Article

Current Biology

Single-trial dynamics of hippocampal spatial representations are modulated by reward value

Highlights

- High reward accelerates reorganization of hippocampal spatial map during learning
- Forward and reverse replay differ in dynamics during learning
- Awake hippocampal replay may not contribute to updating of the spatial map
- Reverse replay relates more strongly to recent experiences than forward replay

Authors

Frédéric Michon, Esther Krul, Jyh-Jang Sun, Fabian Kloosterman

Correspondence

fabian.kloosterman@nerf.be

In brief

Michon et al. show in rats that the hippocampal spatial map is quickly updated during learning of reward-place associations, which is further accelerated by high reward value. They found differing dynamics of forward and reverse replay events during learning but no indication that replay directly influences hippocampal spatial representations.

Michon et al., 2021, Current Biology *31*, 1–13
October 25, 2021
https://doi.org/10.1016/j.cub.2021.07.058

CellPress

Please cite this article in press as: Michon et al., Single-trial dynamics of hippocampal spatial representations are modulated by reward value, Current Biology (2021), https://doi.org/10.1016/j.cub.2021.07.058

Current Biology

CellPress

Article

Single-trial dynamics of hippocampal spatial representations are modulated by reward value

Frédéric Michon,[1,2,3,5] Esther Krul,[1,2] Jyh-Jang Sun,[1,4] and Fabian Kloosterman[1,2,3,4,6,7,*]
[1]NERF, Kapeldreef 75, 3001 Leuven, Belgium
[2]Brain & Cognition, KU Leuven, Tiensestraat 102, 3000 Leuven, Belgium
[3]VIB, Rijvisschestraat 120, 9052 Ghent, Belgium
[4]imec, Remisebosweg 1, 3001 Leuven, Belgium
[5]Present address: Netherlands Institute for Neuroscience, Meibergdreef 47, 1105 Amsterdam, the Netherlands
[6]Twitter: @kloostermanlab
[7]Lead contact
*Correspondence: fabian.kloosterman@nerf.be
https://doi.org/10.1016/j.cub.2021.07.058

SUMMARY

Reward value is known to modulate learning speed in spatial memory tasks, but little is known about its influence on the dynamical changes in hippocampal spatial representations. Here, we monitored the trial-to-trial changes in hippocampal place cell activity during the acquisition of place-reward associations with varying reward size. We show a faster reorganization and stabilization of the hippocampal place map when a goal location is associated with a large reward. The reorganization is driven by both rate changes and the appearance and disappearance of place fields. The occurrence of hippocampal replay activity largely followed the dynamics of changes in spatial representations. Replay patterns became more selectively tuned toward behaviorally relevant experiences over the course of learning via the refined contributions of specific cell subpopulations. These results suggest that high reward value enhances memory retention by accelerating the formation and stabilization of the hippocampal cognitive map and selectively enhancing its reactivation during learning.

INTRODUCTION

Our ability to remember an event depends in part on the behavioral relevance or associated value of the event.[1] Generally, the more rewarding an experience is, the better and the longer it will be remembered.[2–4] Rewarding feedback is known to modulate multiple memory processes. In addition to enhancing memory consolidation,[2,3,5,6] reward has also been shown to promote memory encoding and to accelerate learning during experience.[4,5,7]

The activity of the hippocampus and its interactions with cortical and subcortical brain regions are critical to both the formation and the consolidation of episodic and spatial memory traces.[8] Brain-imaging studies in humans have shown that hippocampal activity and its functional connectivity with other brain regions are modulated by reward magnitude during experience[3,5,7] and are predictive of future memory recall.

In rodents, during active exploration, individual neurons in the hippocampus are preferentially active when the animal crosses a particular location, or place field, in the environment.[9] Collectively, the activity of place cells forms a map-like representation of space. The formation of this cognitive map is experience dependent,[10,11] and it is sensitive to changes in a large range of spatial and other stimuli, a process referred to as "remapping."[12–14] Notably, studies have reported a number of reward-related changes in the firing properties of hippocampal neurons: the accumulation of place fields near newly rewarded locations;[15–17] the modulation of place cell firing rate by reward expectation;[18,19] and the demonstration of a subpopulation of hippocampal neurons that is consistently active at rewarded locations in different environments.[20] These studies also reported that place cells in the dorsal Cornu Ammonis area 1 (CA1) are more sensitive to reward-related changes than place cells in area CA3.[15,19] In addition, reward magnitude modulates the reactivation of experience-related hippocampal neural activity patterns ("replay") during experience.[21,22]

However, few studies have investigated the effect of reward magnitude on the hippocampal place code, and none have reported an effect,[23,24] possibly because rewarded locations were kept stationary. Moreover, the impact of reward on the formation and updating of the hippocampal map remains poorly understood. In this study, we take advantage of a recently developed paradigm[25] to compare the dynamics of the hippocampal code during repeated learning of two rewarded locations associated with different reward size.

RESULTS

Rapid reorganization of spatial representations

Six rats were trained on a dual reward-place association task,[6,25] in which they learned across 5 instruction trials that two locations were associated with small or large reward (Figure 1A). The run speed profile across the 5 instruction trials

Please cite this article in press as: Michon et al., Single-trial dynamics of hippocampal spatial representations are modulated by reward value, Current Biology (2021), https://doi.org/10.1016/j.cub.2021.07.058

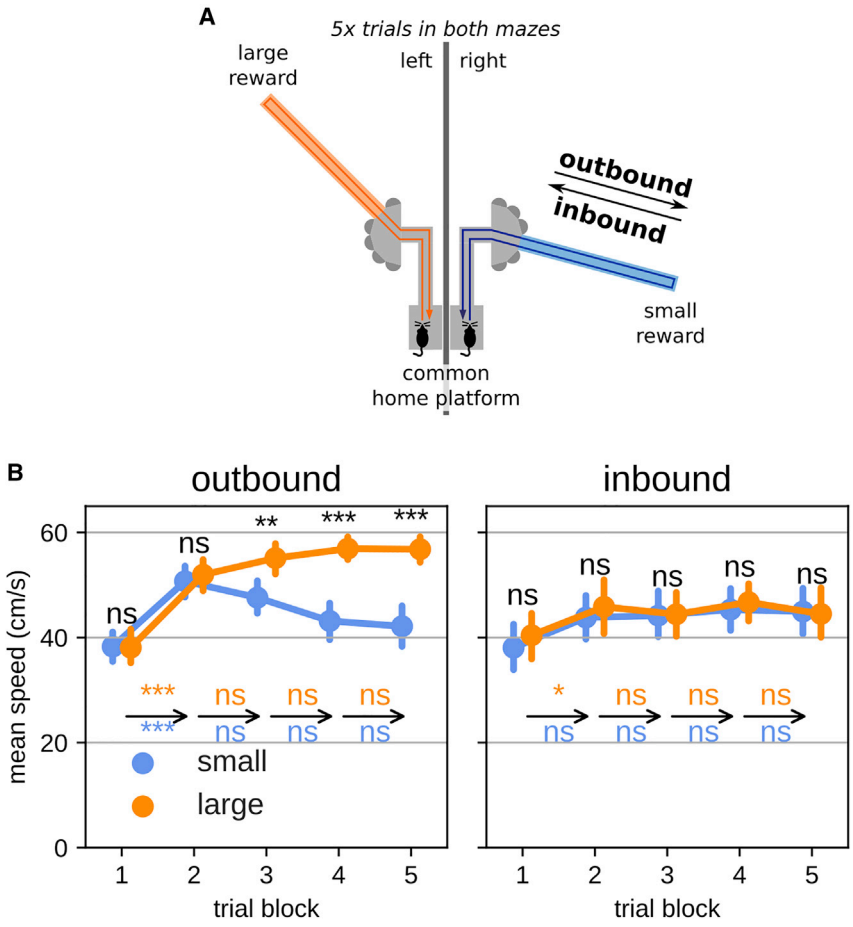

Figure 1. Fast learning of the reward-place associations during training

(A) Schematic representation of the apparatus. Note that the home platform and short piece of track from home to the first bifurcation are shared between left and right environments but are duplicated in the drawing to highlight the separate trials. During training, rats learn to associate a small reward (blue) or a large reward (orange) with a specific arm in the left and right environment. Both arms are visited alternatingly in five trial blocks. Every day, the target arm location, the reward-environment associations, and visiting order were pseudo-randomly assigned.

