serves as an evolving background shaped by the sequential flow of events, indirectly acting as a bridge between events close in time.

Behavioral and neural evidence for *retrieved temporal context* primarily comes from free recall tasks. Kahana (1996) showed that free-recall transitions were more likely to occur between neighboring list items, with a pronounced preference for forward-order transitions over backward ones. TCM was developed to provide a parsimonious account of such *temporal contiguity* effects. However, since contiguity could arise from other mechanisms, such as direct associative chaining between adjacent items, it was important to provide neural evidence for retrieved temporal context. Manning et al. (2011) analyzed brain recordings from 69 patients during free recall of word lists and demonstrated a *neural contiguity effect*, thought to be evidence for retrieved temporal context.

However, the order of retrieval in a free recall experiment is not random and confounds neural contiguity analyses (Folkerts et al., 2018). Specifically, *lack of experimental control over recall order*, *residual activation from previously recalled items* (Folkerts et al., 2018), and *semantic or narrative-based strategies* that participants often employ (Hintzman, 2011) can all lead to illusory behavioral and neural contiguity effects. Note again that free recall exhibits a forward-asymmetric contiguity effect (Kahana, 1996) such that most words are recalled in approximately the same order

In D. Barner, N.R. Bramley, A. Ruggeri and C.M. Walker (Eds.), *Proceedings of the 47th Annual Conference of the Cognitive Science Society*

as they were presented originally. Therefore, brain activity when any word is recalled tends to have residual activation from the just recalled words, which happen to be temporally proximal words from the encoding phase. This confound, artifactually, makes the similarity between brain activity at recall and encoding fall off with temporal lag (Folkerts et al., 2018) and challenges our ability to isolate retrieval of gradually changing temporal context representations from short-term item persistence effects. Moreover, Hintzman (2011) argued that tasks like free recall encourage explicit memory strategies that link adjacent items through narrativization that, in turn, result in patterns that look like temporal context retrieval. Therefore, we analyze data from a continuous recognition task with a random order of stimuli; this does not encourage strategies that associate nearby items, addressing both drawbacks identified with free recall tasks.

While Schwartz et al. (2005) demonstrated temporal context effects in recognition using short-range lags within a session, we extend this by leveraging the Natural Scenes Dataset (NSD) to study context reinstatement across sessions and weeks, offering a more naturalistic and large-scale validation. Crucially, NSD includes neural data for each trial, enabling us to directly examine the neural signatures of temporal context reinstatement in an image recognition task (though the neural analysis is not the focus of the current paper).

In this study, we examine whether recognition accuracy for an image (*target image*) improves when it is preceded by successful recognition of another image (*cue image*) that was temporally proximal to it during encoding. We further investigate whether recognition accuracy drops as a function of temporal distance between the cue and target image at *encoding*, i.e., whether there is a temporal gradient to the influence of retrieved temporal context, as predicted by TCM. Finally, we test whether the predicted temporal gradient is stronger if the target image follows the cue image immediately, compared to when the target image appears farther away from the cue image at retrieval. This prediction is based on the idea that retrieved temporal context also decays with time, and therefore its influence (measured as a temporal gradient at encoding) also weakens with time since retrieval. Our results indicate that temporal context is retrieved *obligatorily* even when the task does not encourage strategies that would mimic temporal context retrieval and contribute to settling a longstanding debate about *retrieved temporal context* during episodic memory.

Methodology

Dataset

We use the openly accessible Natural Scenes Dataset (NSD) [1] (Allen et al., 2022). The NSD acquired fMRI data from eight healthy adult participants (2M, 6F, age range $= 19 - 32yrs$) during a continuous recognition memory task. The data were collected over the course of ≈ 1 year. Each participant did ≈ 40 sessions separated by roughly one week. We focus on temporal context retrieval across sessions but not within sessions to avoid recency effects in our analyses. Each session consisted of twelve 5–minute runs, with functional Magnetic Resonance Imaging (fMRI) brain activity recorded continuously during each run. In each run, participants viewed $62 - 63$ natural scene images presented sequentially, and they had to press a button to indicate whether they had previously seen an image in any previous or current session of the experiment. For more details on NSD, please see Allen et al. (2022).

[1] https://naturalscenesdataset.org/

Hypotheses

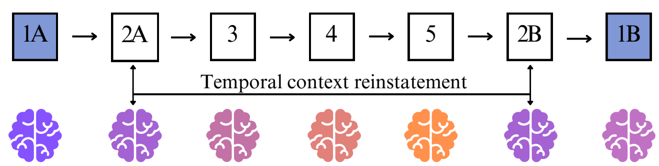

Figure 1: The gradually changing color represents slow drift in temporal context in the brain during the sequential presentation of images in a continuous recognition task. The brain state at 2A is more similar to the brain state at 3 compared to 4, 5, etc. Features of 1A are present in the temporal context of 2A and come to be associated with the features of 2A. When image 2 is repeated (2B), we reinstate brain activity from 2A, which contains not just features corresponding to 2A but also to 1A, facilitating the recognition of image 1B if it happens to be presented soon after 2B. *Target* images of the sort shaded in blue, contingent on successfully recognized *cue* images such as 2B, are selected for further analyses.

Images experienced close together in time are encoded within overlapping contextual representations. During the encoding phase, when image 2A (Figure 1) is presented for the first time, brain activity at that moment captures not only the features of image 2A but also the lingering context from previous images, such as 1A. As a result, image 2A is encoded with a context that includes both its own features and residual information from nearby images like 1A.

Later, when the same image (now labeled 2B in Figure 1) is presented again, TCM posits that successful recognition of the image not only reinstates 2B-related features but also the temporal context that was active during its initial encoding, which contains features of 1A and other items temporally adjacent to 2A during encoding. We condition our analyses of subsequent images (*target images*) on such successfully recognized images (*cue images*).

This TCM-based idea that temporal context is retrieved during the successful recognition of a cue image leads us to our three hypotheses:

H1: The recognition accuracy for a *target* image (1B) will be higher when preceded by the successful recognition of a temporally proximal image (2B). This pattern would be consistent with the assumption that temporal context is retrieved during successful recognition of 2B and that it will contain features of 1A.

H2: The facilitation from the retrieved temporal context of 2B for an ensuing image 1B will be positively correlated with temporal proximity (or negatively correlated with temporal *LAG*) between the images (1A and 2A) during encoding. Furthermore, this correlation between target recognition accuracy and *LAG* should disappear if the cue image 2B was not recognized successfully.

H3: The negative correlation between the recognition success and temporal lag at encoding as stated in **H2** will become weaker if TDR (temporal distance at retrieval) is increased. This is because the retrieved temporal context of image 2B itself decays with time.

Selection criteria for images to be analyzed

Since temporal context is hypothesized to be retrieved during successful recognition, we first identify successfully recognized images. Then we analyze the influence of the successfully retrieved temporal context on subsequent images within a short window. The specific criteria for the selection of images and the time elapsed between them at both encoding and retrieval are as follows:

Criterion 1: Selecting instances of image presentations for analysis. Each image was repeated three times in the experiment. We focus on the first two instances: the first instance establishes the original temporal context for a stimulus, and the second instance is used to retrieve and test the established temporal context. For example, in Figure 2 (left), when image 2A is encoded for the first time, the temporal context associated with it also includes some information about 1A that was encountered recently. Then, when the image is repeated later (2B), if it is successfully recognized, TCM posits that its associated temporal context is also retrieved. This retrieved temporal context, due to having features of 1A and not just 2A, should facilitate the subsequent recognition of 1B if 1B happens to appear within a short time window of 2B.

Criterion 2: The distance between the cue and target image must be $<= 10$ *images.* The retrieved temporal context during the successful recognition of the cue image 2B must not decay significantly if it is to have a measurable influence on a subsequent image 1B. Thus, the two images, which were initially close during encoding to form an associated temporal context, must also be presented close to each other during retrieval. Our results are robust to minor changes in this parameter, *temporal distance at retrieval* (TDR).