(B) Average run speed over trial blocks separated by reward conditions and for journeys toward (outbound, left panel) or from (inbound, right panel) the reward platform. Error bars indicate 95% bootstrapped confidence interval. Text annotations: arrows represent the trial-to-trial transitions annotated with the results of statistical comparison (Wilcoxon signed-rank test) separately for large (orange) and small (blue) reward. Per-trial block comparison between reward conditions was performed with Wilcoxon signed-rank test followed with Holm-Sidak correction for multiple tests. ***$p < 0.001$; **$p < 0.01$; *$p < 0.05$; ns, not significant.

indicated that the rats quickly distinguished the small- and large-reward destinations (Figure 1B). On outbound journeys toward the large and small reward, the average run speed steeply increased from trial block 1→2 (Figure 1B, left; Wilcoxon signed-rank test for equal run speed in trial blocks 1 and 2; small: W = 29.00, $p^* = 4.6\times10^{-6}$; large: W = 18.00, $p^* = 1.9\times10^{-6}$). However, although the average speed toward the large reward remained stable, it progressively decreased for the small-reward condition such that the average run speed toward the large reward was significantly higher from trial block 3 onward (Wilcoxon signed-rank test for equal run speed in small and large conditions: trial block 3: W = 140.00, $p^* = 0.0043$; trial block 4: W = 40.00, $p^* = 1\times10^{-5}$; trial block 5: W = 37.00, $p^* = 1\times10^{-5}$). The average run speed on inbound journeys back from the rewarded locations to home mostly remained stable across trial blocks and was not different between reward conditions (Figure 1B, right).

To characterize the dynamics of the hippocampal spatial representation throughout the learning phase, we recorded cell activity from hippocampal subregion CA1 (Figures 2A and 2B). Across sessions, there was no difference between small- and large-reward conditions in the number of place cells and place fields (Figure 2C; Wilcoxon signed-rank test, place cells: $Z = 275.50$, p = 0.71; place fields: $Z = 334.00$, p = 0.79).

We first quantified the firing rate changes at the population level across trial blocks. Given the changes in run speed across trial blocks and the positive relation between speed and firing rate in hippocampal place cells,[26] we corrected the in-field firing rates to remove the contribution of speed (Figures S1A and S1B; see STAR Methods). After correction, partial correlation analysis confirmed that no speed-rate association remained when accounting for reward and trial block (Figure S1C).

For each trial block, we constructed a vector of in-field (speed-corrected) firing rates from all place fields and computed the Spearman correlation coefficient for each pair of trial blocks (Figure 3A). We observed that the population activity for trial blocks 2–5 is highly correlated (r in range 0.45–0.72), whereas the activity in the first trial block has systematically lower correlation with the other trial blocks (r in range 0.14–0.33). We then focused on the population vector correlations for consecutive trial block pairs (Figure 3B). Across trial blocks, the correlation is lowest between trial blocks 1 and 2 (correlation coefficient r and 95% bootstrapped confidence interval; $r_{1\to2} = 0.33$ [0.22, 0.43]) and highest between trials blocks 4 and 5 ($r_{4\to5} = 0.72$ [0.68, 0.76]). The population activity pattern is already highly similar between trial blocks 2 and 3 ($r_{2\to3} = 0.61$ [0.56, 0.66]). We fitted the data with a sigmoid growth curve and determined the total change Δ_r in correlation across trials and the relative change $rel\Delta_r$ between $r_{1\to2}$ and $r_{2\to3}$ as measures of the correlation dynamics (see Figure 3B). The total change Δ_r (mean and 95% confidence interval 0.39 [0.29, 0.52]) was significantly higher than the chance distribution, computed by shuffling for each field the in-field firing rates across trial blocks (Figure 3B, right; Monte-Carlo p = 0.002). Most of the increase in correlation was already attained at $r_{2\to3}$ ($rel\Delta_r = 0.71$ [0.57, 0.83]).

Please cite this article in press as: Michon et al., Single-trial dynamics of hippocampal spatial representations are modulated by reward value, Current Biology (2021), https://doi.org/10.1016/j.cub.2021.07.058

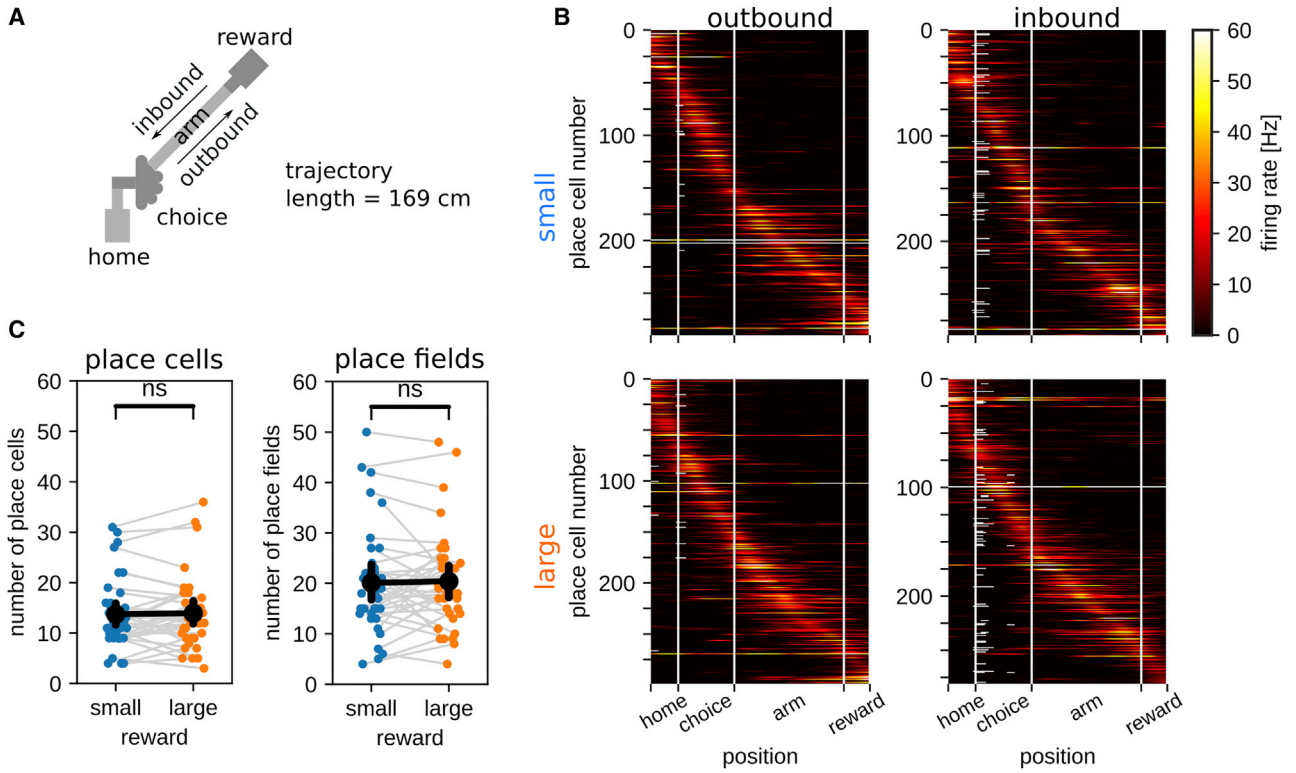

Figure 2. Spatial rate maps

(A) Schematic representing the different maze segments.

(B) Spatial rate maps for all place cells in all recording sessions that are active during locomotion. Cells are ordered by the location (from home to reward platform) of the peak firing rate and separated for outbound (left) and inbound (right) runs in the small-reward environment (top row) and large-reward environment (bottom row).

(C) The number of place cells (left) and place fields (right) for all animals and sessions separated by reward condition. Black points represent the group means. Error bars represent the 95% confidence interval.

These results indicate that the hippocampal spatial representation quickly reorganizes after animals visited the target arm and experienced the associated reward size for the first time.

We next looked at the trial-to-trial population-level activity changes separately for small- and large-reward trials (Figure 3C). The correlation between the first two trials is low for both small and large reward conditions (correlation coefficient and 95% bootstrapped confidence interval; small: $r_{1\rightarrow2} = 0.36$ [0.23, 0.47], large: $r_{1\rightarrow2} = 0.28$ [0.14, 0.47]). However, we observed a difference in the dynamics of the correlations across trials for the two reward conditions. Whereas, for small reward, the trial-to-trial correlation gradually increased toward the end of the instruction phase, for large reward, we found a step change between $r_{1\rightarrow2}$ and $r_{2\rightarrow3}$ and little further increase in later trials. Partial regression analysis confirmed that the population vector correlation changes across trials were not related to variations in run speed that differ between reward conditions (Figures S1D and S1E).

We fitted the data with a sigmoid growth curve and compared Δ_r and $rel\Delta_r$ between reward conditions. For both small and large reward, the total change in correlation Δ_r was significantly higher than chance (Figure 3D; Monte-Carlo p value; small: p = 0.002, large: p = 0.002), and there was no difference between the reward conditions. The relative change in correlation between the first two trial block pairs was, however, significantly higher for large than small reward (Figure 3D; bootstrapped distribution of $rel\Delta_r^{large} - rel\Delta_r^{small}$; 95% confidence interval: [0.16, 0.52]; probability that distribution is different from zero: p = 0.008).

These results indicate that, although spatial representations reorganize and stabilize for both small- and large-reward conditions, this process occurs faster when the reward is large.

We performed the same analysis separately for place fields on the outbound and inbound trajectories (Figure 3E). For both small and large reward, we find a significant increase in correlation over the trial pairs for both outbound and inbound trajectories (Monte-Carlo p value for Δ_r different from chance; outbound small: p = 0.002; outbound large: p = 0.002; inbound small: p = 0.002; inbound large: p = 0.002). There is no difference between outbound and inbound trajectories in how quickly the correlation increases between the first two trial block pairs.