Criterion 3: Selecting image pairs encoded together within the same run Temporal context drifts in time such that the similarity between contextual states decays with time. Therefore, we select image pairs that were seen in the same run as they occur within 5 minutes of each other, making it likely that signatures of one image are present in the temporal context of another.

Criterion 4: Mitigating the recency effect by separating the first and second instances of images by one session. Lastly, we aim to isolate the effect of temporal context from the influence of recency. Our analysis found out that images repeated within the same session were recognized with $\sim 90\%$ accuracy, suggesting a strong recency effect. The NSD experiment was designed such that an image presented in one session would typically repeat either in the same session or in the very next session. So to minimize the influence of recency and better isolate the effects of temporal context, we focus on images that were repeated across consecutive sessions separated by 1-2 weeks and eliminate images that were repeated within the same session from our analysis.

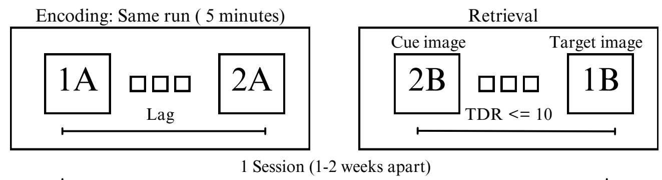

Figure 2: Images selected for analysis. (**Left**) Target images of type 1A that occurred at different LAGs within the same encoding run as the cue image 2A were selected for further analyses. (**Right**) The cue and target image must be within 10 images of each other ($TDR \leq 10$). The first and second presentations of both cue and target images must be separated by one session to minimize recency effects.

We identified a subset of 1200 cue-target image pairs that satisfy the four criteria identified above. Next, we identify the different variables that are critical for testing the predicted role of retrieved temporal context on the recognition probability of target images.

Variables influencing recognition of the target image

Var 1: Retrieved Temporal Context (RTC). We assume $RTC = 1$ when a *cue image* at retrieval is recognized correctly, signaling the successful retrieval of both the image and its temporal context. If the cue image is not recognized correctly, $RTC = 0$, since we assume that temporal context is not retrieved to the same extent as it is during successful recognition. Our assumptions about RTC, while consistent with the memory accuracy results comparing $RTC = 0$ and $RTC = 1$ (see Results), are tested further through the temporal lag analyses in Hypotheses 2 and 3.

Var 2: Temporal Lag at Encoding (LAG). *LAG* refers to the number of trials separating the two images during the encoding phase (see Figure 2 (left)). Since the presence of an image in the temporal context of another image is stronger if they are separated by a small *LAG* compared to a large *LAG*, the influence of retrieved temporal context on a target image will vary as a function of *LAG*. LAG can be positive or negative—negative when the target image appeared before the cue during encoding, and positive when it appeared after. Positive LAGs are relevant because cue images can still be part of the target's temporal context, allowing retrieval of the target even if it followed the cue during encoding.

Var 3: Temporal Distance at Retrieval (TDR). *TDR* (see Figure 2 (right)) represents the number of trials between the presentation of the cue image and the target image during the retrieval phase. The cue image triggers the retrieval of its associated temporal context. However, this retrieved context drifts away with time, potentially weakening its influence on target images presented farther away. Therefore, TDR is a relevant variable that can modulate the influence of retrieved temporal context on the recognition success of a target image.

Statistical models to test the hypotheses

We adopt a mixed-effect logistic regression approach to analyze the effects of the variables identified in the previous section on recognition accuracy. This approach leverages the hierarchical structure of the data, where observations were nested within participants (Jaeger, 2008). This statistical method incorporates both fixed effects, representing the systematic influence of predictors across participants, and random effects, which account for individual differences between participants.

Random effect structure. Including all factors as random effects led to a singular fit, suggesting over-parameterization or low data variability. We resolved this by treating the factors as fixed effects while retaining subject-level random intercepts, improving model convergence (Bates et al., 2015).

Base model for analyzing the influence of *RTC* on target recognition success. We use the following mixed-effects logistic regression model to test Hypothesis **H1** that retrieved temporal context (RTC) has a positive influence on the probability of successful target recognition:

$$\text{logit}(P(\text{Recognition}_{ij} = 1)) \sim \beta_0 + \beta_1 \cdot \text{RTC}_{ij} + u_i, \quad (1)$$

where $P(\text{Recognition}_{ij} = 1)$ represents the probability of correct recognition for trial j by participant i, β_0 is the intercept capturing the baseline log-odds of recognition when $RTC = 0$, β_1 quantifies the effect of RTC on recognition accuracy with a positive value indicating that $RTC = 1$ increases recognition probability of an ensuing target image, and $u_i \sim N(0, \sigma_u^2)$ is a random intercept for each participant accounting for individual differences in recognition accuracy.

Adding the effects of *LAG*, *TDR*, and their interactions. We further modify our model (as shown in Equation 2) to test Hypotheses **H2** and **H3**. Specifically, we predict a negative effect of LAG because an increase in LAG will lead to a weaker association between the images, resulting in reduced recognition accuracy. Furthermore, we hypothesize (**H3**) that this negative effect of LAG will be less strong with increasing TDR because retrieved context drifts away over larger windows (TDR). To understand this prediction of an interaction between LAG and TDR, imagine the retrieved temporal context has decayed to an extent that there is no residual activation related to the target image anymore. When the target image is presented at this point, we should not expect recognition success to have any relationship, either positive or negative, with LAG. For the same reason, we would also expect a negative main effect of TDR on target recognition accuracy.

$$\begin{aligned} &logit(P(\text{Recog.}_{ij} = 1)) \sim \beta_0 + \beta_1 \cdot \text{RTC} + \beta_2 \cdot \text{TDR} + \\ &\beta_3 \cdot \text{LAG}_{abs} + \beta_4 \cdot \text{TDR} \cdot \text{LAG}_{abs} + \beta_5 \cdot \text{RTC} \cdot \text{LAG}_{abs} + u_i \end{aligned} \quad (2)$$

We use absolute values of LAG and constrain LAG_{abs} to 20 trials ($\sim$ 100 seconds), which provides a reasonable range of time to test the decaying influence of temporal context. However, results are robust to minor variations in this constraint. We also include an interaction term between LAG_{abs} and RTC based on the expectation that the influence of LAG_{abs} on recognition accuracy is only significant when the temporal context is retrieved in the first place.

Finally, we modify the model to account for both positive and negative values of LAG (Equation 3), to investigate potential asymmetries in the influence of temporal context (Howard & Kahana, 2002). However, this asymmetric model failed to converge when RTC was retained in the model. Therefore, we reran the asymmetric model separately for $RTC = 0$ and $RTC = 1$ (which solved the convergence issues) solely for the purpose of visually comparing temporal gradients of recognition accuracy in both directions of LAG. While we use model fits from Equation 3 for visualizing the effects of positive and negative LAG, support for hypotheses **H2** and **H3** will be evaluated based on fitting the model in Equation 2.

$$\begin{aligned} &logit(P(\text{Recog.}_{ij} = 1)) \sim \beta_0 + \beta_1 \cdot \text{TDR} + \beta_2 \cdot \text{LAG}_{pos} + \\ &\beta_3 \cdot \text{LAG}_{neg} + \beta_4 \cdot \text{TDR} \cdot \text{LAG}_{neg} + \beta_5 \cdot \text{TDR} \cdot \text{LAG}_{pos} + u_i \end{aligned} \quad (3)$$

Results

H1: Impact of RTC on recognition accuracy

We first tested whether a target image was more likely to be successfully recognized when the cue image was recognized successfully, presumably triggering retrieved temporal context ($RTC = 1$). The comparison in Figure 3 shows that the recognition accuracy for a target was higher when the cue image was recognized correctly. This effect was consistently observed across participants (paired t-test; $t(7) = 3.79$, $p = 0.0068$), indicating that the successful recognition of the cue image also activates some information about the target image, which is consistent with our assumption that temporal context is retrieved during recognition.