These data suggest that place field activity reorganizes on both inbound and outbound trajectories.

Reorganization is accompanied by predominantly positive rate changes

To identify how the activity in individual place fields evolves, we looked at the in-field rate changes across trial blocks. At the population level, there is a general rate increase toward later trial

Please cite this article in press as: Michon et al., Single-trial dynamics of hippocampal spatial representations are modulated by reward value, Current Biology (2021), https://doi.org/10.1016/j.cub.2021.07.058

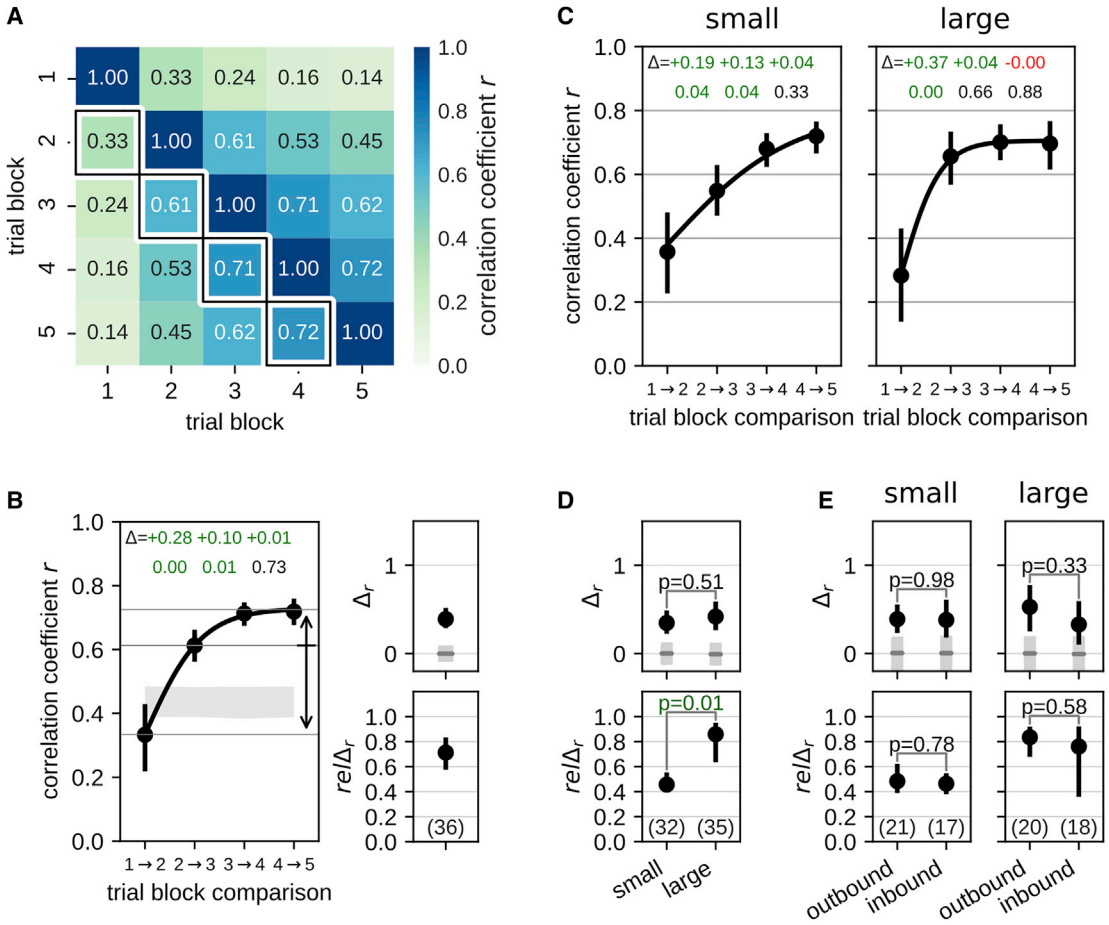

Figure 3. Trial-to-trial correlations show fast stabilization of the spatial code

(A) Pairwise correlations of the in-field firing rate vectors for all place fields and for all trial-block combinations. Note that pairwise correlations between trial blocks 2–5 are higher than correlations involving trial block 1.

(B) (Left) Correlations of the in-field firing rate vectors for sequential trial block pairs (highlighted by black outlines in A). Error bars represent 95% bootstrapped confidence interval. Black line represents fit of growth curve. Note that the correlation between trial blocks 1 and 2 is low, and the correlation in subsequent trial block pairs is high. A one-way repeated-measures analysis of variance by ranks (Friedman test) showed significant differences between trial-block pairs (statistic = 37.60; $p = 3.4\times10^{-8}$). Annotations at the top of the plot indicate the change between adjacent trial pairs and p value of significant pairwise differences (Conover post hoc test with Holm-Sidak p value correction). Coefficients and 95% confidence interval of the fitted growth curve are as follows: $\beta_1 = -0.06$ [−0.30, 0.13]; $\beta_2 = 0.79$ [0.58, 1.04]; and $\lambda = 1.76$ [1.23, 2.39]. Light gray region represents the 95% confidence interval of the population vector correlations after random per-field shuffling of the in-field firing rates across trial blocks (500 permutations). (Right) Mean and bootstrapped 95% confidence interval of two parameters of the fitted curve, total change Δ_r (top), and relative change $rel\Delta_r$ (bottom) are shown. Gray line and light gray region represent the mean and 95% confidence interval of the chance distribution after random permutations of the in-field firing rates across trial blocks. Annotation at the bottom indicates the number of sessions that meet the minimum requirement of 8 place fields.

(C) Correlations of the in-field firing rate vectors separately for small- (left) and large (right)-reward conditions. For both reward conditions, Friedman test indicated significant differences between trial-block pairs (small: statistic = 23.74, $p = 2.8\times10^{-5}$; large: statistic = 20.78, p = 0.00012). Annotations at the top of the plots indicate the change between adjacent trial pairs and p value for significant pairwise differences (Conover post hoc test with Holm-Sidak p value correction). Coefficients and 95% confidence interval of the fitted growth curve are as follows: small: $\beta_1 = -0.04$ [−0.27, 0.17], $\beta_2 = 0.83$ [0.61, 1.07], $\lambda = 0.80$ [0.54, 1.15]; large: $\beta_1 = -0.14$ [−0.46, 0.16], $\beta_2 = 0.84$ [0.53, 1.19], $\lambda = 2.58$ [1.44, 3.70].

(D) Mean and bootstrapped 95% confidence interval of two parameters of the fitted curve, total change Δ_r (top), and relative change $rel\Delta_r$ (bottom), separated for small and large reward. Gray lines and light gray regions represent the mean and 95% confidence interval of the chance distribution after random permutations of the in-field firing rates across trial blocks. Annotation at the bottom indicates the number of sessions that meet the minimum requirement of 8 place fields.

(E) Similar to (D) with comparison of outbound and inbound trajectories for both reward sizes.

See also Figures S1 and S2.

blocks for both large- and small-reward conditions (Figure 4A). Because we corrected the in-field rates for running speed, the increased rates are not due to the animals running faster in later trial blocks, which we confirmed with partial regression analysis (Figure S1C).

Pairwise comparison of in-field firing rate between consecutive trials showed a significant step increase between trial block 1 and 2 in the large-reward condition. For the small reward, the in-field rate appears to increase toward trial block 5 but without significant pairwise differences.

Please cite this article in press as: Michon et al., Single-trial dynamics of hippocampal spatial representations are modulated by reward value, Current Biology (2021), https://doi.org/10.1016/j.cub.2021.07.058