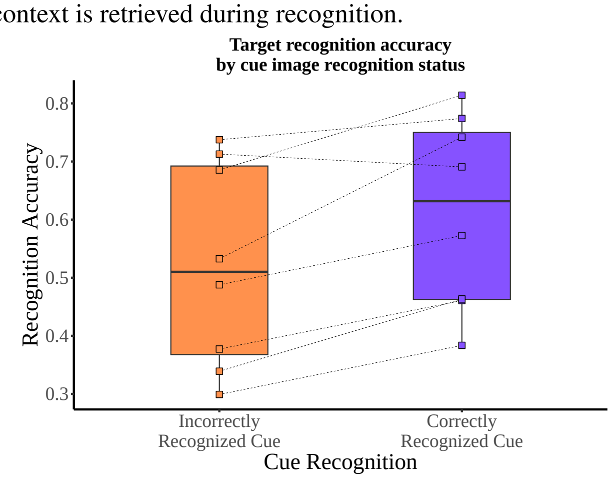

Figure 3: Recognition of a target image improved when it was preceded by a correctly recognized cue image that was previously presented in close proximity to the target. This suggests that successful recognition reinstates temporal context (assumed RTC = 1), thereby enhancing the likelihood of recognizing associated images.

To confirm whether improved target recognition was driven by retrieval of temporal context rather than low-level visual or semantic similarity, we compared CLIP-based representations of cue and target images (Radford et al., 2021). We then checked whether these similarities were different for correctly and incorrectly recognized targets. A paired t-test on subject-level mean similarity scores revealed no significant difference ($t(7) = -0.6092, p = 0.5617$) suggesting that the facilitation effect is unlikely to stem from visual or semantic overlap between cue and target images.

To further quantify this influence of retrieved temporal context (RTC) on target recognition, we used the model in Equation 1. Retrieved temporal context (RTC) had a statistically significant and positive effect ($\beta_1 = 0.514, p < 0.0001$) on the recognition probability of the target images. The estimated odds ratio for RTC $= 1$ was $e^{0.518} \approx 1.68$, suggesting that whenever temporal context is assumed to be retrieved during successful recognition of an image, the likelihood of correct recognition of a subsequent image that was originally encoded in the temporal vicinity of this image increased by approximately 68%, as predicted by TCM.

To further test our assumptions about successful recognition of cue image (RTC $= 1$) reinstating temporal context, we analyzed this boost in recognition accuracy for RTC $= 1$ for different *LAG*s in our contiguity analyses below.

H2 & H3: Impact of RTC, LAG, & TDR and their interactions on target recognition accuracy

In the subsequent model (Equation 2), we incorporated additional covariates (TDR, LAG_{abs}, and their interactions). A likelihood ratio test showed that the inclusion of these additional variables significantly improved the model fit ($\chi^2 = 11.74$, $p = 0.0028$) over the base model (Equation 1).

Table 1: Random Effects Model Results (Equation 2).

Effect	Coef β	SE (β)	z	p-value
Intercept	0.55	0.40	1.38	0.17
RTC	0.93	0.24	3.8	0.000174***
TDR	-0.13	0.05	-2.8	0.00585**
LAG_{abs}	-0.04	0.03	-1.4	0.1627
$TDR : LAG_{\text{abs}}$	0.01	0.004	2.5	0.0138*
$RTC : LAG_{\text{abs}}$	-0.05	0.02	-2.1	0.0384*

Statistically significant terms are marked in Green ; $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$.

In this model, RTC remained a significant predictor ($\beta_1 = 0.931, p < 0.001$), with an odds ratio of ≈ 2.54, indicating a 154% increase in recognition probability. The interaction $RTC : LAG_{abs}$ was also significant, supporting Hypothesis **H2**, i.e., recognition accuracy declines more steeply with LAG_{abs} when $RTC = 1$. While the main effect of LAG_{abs} (n.s.) suggested a general negative trend, its impact was notably stronger in the presence of retrieved temporal context, in line with Hypothesis **H2**. See Table 1 for full results.