CellPress

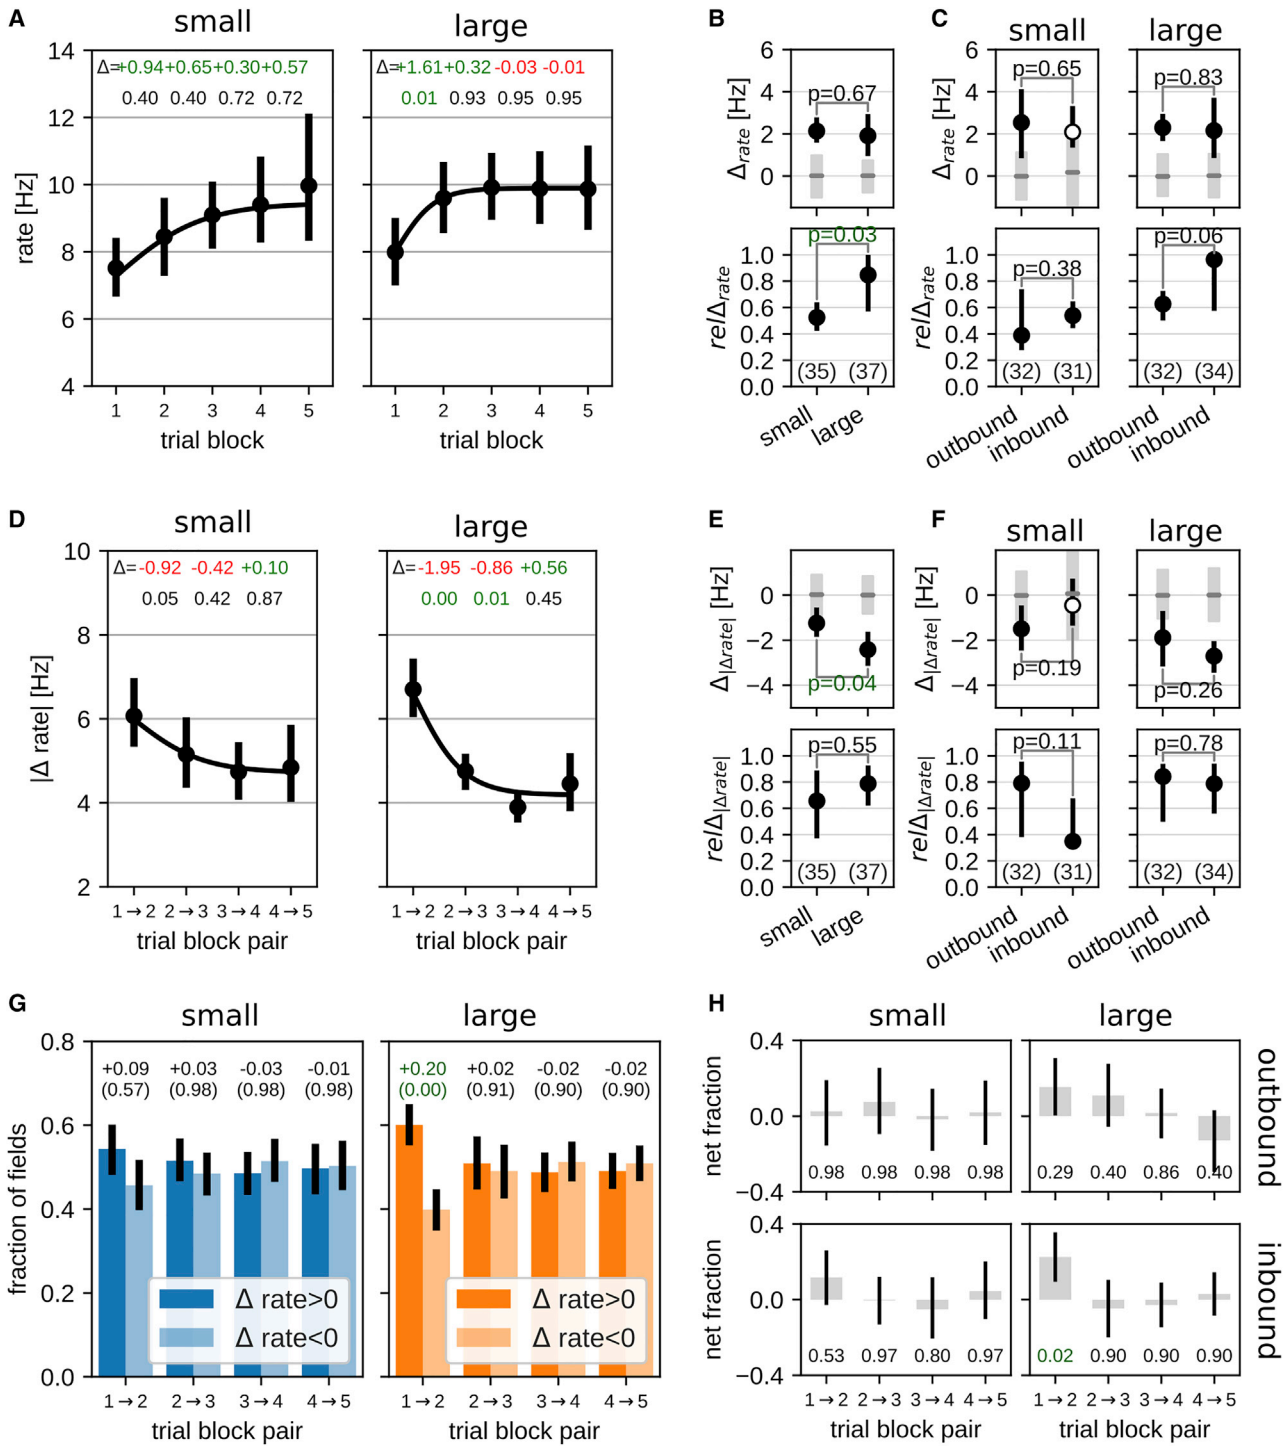

Figure 4. Trial-to-trial changes of in-field firing rate

(A) Average per-session in-field firing rate across trial blocks for small and large reward. Error bars indicate 95% bootstrapped confidence interval. For both reward conditions, Friedman test indicated significant differences between trial-block pairs (small: statistic = 6.67, p = 0.083; large: statistic = 10.30, p = 0.016). Annotations at the top of the plots indicate the change between adjacent trials and p value for significant pairwise differences (Conover post hoc test with Holm-Sidak p value correction).

(B) Mean and bootstrapped 95% confidence interval of two parameters of the fitted curve, total change Δ_{rate} (top), and relative change $rel\,\Delta_{rate}$ (bottom), separated for small and large reward. Gray lines and light gray regions represent the mean and 95% confidence interval of the chance distribution after random permutations of the in-field firing rates across trial blocks. Annotation at the bottom indicates the number of sessions that meet the minimum requirement of 4 place fields.

(legend continued on next page)

Please cite this article in press as: Michon et al., Single-trial dynamics of hippocampal spatial representations are modulated by reward value, Current Biology (2021), https://doi.org/10.1016/j.cub.2021.07.058

The total change in rate across trials Δ_{rate} determined by fitting a sigmoid growth curve was similar for both reward sizes and significantly higher than expected by chance (Figure 4B; Monte-Carlo p value; small: p = 0.002; large: p = 0.002). However, the maximum rate was reached faster for large reward than small reward, as was clear from a comparison of $rel\Delta_{rate}$ (Figure 4B; bootstrapped distribution of $rel\Delta_{rate}^{large} - rel\Delta_{rate}^{small}$; 95% confidence interval: [0.04, 0.52]; probability that distribution is different from zero: p = 0.028).

For both reward groups, there was no difference between outbound and inbound trajectories in the total firing rate change, although $\Delta_{rate}^{small,inbound}$ was not significantly different from the chance distribution. For the large reward condition, we further found a tendency for the total rate change to be larger in the inbound direction than the outbound direction (Figure 4C).

These results indicate an overall increase in in-field firing rate across trials at the population level, but we noticed a high variability that could indicate a larger heterogeneity of rate dynamics at the single-cell level. For this reason, we analyzed separately the magnitude of the rate changes $|\Delta rate|$ and the fraction of fields with rate increase or decrease across trial blocks.

Consistent with the population vector analysis, the largest per-session absolute rate changes are observed from trial block 1 to 2 for small- and large-reward conditions (Figure 4D). Partial regression analysis confirmed that the absolute rate changes across trials were not related to variations in run speed that differ between reward conditions (Figures S1D and S1F).

For subsequent trial block pairs, the absolute rate changes progressively declined but at different extents for small and large reward. We fitted the data with a sigmoid growth curve and compared the total and relative changes between reward conditions. The total decrease in absolute rate change across trial block pairs, $\Delta_{|\Delta rate|}$, was more pronounced for large than small reward (Figure 4E), which was mainly driven by the absence of a reduction in absolute rate change for small reward in the inbound direction (Figure 4F).

We next looked at whether the absolute rate changes reflect both increases and decreases of in-field firing rate. For this, we computed the per-session fractions of place fields with rate increase or decrease (Figure 4G). For large reward, but not small reward, we found a significant bias toward rate increases from trial block 1 → 2 (Wilcoxon signed-rank test between increasing and decreasing rate fractions; small: W = 220.50, p* = 0.57; large: W = 80.00, p* = 0.0023). For other trial block pairs, we observed an equal proportion of rate increase and rate decrease. The bias for rate increases was present for both outbound and inbound directions in the larger reward condition although reached significance only for inbound (Figure 4H).

A subset of place fields emerges or vanishes after initial trials

We noted that a subset of place fields was not active (i.e., zero in-field rate) during one or more trials. We found a clear overrepresentation of place fields that were not active in the first 1–3 trials but became active afterward (emerging fields) and place fields that were initially active but then disappeared after trials 1–3 (vanishing fields; Figure 5A).

For both small- and large-reward condition, there was a significantly larger fraction of emerging place fields than vanishing place fields (Figure 5B; Wilcoxon test for equal fractions of emerging and vanishing fields; small: W = 77.00, p = 0.012; large: W = 98.00, p = 0.029).

For small reward size, the larger fraction of emerging fields was only present for the outbound trajectory. For large reward size, there was a trend for more emerging fields on both outbound and inbound trajectories, although it only reached significance for the inbound direction (Figure 5C).