TDR showed a significantly negative association with recognition probability ($\beta_2 = -0.131, p = 0.006$), suggesting that as the temporal distance at retrieval (TDR) increases, the retrieved temporal context decays further, resulting in decreased recognition accuracy. Critically, the interaction $TDR : \text{LAG}_{\text{abs}}$ ($\beta_4 = 0.011, p = 0.014$) indicates that the negative effect of LAG_{abs} on recognition accuracy was attenuated as TDR increased, supporting Hypothesis **H3**.

To visualize the effect of *LAG* on recognition accuracy, we modeled the effect of LAG separately for cue images whose temporal context is successfully retrieved (RTC $= 1$). Results for $RTC = 1$ (Equation 3) are shown in Table 2. The coefficients indicate a slightly stronger temporal gradient of target recognition accuracy for negative *LAG* compared to positive *LAG*. Figure 4 shows how the predicted recognition accuracy decreases with *LAG* in both directions. However, this effect is observed only in the presence of successfully retrieved temporal context. None of the terms is statistically significant when the model is rerun for $RTC = 0$. Dotted lines in Figure 4 show the average recognition accuracy of target images, calculated for *LAG* buckets of size 5, averaged across subjects.

Table 2: Asymmetric model results (for correct cue image recognition, assumed RTC = 1; Equation 3).

Effect	Coef β	SE (β)	z	p-value
Intercept	1.88	0.50	3.77	0.000163***
TDR	-0.20	0.07	-2.96	0.003103**
LAG_{neg}	0.13	0.04	3.07	0.002165**
LAG_{pos}	-0.10	0.04	-2.63	0.008553**
$TDR : LAG_{\text{neg}}$	-0.019	0.007	-2.77	0.005593**
$TDR : LAG_{\text{pos}}$	0.012	0.006	1.88	0.059666

Statistically significant terms are marked in Green ; $^{***}p < 0.001$, $^{**}p < 0.01$.

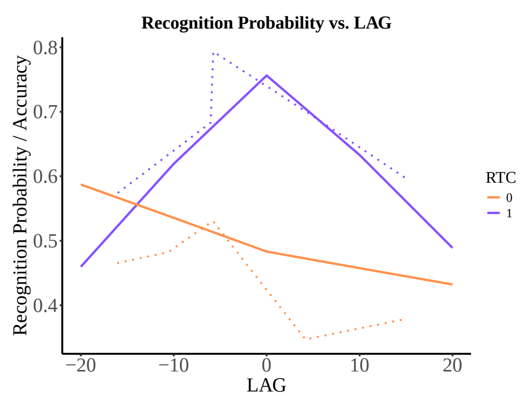

Figure 4: The effect of *LAG* (temporal distance at encoding) on recognition probability for $RTC = 1$ and $RTC = 0$. The solid lines correspond to how the model predicts the effect of *LAG*, and the dotted lines correspond to the average accuracy across subjects in that *LAG* range.

Finally, in Figure 5, we visualize the interaction between TDR and LAG for both positive and negative LAG using the asymmetric model that was run only for $RTC = 1$ images. The coefficient for $TDR : LAG_{neg}$ is negative and statistically significant (Table 2), indicating that the pattern of decreasing recognition accuracy with negative LAG is stronger for smaller compared to larger TDR values (seen as decreasing slopes or flatter fits on the negative side of Figure 5 with increasing TDR). $TDR : LAG_{pos}$ has a positive coefficient and, while not statistically significant, indicates that the negative slope of recognition accuracy for positive LAG decreases with TDR (seen as increasing slopes or flatter fits for positive LAG with increasing TDR).