We next looked at when place fields emerged or vanished. Emerging fields did not appear uniformly on trial blocks 2–4 (chi-square test for null hypothesis of uniform distribution; small: statistic = 28.07, p = 8×10^{-7}; large: statistic = 52.22, p = 4.6×10^{-12}), but rather most emerging fields appeared after the first trial block (Figure 5D). This pattern of field emergence was the same for small and large reward. For vanishing fields in the large-reward condition, a similar bias toward disappearing after the first trial was observed (Figure 5E; chi-square test; statistic = 17.29; p = 0.00018). For small reward, however, the distribution is not statistically different from uniform (chi-square test; statistic = 2.36; p = 0.31), and the highest percentage of vanishing place fields disappears in trial block 3.

(C) Similar to (B) with comparison of outbound and inbound trajectories for both reward sizes. An open circle indicates that the parameter value is not different from the chance distribution.

(D) Average per-session magnitude of in-field rate changes for pairs of trial blocks. Error bars indicate 95% bootstrapped confidence interval. For both reward conditions, Friedman test indicated significant differences between trial-block pairs (small: statistic = 17.85, p = 0.00047; large: statistic = 25.05, p = 1.5×10^{-5}). Annotations at the top of the plots indicate the change between adjacent trial pairs and p value for significant pairwise differences (Conover post hoc test with Holm-Sidak p value correction). Coefficients and 95% confidence interval of the fitted growth curve are as follows: small: β_1 = 7.26 [6.17, 8.57], β_2 = −2.54 [−3.83, −1.40], λ = 1.53 [0.42, 2.83]; large: β_1 = 9.04 [7.62, 10.34], β_2 = −4.85 [−6.40, −3.26], λ = 2.12 [1.39, 3.25].

(E) Mean and bootstrapped 95% confidence interval of two parameters of the fitted curve, total change $\Delta_{|\Delta rate|}$ (top), and relative change $rel\Delta_{|\Delta rate|}$ (bottom), separated for small and large reward. Gray lines and light gray regions represent the mean and 95% confidence interval of the chance distribution after random permutations of the in-field firing rates across trial blocks. Annotation at the bottom indicates the number of sessions that meet the minimum requirement of 4 place fields.

(F) Similar to (E) with comparison of outbound and inbound trajectories for both reward sizes. An open circle indicates that the parameter value is not different from the chance distribution.

(G) Average fraction of fields with increasing (dark color) or decreasing (light color) in-field firing rate. Error bars indicate 95% bootstrapped confidence interval. Text annotations: top row, fraction difference between fields with increasing and decreasing in-field firing rate; bottom row, Holm-Sidak corrected p value for Wilcoxon signed-rank test between increasing and decreasing rate fractions.

(H) Net fraction of fields with increasing versus decreasing in-field firing rate across trial pairs for small or large reward and outbound or inbound trajectories. Error bars indicate 95% bootstrapped confidence interval. Text annotation at the bottom indicates Holm-Sidak corrected p value for Wilcoxon signed-rank test between increasing and decreasing rate fractions.

See also Figures S1 and S3.

Please cite this article in press as: Michon et al., Single-trial dynamics of hippocampal spatial representations are modulated by reward value, Current Biology (2021), https://doi.org/10.1016/j.cub.2021.07.058

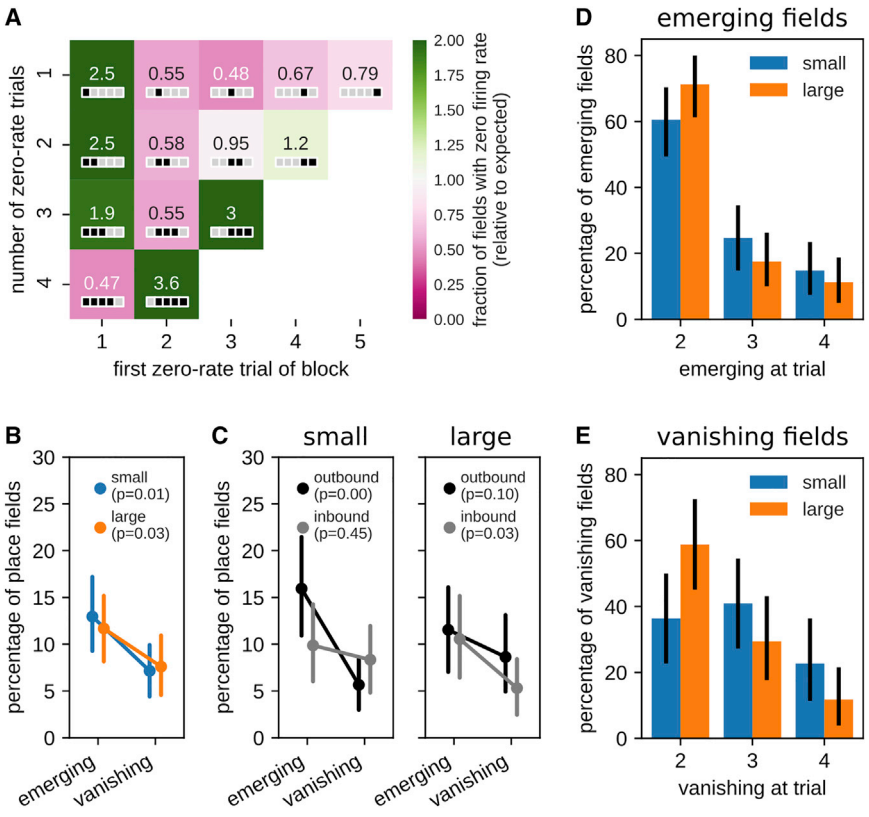

Figure 5. A subset of place fields emerges or vanishes after the initial three trial blocks

(A) The number of place fields with a string of *n* zero-rate trials starting at trial block *t* was expressed relative to the expected number place fields with *n* zero-rate trial blocks (assuming a uniform distribution of zero-rate occurrences across the 5 trial blocks). For example, the value of 2.5 at (1,1) indicates that it is more likely than expected that a single zero-rate trial occurs in the first trial block. For each cell in the matrix, the inset depicts the pattern of zero-rate trial blocks (black and gray squares represent zero-rate and non-zero-rate trial blocks, respectively).

(B) Percentage of emerging and vanishing place fields for small- and large-reward conditions. Note that there are significantly more emerging place fields than vanishing place fields for both reward conditions, but there is no difference between small and large reward.

(C) Percentage of emerging and vanishing place fields separately for outbound and inbound trajectories for small (left) and large (right) reward sizes.

(D and E) Distribution of trial blocks at which place fields emerge (D) or vanish (E). For both small- and large-reward conditions, place fields emerge predominantly in trial block 2. For the large-reward condition, place fields predominantly disappear in trial block 2. For the small-reward condition, the distribution is not statistically different from uniform, and the highest percentage of place fields disappear in trial block 3.

See also Figures S2 and S3.

Overall, these results indicate that the appearance of place fields occurs rapidly after the first trial. Furthermore, the slower dynamics of place-field disappearance in the small-reward condition could contribute to the overall slower stabilization of the small-reward spatial representation.

To test whether the spatial representation changes are explained by the emerging and vanishing fields, we repeated the population vector correlation and rate-change analysis but now excluding emerging and vanishing fields (Figures S2 and S3). Overall, the trial-to-trial dynamics of the population vector correlations and rate changes persist and are thus not solely explained by the emergence or disappearance of place fields.

Place-cell activity in SWRs is modulated by experience, directionality, and reward size

Periods of consummatory behavior are characterized by the reactivation of hippocampal place cell assemblies (replay).[27] Hippocampal replay is modulated by both experience and reward.[22,28] In the following analyses, we ask how replay activity is modulated by experience at the trial level during learning of the reward-place associations and whether replay could be a driving factor for the changes in the spatial representations as shown above.

First, we examined the occurrence of sharp-wave ripples (SWRs), electrophysiological markers of putative replay. The large majority of SWRs occurred while the animal was consuming the large reward, and only very few occurred when consuming the small reward or at other locations (Figure 6A; Table S1; mean ± SEM [percentage ± SEM] number of SWRs per session: large-reward platform, 78.6 ± 5.1 [88.0% ± 1.4%]; small-reward platform, 2.1 ± 0.7 [2.1% ± 0.6%]). The higher incidence of SWRs at the large-reward site can be explained by the fact that the animals spent more time at this location than at the small-reward site (mean [99% confidence interval]; large: 49.75 s [48.10, 51.50]; small: 6.44 s [5.92, 6.99]).

Across trials, the number of SWRs trended higher for later trial blocks (Figure 6B; linear regression; slope = 0.89 events per trial block; R = 0.15; p = 0.065) without increase in total time spent immobile (Figure S4A). Pairwise comparison of consecutive trial blocks did not reveal significant trial-to-trial changes in the number of SWRs.

We next asked how individual place cells change their activity within SWRs over trial blocks. We computed the proportion (p) of SWRs in which a place cell fired at least one spike. Only contextual place cells with place fields on the path to either large or small reward (p_{small} and p_{large}) were considered, and cells with fields on both paths or with fields in the common home were excluded. Overall, cells with fields on the path to the large reward were active in a higher proportion of SWRs than cells with fields on path to small reward (Figure 6C), consistent with a replay bias for the large reward environment. However, p_{large} decreased across trials (linear regression; slope = −0.01 per trial block; R = −0.10; p = 0.0028), such that p_{large} and p_{small} significantly differed in all trial blocks except the last.