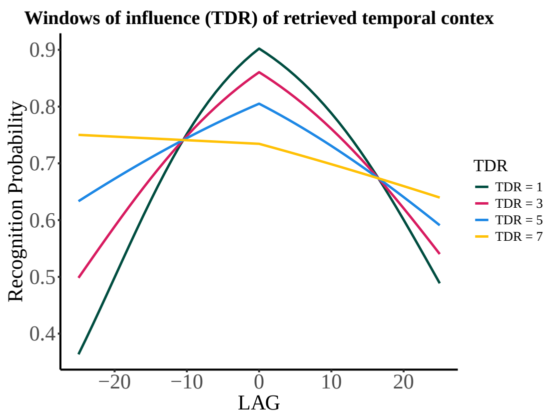

Figure 5: Comparison of the effect of LAG on recognition accuracy for different TDR (temporal distance at retrieval) values. Slopes of the effect of LAG on accuracy are flatter on both sides of $LAG = 0$ with increasing TDR.

Discussion

This study provides behavioral evidence for the influence of temporal context on memory in a continuous image recognition task. Specifically, our results indicate that successful recognition of an image retrieves the temporal context associated with that image, and this retrieved temporal context, in turn, facilitates the recognition of other images that were encoded in temporal proximity to the retrieved image. Furthermore, we show that this facilitatory effect of retrieved temporal context drops off as a function of the temporal lag between the images at encoding. A similar effect of drifting temporal context is also observed at retrieval, where the influence of retrieved temporal context weakens with time.

Our analysis of continuous recognition data addresses several confounds identified previously with free recall data for demonstrations of temporal context retrieval (Folkerts et al., 2018; Hintzman, 2011), and therefore, provides stronger evidence for *obligatory* temporal context retrieval. Our results align with those of Healey (2018), showing that incidental encoding followed by surprise free-recall tests (to eliminate the possibility of specific task-driven memory strategies during encoding) attenuate, but do not completely eliminate, contiguity effects. Cued-recall tasks provide further evidence by showing that intrusions follow similar temporal gradients (Davis et al., 2008). In cued recall, forming associations across word pairs is detrimental to the task, and therefore, any signatures of temporal context would suggest an almost automatic role for temporal context retrieval. Such results indicate that contiguity is not a mere artifact of strategy and that temporal context retrieval might be obligatory.

While contiguity effects in free recall often show a forward asymmetry—where transitions to later items are more probable—our results revealed a qualitatively stronger influence of *negative lag at encoding* (i.e., retrieval benefiting more from earlier temporal neighbors). This pattern contrasts with the canonical forward bias seen in free recall but is more consistent with (Schwartz et al., 2005), who reported symmetric temporal effects in recognition memory. One explanation is that *explicit* free recall tasks really do encourage memory strategies that promote direct associations between items in the forward direction, as Hintzman (2011) argued. Consistent with this suggestion, Healey (2018)'s temporal contiguity plots for *incidental* encoding conditions also show weaker forward asymmetry. Further analyses of larger recognition datasets are required to obtain statistically robust results for the comparison between negative and positive lags. If forward asymmetry is merely an artifact of the explicit free recall task, the temporal context model may need revision while keeping its mechanisms that enable the influence of obligatorily retrieved temporal context.

To summarize, we extend the findings from (Schwartz et al., 2005) by providing evidence that temporal context is retrieved even during image recognition. Whereas Schwartz reported temporal context reinstatement only at high confidence levels, we demonstrate that correct recognition alone is sufficient to trigger temporal context retrieval. Moreover, we analyze a time window of influence at retrieval, showing how retrieved context fidelity decays over time. Our results also offer evidence for a broader temporal window of temporal contextual linkage at encoding.

A key limitation of this study is that it utilized the Natural Scenes Dataset (NSD), which was not specifically designed to investigate temporal context effects. As a result, we applied specific criteria and filtered the dataset, which reduced its size and constrained the scope of our analyses. Another drawback of NSD is the absence of confidence ratings, which turn out to be important for some contiguity analyses (Folkerts et al., 2018). Therefore, this small-scale dataset, while informative, limits the generalizability of our findings. Future work will address some of these limitations by incorporating fMRI data available in the Natural Scenes Dataset to obtain neural evidence for retrieved temporal context, providing a richer understanding of the mechanisms underlying temporal context in episodic memory.