Please cite this article in press as: Michon et al., Single-trial dynamics of hippocampal spatial representations are modulated by reward value, Current Biology (2021), https://doi.org/10.1016/j.cub.2021.07.058

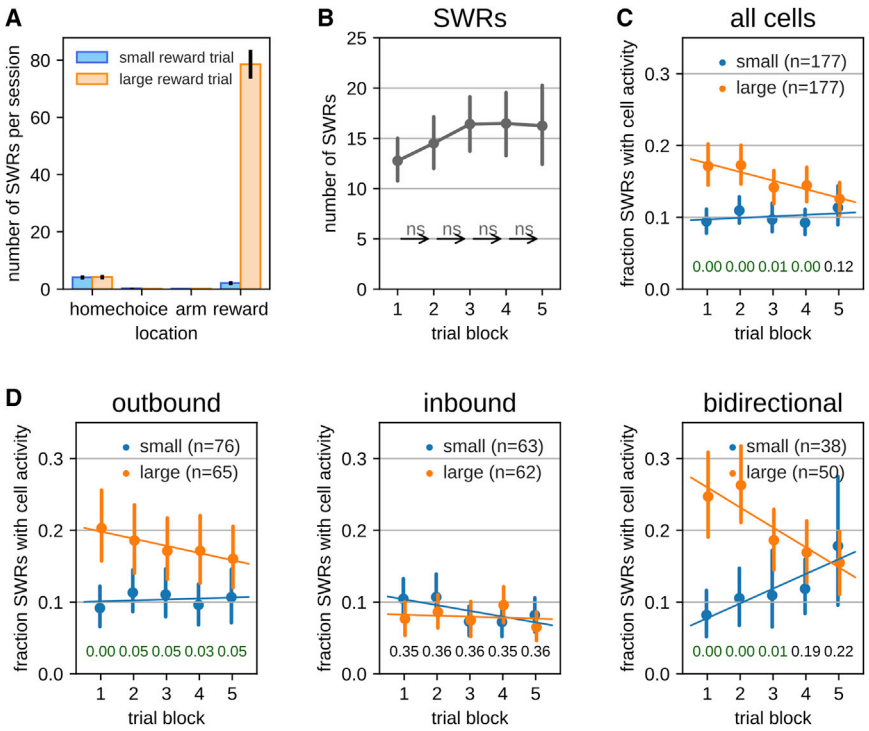

Figure 6. Place-cell activity during SWRs changes over trial blocks

(A) Average number of SWRs per session as a function of location in the large-reward and small-reward environments. Note that the large majority of SWRs occur when rats visit the large-reward platform.

(B) The number of SWRs per trial block. Arrows represent the trial-to-trial transition annotated with the significance level. None of the trial-to-trial changes were significant (Wilcoxon signed-rank test, trial block 1→2: statistic = 160.00, p* = 0.51; trial block 2→3: statistic = 134.50, p* = 0.26; trial block 3→4: statistic = 217.50, p* = 1; trial block 4→5: statistic = 216.00, p* = 1).

(C) The fraction of SWRs during which place cells are active, separated for cells with fields in the small- and large-reward environments. Text annotations at the bottom indicate the corrected p value for Mann-Whitney U tests between large- and small-reward conditions. Error bars indicate 95% bootstrapped confidence intervals. Lines represent least-squares linear fit.

(D) Same as (C) for place cells with fields in the outbound direction (left), in the inbound direction (middle), or in both directions (right).

See also Figure S4 and Table S1.

We next looked at the activity in SWRs separately for unidirectional and bidirectional place cells. For unidirectional place cells, p_{large} is significantly higher than p_{small} across all trial blocks only for cells that are active in the outbound run (i.e., toward the reward), but not for cells active only in the inbound run (Figure 6D). There was a negative trend in p_{large} across trial blocks for the outbound direction, but this was not significant (linear regression; slope = −0.01 per trial block; R = −0.07; p = 0.18). For place cells active in the inbound direction in the small-reward environment, p_{small} modestly but significantly decreased across trial blocks (linear regression; slope = −0.01 per trial block; R = −0.11; p = 0.049). Interestingly, the largest changes across trial blocks were seen for bidirectional cells with place fields in both outbound and inbound directions (Figure 6D). We found an increase of p_{small} and decrease of p_{large} across trial blocks (linear regression, p_{large}: slope = −0.03 per trial block, R = −0.22, p = 0.00053; p_{small}: slope = 0.02 per trial block, R = 0.16, p = 0.026), with a significant difference between p_{small} and p_{large} in the first three, but not the last two, trial blocks.

The changes observed in the contribution of bidirectional cells to SWR activity was not directly linked to their activity during run, as virtually none of the place cells active on both running directions had emerging or vanishing fields throughout learning (Figure S4B).

Overall, these results are consistent with a bias toward reverse replay (i.e., activation of cells with fields on the outbound trajectory) in the large-reward condition and point to a refinement of SWR cell activation that increases spatial and directional specificity.

Replay activity is modulated by experience and reward size

To quantify replay events during periods of immobility, we used Bayesian neural decoding to detect spatial representations during SWR bursts. Previously, we used an encoding model that incorporates hippocampal activity from all trials (trial-average model).[6] That analysis showed a strong bias for reverse replay of the outbound trajectory toward the large reward. Here, we investigated how the expression of replay changes across trial blocks.

The trial-average model assumes that the spatial tuning of hippocampal place cells is constant and reduces error in the model due to variations in spiking. Because we showed that place-field activity systematically changes across trial blocks with different dynamics for small- and large-reward conditions, the trial-average model may not be the best model to analyze replay events on individual trial blocks. For this reason, we performed analysis of replay with per-trial block encoding models (single-trial model; Figures 7, S5, and S6).

We analyzed how the number of replay events varied across trial blocks. Overall, the number of $replay_{large}$ events is higher than the number of $replay_{small}$ for every trial block (Figure 7B; Table S1). Because the large majority of SWRs and replay events occurred at the large-reward platform (Table S1), by definition, most $replay_{small}$ are remote and $replay_{large}$ are predominantly local.[29,30] We find a gradual and significant increase in the number of replay events that express trajectories in the small- and large-reward environments (linear regression, $replay_{large}$: slope = 0.30 events per trial block, R = 0.20, p = 0.0076; $replay_{small}$: slope = 0.18 events per trial block, R = 0.19, p = 0.013). In addition, we observed a small significant jump in $replay_{large}$ events from trial block 1→2 (Wilcoxon signed-rank test: W = 128.00, p* = 0.014; Figure 7B).

We further explored the dynamics across trials of reverse and forward replay (characterized as having a decoded direction

Please cite this article in press as: Michon et al., Single-trial dynamics of hippocampal spatial representations are modulated by reward value, Current Biology (2021), https://doi.org/10.1016/j.cub.2021.07.058

CellPress

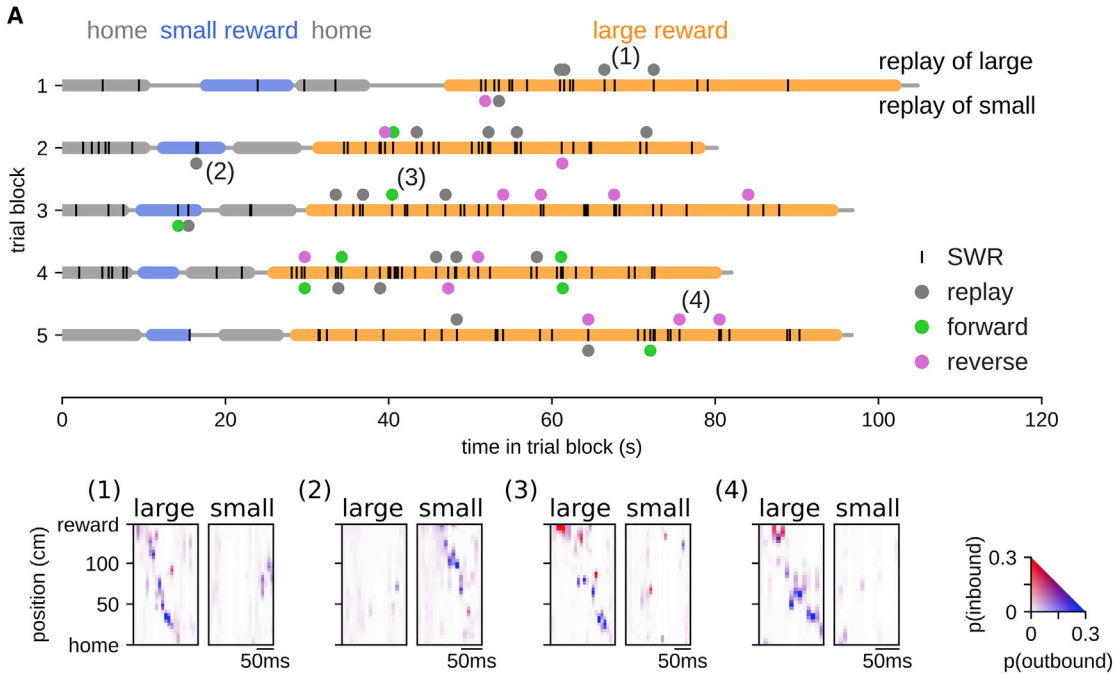

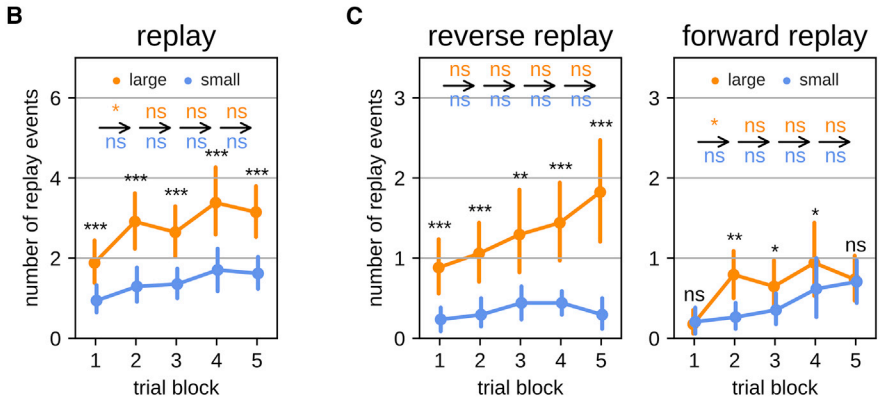

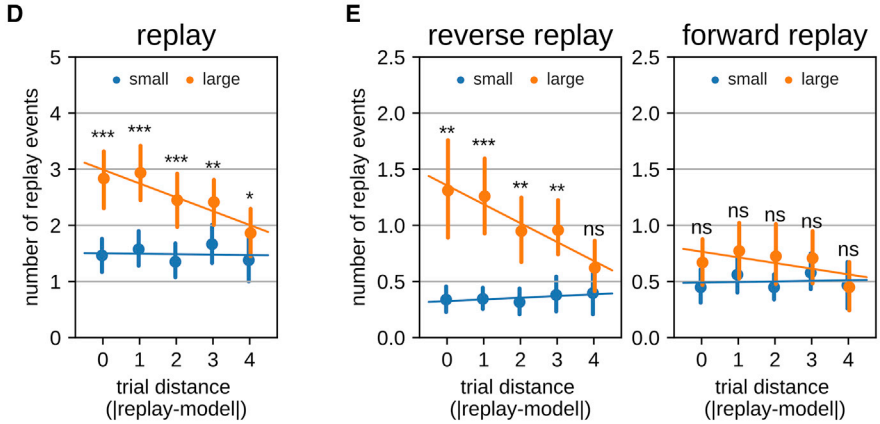

(legend on next page)

Please cite this article in press as: Michon et al., Single-trial dynamics of hippocampal spatial representations are modulated by reward value, Current Biology (2021), https://doi.org/10.1016/j.cub.2021.07.058

opposite or similar to that of the trajectory fit, respectively; Figure 7C). We found that the number of $replay^{reverse}_{large}$ events, but not $replay^{reverse}_{small}$ events, significantly increased across trial blocks (linear regression, $replay^{reverse}_{large}$: slope = 0.226 events per trial block, R = 0.21, p = 0.0053; $replay^{reverse}_{small}$: slope = 0.03 events per trial block, R = 0.07, p = 0.37) without sudden trial-block-to-trial-block changes. The average occurrence of $replay^{reverse}_{large}$ events was consistently higher than $replay^{reverse}_{small}$ across all trial blocks (Wilcoxon ranked-test, trial block 1: W = 105.00, p* = 0.00013; trial block 2: W = 119.00, p* = 0.00022; trial block 3: W = 183.00, p* = 0.0021; trial block 4: W = 117.50, p* = 0.00022; trial block 5: W = 77.50, $p^* = 3.6\times10^{-5}$).

The number of forward replay events increased for both reward conditions throughout the trial blocks (linear regression, $replay^{forward}_{large}$: slope = 0.13 events per trial block, R = 0.19, p = 0.014; $replay^{forward}_{small}$: slope = 0.14 events per trial block, R = 0.25, p = 0.0011). The average occurrence of $replay^{forward}_{small}$ did not show large changes on a trial-block-to-trial-block basis. In contrast, the average number of $replay^{forward}_{large}$ events strongly increased from trial block 1→2 (Wilcoxon signed-rank test; trial block 1 → 2: W = 130.50; p* = 0.015). As a result, the difference between the number of $replay^{forward}_{large}$ and $replay^{forward}_{small}$ events was significant on trial block 2 (Wilcoxon signed-rank test: W = 155.50; p* = 0.0028).

Reverse replay of the large environment is dominated by inbound trajectories (i.e., from reward platform toward home) and reactivation of cells with outbound place fields (Figure S6, left). The gradual increase across trials was only observed for the reverse inbound replay trajectories and not the few reverse outbound replay trajectories. Forward replay of the large environment is also dominated by inbound trajectories that correspond to reactivation of cells with inbound place fields (Figure S6, right). The step-like increase between trial blocks 1 and 2 was only present for forward inbound trajectories and not forward outbound trajectories.

The bias for inbound replay trajectories was less pronounced for the small-reward environment but still present for forward replay (Figure S6).

The dynamics of replay occurrence reported above for the single-trial model were similar to those observed with the trial-average encoding model, except that, for the latter, there was no significant increase of $replay^{reverse}_{large}$ events across trial blocks (Figure S7). To shed more light on this difference, we quantified how the number of replay events depended on the trial that was used to build the single-trial encoding model (Figures 7C and S5A–S5C). For the large-reward condition, we noted that the number of detected replay events on a given trial was highest when the place-cell activity on the same trial was used to build the model. The number of detected replay events gradually declined when more distant trials in the trial sequence were used to build the model (linear regression; slope = −0.25 events per trial block; R = −0.26; p = 0.0018). In contrast, the number of detected $replay_{small}$ events remained stable (linear regression; slope = −0.01 events per trial block; R = −0.01; p = 0.89).

The observed decrease in number of detected events as a function of trial distance between decoding model and replay occurrence of the large environment was specific to $replay^{reverse}_{large}$ (linear regression, $replay^{reverse}_{large}$ slope = −0.17 events per trial block, R = −0.27, p = 0.00089; $replay^{forward}_{large}$ slope = −0.05 events per trial block, R = −0.11, p = 0.19). These results suggest that reverse replay events may be more tightly related to single experiences than forward replay events.

The changes in contribution of single units to SWRs across trials did not directly relate to changes in the overall characteristics of hippocampal replays. On average, event duration, replay trajectory span, and trajectory fit scores remained stable throughout trial blocks for all conditions, apart from a significant increase of the fit scores for replays of the small-reward environment (Figures S5D–S5F).

Taken together, these results suggest that subsets of hippocampal replay events are differently modulated by reward and experience. Reverse replay activity was strongly modulated by reward size, and forward replay activity increased with experience, with a distinct reward effect in the second trial block following the large changes in the place-field activity that we showed above.

DISCUSSION

In this study, we explored the influence of reward value during learning on hippocampal representations in a paradigm where rats repeatedly learn reward-place associations in a familiar setting. We demonstrated quick reorganization of hippocampal place-cell activity over the course of a few trials, which was accompanied by both place field turnover and rate changes. A novel finding in our study is that the updating of the spatial representation was accelerated if animals received a large reward, which may be a neural correlate of the known enhancement of learning speed by reward value.[4,5,7] Hippocampal replay during reward consumption was dominated by a gradually increasing number of reverse replays of the path to the large reward and may have contributed to the faster reorganization of the place code. Detailed analysis of the trial dynamics of reverse and forward replay events indicated, however, that multiple factors

Figure 7. Changes in replay over training and between reward conditions

(A) Examples of identified replay events in one session.

(B) Number of $replay_{large}$ and $replay_{small}$ events across trial blocks. There is a significantly higher number of $replay_{large}$ events for all trial blocks (Wilcoxon signed-rank test: trial block 1: W = 138.00, p* = 0.014; trial block 2: W = 87.00, $p^* = 6.5\times10^{-5}$; trial block 3: W = 117.50, p* = 0.0034; trial block 4: W = 126.00, p* = 0.014; trial block 5: W = 92.00, p* = 0.0025). Arrows represent the trial-to-trial transitions annotated with the significance level separately for the two reward groups.

(C) Similar as (B) for the number of reverse (left) and forward (right) replay events across trial blocks.

(D) Number of $replay^{forward}_{large}$ and $replay^{forward}_{small}$ events across trial blocks as a function of the distance between trial block used to create the encoding model and the trial block of replay occurrence.

(E) Similar as (D) for reverse (left) and forward (right) replay.

For (B)–(E), error bars indicate 95% bootstrapped confidence interval. Lines represents the least squares regression fit. For each trial block, significance level for the difference between the two reward groups is indicated above the data points. All pairwise comparisons were performed with Wilcoxon signed-rank test followed by Holm-Sidak correction for multiple tests. ***p < 0.001; **p < 0.01; *p < 0.05. See also Figures S5–S7 and Table S1.