Acknowledgments

We thank Yash Agrawal, Prayush Rathore, Gargi Shukla, Anuska Maity, and Pritha Ghosh for helpful comments on previous versions of this manuscript. We thank IIIT Hyderabad for the Faculty Seed Fund (VS, IIIT/R&D Office/Seed-Grant/2021-22/018). This work makes use of the Natural Scenes Dataset (NSD), a large-scale fMRI dataset collected at ultra-high-field (7T) strength at the Center for Magnetic Resonance Research (CMRR), University of Minnesota. We thank the NSD team for making this resource publicly available. Figures 1 and 2 were created using *Canva*, and the icons used in these figures are also sourced from *Canva*.

References

Allen, E. J., St-Yves, G., Wu, Y., Breedlove, J. L., Prince, J. S., Dowdle, L. T., Nau, M., Caron, B., Pestilli, F., Charest, I., et al. (2022). A massive 7t fmri dataset to bridge cognitive neuroscience and artificial intelligence. *Nature neuroscience*, *25*(1), 116–126.

Bates, D., Maechler, M., Bolker, B., Walker, S., Christensen, R. H. B., Singmann, H., Dai, B., Grothendieck, G., Green, P., & Bolker, M. B. (2015). Package 'lme4'. *convergence*, *12*(1), 2.

Caplan, J. B., Ardebili, A. S., & Liu, Y. S. (2022). Chaining models of serial recall can produce positional errors. *Journal of Mathematical Psychology*, *109*, 102677.

Davis, O. C., Geller, A. S., Rizzuto, D. S., & Kahana, M. J. (2008). Temporal associative processes revealed by intrusions in paired-associate recall. *Psychonomic Bulletin & Review*, *15*(1), 64–69.

DuBrow, S., Rouhani, N., Niv, Y., & Norman, K. A. (2017). Does mental context drift or shift? *Current opinion in behavioral sciences*, *17*, 141–146.

Folkerts, S., Rutishauser, U., & Howard, M. W. (2018). Human episodic memory retrieval is accompanied by a neural contiguity effect. *Journal of Neuroscience*, *38*(17), 4200–4211.

Healey, M. K. (2018). Temporal contiguity in incidentally encoded memories. *Journal of Memory and Language*, *102*, 28–40.

Hintzman, D. L. (2011). Research strategy in the study of memory: Fads, fallacies, and the search for the "coordinates of truth". *Perspectives on Psychological Science*, *6*(3), 253–271.

Howard, M. W., & Kahana, M. J. (2002). A distributed representation of temporal context. *Journal of mathematical psychology*, *46*(3), 269–299.

Jaeger, T. F. (2008). Categorical data analysis: Away from anovas (transformation or not) and towards logit mixed models. *Journal of memory and language*, *59*(4), 434–446.

Kahana, M. J. (1996). Associative retrieval processes in free recall. *Memory & cognition*, *24*(1), 103–109.

Manning, J. R., Polyn, S. M., Baltuch, G. H., Litt, B., & Kahana, M. J. (2011). Oscillatory patterns in temporal lobe reveal context reinstatement during memory search. *Proceedings of the National Academy of Sciences*, *108*(31), 12893–12897.

Manns, J. R., Howard, M. W., & Eichenbaum, H. (2007). Gradual changes in hippocampal activity support remembering the order of events. *Neuron*, *56*(3), 530–540.

Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., & Sutskever, I. (2021). Learning transferable visual models from natural language supervision. *Proceedings of the 38th International Conference on Machine Learning*, *139*, 8748–8763.

Schwartz, G., Howard, M. W., Jing, B., & Kahana, M. J. (2005). Shadows of the past: Temporal retrieval effects in recognition memory. *Psychological Science*, *16*(11), 898–904.

Tulving, E. (1972). Episodic and semantic memory. *Organization of memory/Academic Press*.