Please cite this article in press as: Michon et al., Single-trial dynamics of hippocampal spatial representations are modulated by reward value, Current Biology (2021), https://doi.org/10.1016/j.cub.2021.07.058

govern the expression of awake replay and that subsets of replay events may follow the spatial map dynamics rather than actively contribute to the reorganization.

High reward value accelerates reorganization and stabilization of the spatial map

For the environment that was associated with large reward, we observed a faster reorganization and stabilization of the spatial map. Notably, changes in hippocampal spatial representation resulted from the appearance and disappearance of place fields and the modulation of firing rate. The largest change occurred between the first two trials in the large reward environment, with an overall increase in within-field firing rate associated with a higher turnover of active fields. In contrast, the spatial representation of the small-reward environment changed more gradually throughout learning.

Our observations are consistent with reports in the literature showing that environmental changes,[31,32] exposure to another environment,[12,13,33] and spatial learning[15,34,35] induce a fast partial remapping of the spatial representation in the hippocampus,[32] with the modulation of within-field firing rate (rate remapping) and the formation of new fields and disappearance of existing ones (global remapping) for different subgroups of neurons.[12–14] Consistent with the accumulation of place-field activity at novel or changed goal locations,[15,16,36] the ratio of emerging to vanishing fields was highest at the home and reward locations. However, we did not observe a difference in the final representations of the environments associated with large and small rewards, similar to what has been reported previously.[24] This suggests that salient goal locations are overrepresented in the dorsal CA1 place code during learning, regardless of their associated reward value.

Dynamics of reverse and forward replay events

Accompanying the rapid spatial map dynamics in the large-reward environment, we observed a consistent bias of reverse replay. The number of reverse replay events progressively increased throughout training, suggesting a combined effect of reward and experience on reverse replay activity. The bias arises in part because of the reward size[22] but also because animals spent much longer at the large-reward platform and awake replay is strongly linked to the current location of the animal.[29] Therefore, most replay events occurred while the animals were consuming the large reward and represented trajectories to or from the large-reward platform, although a smaller fraction of these replay events represented remote replays[30] of the small-reward environment. Replay of place-cell sequences has been linked to enhanced synaptic plasticity and long-term storage of the sequential structure in the network, and thus the higher incidence of reverse replay of the large-reward environment would be consistent with a contribution to the faster updating and stabilization of the spatial representations.

When considering the dynamics of reverse replay of the outbound and inbound maps separately, a more nuanced picture emerges. Reverse replay almost exclusively involves the outbound place fields, even though no difference was found between the dynamics of the place fields in outbound and inbound directions. This discrepancy suggests that the updating of the spatial representations does not require reverse replay, although we cannot exclude that different mechanisms are at play for inbound and outbound map changes.

Forward replay events predominantly used the inbound map, and the trial-to-trial dynamics followed the spatial map changes for both reward sizes. Forward replay of the large-reward environment showed a step-like increase from trial 1 to 2, and forward replay of the small-reward environment showed a gradual increase across trials. These results are consistent with the fact that replay activity is dependent on experience[28,37] and do not support a model in which awake replay activity directly influences hippocampal representations during learning.

The different dynamics of reverse and forward replay occurrence indicate that likely different mechanisms contribute to their generation and support the hypothesis that subsets of hippocampal replay play separate roles in memory processes.[38–42] Our results are consistent with a model in which reverse replay enhances reward-related learning by retroactively strengthening cell assemblies active prior to obtaining a larger reward.[22] Indeed, reverse replays more closely relate to single-trial maps rather than more "matured" representations, as illustrated by their occurrence after a unique lap[43] and uniquely modulated by reward.[22] These observations also suggest that reverse replays are more strongly driven by external inputs, such as from cortical areas[44,45] or reward-responsive neurons present in the ventral tegmental area (VTA).[46] Forward replay may contribute to several memory processes, such as memory consolidation and evaluation of future options,[37,40,47] both of which require prior encoding of a memory trace. Our observation that forward replay is more strongly related to the experience-dependent formation and modification of hippocampal spatial representations is consistent with this view. Further investigation is needed to decipher the roles of subsets of hippocampal replay to memory, for example, via their specific online manipulation during behavior.[48,49]

Changes in cellular contribution to replay events

Surprisingly, the increase in replay activity over trial blocks was accompanied by a progressive decrease in single-cell contributions to sharp-wave ripple events (i.e., putative replay events). This effect was mostly observed for bi-directional place cells that were active on both outbound and inbound trajectories. The contribution of cells that were active on journeys toward the large reward was elevated more strongly than the contribution of cells that were active in the small-reward environment, which could not be explained by changes in replay properties. Rather, we reason that bidirectional place cells, which may be active at different locations in the two running directions, lower the ability to accurately decode both run direction and position. Thus, a lower participation of bidirectional place cells together with a stable contribution of unidirectional place cells would result in the better estimation of position and direction and, consequently, in the improved detection of replay content.

Reverse replay relates more strongly to the most recent experience

To decode replay activity, we used a single-trial encoding model to control for the trial-to-trial changes in spatial representations instead of a trial-average model that is generally used. The encoding model we selected for decoding had an impact on the occurrence of reverse replay of events from the large-reward

environment. We found a higher number of identified events when using the run activity that occurred close to the time of the events occurrence, which suggests that reverse replay corresponds to the reactivation of CA1 activity patterns that relate to a single experience, more so than forward replay events.

Taken together, our results suggest that the reward-related enhancement of awake hippocampal replay activity is finely tuned toward behaviorally relevant experiences, in this case, the journeys leading to the large reward. The bias within replay events for representations of trajectories leading to the large reward is learning dependent and engages a selective refinement of the contribution of cell subpopulations.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- RESOURCE AVAILABILITY
  - Lead contact
  - Materials availability
  - Data and code availability
- EXPERIMENTAL MODEL AND SUBJECT DETAILS
- METHOD DETAILS
  - Behavioral procedure
  - Electrophysiological recordings
- QUANTIFICATION AND STATISTICAL ANALYSIS
  - Data Analysis
  - Behavior
  - Detection of sharp-wave ripple (SWR) events
  - Place cell analysis
  - Speed-corrected in-field firing rates
  - Population vector correlation
  - Emerging and vanishing place fields
  - Bayesian neural decoding
  - Replay analysis
  - Statistics
  - Curve fitting

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at https://doi.org/10.1016/j.cub.2021.07.058.

ACKNOWLEDGMENTS

F.K. is funded through Flemish Research Project FWO G0D7516N and KU Leuven C1 grant C14/17/042.

AUTHOR CONTRIBUTIONS

Conceptualization, F.M. and F.K.; formal analysis, F.M., E.K., and F.K.; investigation, F.M. and J.-J.S.; writing – original draft, F.M. and F.K.; writing – review & editing, F.M., J.-J.S., and F.K.; visualization, F.M. and F.K.; supervision, F.K.; funding acquisition, F.K.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: October 23, 2020
Revised: April 26, 2021
Accepted: July 23, 2021
Published: August 19, 2021

REFERENCES

1. Stickgold, R., and Walker, M.P. (2013). Sleep-dependent memory triage: evolving generalization through selective processing. Nat. Neurosci. *16*, 139–145.
2. Salvetti, B., Morris, R.G.M., and Wang, S.H. (2014). The role of rewarding and novel events in facilitating memory persistence in a separate spatial memory task. Learn. Mem. *21*, 61–72.
3. Gruber, M.J., Ritchey, M., Wang, S.F., Doss, M.K., and Ranganath, C. (2016). Post-learning hippocampal dynamics promote preferential retention of rewarding events. Neuron *89*, 1110–1120.
4. Miendlarzewska, E.A., Bavelier, D., and Schwartz, S. (2016). Influence of reward motivation on human declarative memory. Neurosci. Biobehav. Rev. *61*, 156–176.
5. Igloi, K., Gaggioni, G., Sterpenich, V., and Schwartz, S. (2015). A nap to recap or how reward regulates hippocampal-prefrontal memory networks during daytime sleep in humans. eLife *4*, e07903.
6. Michon, F., Sun, J.J., Kim, C.Y., Ciliberti, D., and Kloosterman, F. (2019). Post-learning hippocampal replay selectively reinforces spatial memory for highly rewarded locations. Curr. Biol. *29*, 1436–1444.e5.
7. Wolosin, S.M., Zeithamova, D., and Preston, A.R. (2012). Reward modulation of hippocampal subfield activation during successful associative encoding and retrieval. J. Cogn. Neurosci. *24*, 1532–1547.
8. Squire, L.R. (2004). Memory systems of the brain: a brief history and current perspective. Neurobiol. Learn. Mem. *82*, 171–177.
9. O'Keefe, J., and Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. Brain Res. *34*, 171–175.
10. Bostock, E., Muller, R.U., and Kubie, J.L. (1991). Experience-dependent modifications of hippocampal place cell firing. Hippocampus *1*, 193–205.
11. Navratilova, Z., Hoang, L.T., Schwindel, C.D., Tatsuno, M., and McNaughton, B.L. (2012). Experience-dependent firing rate remapping generates directional selectivity in hippocampal place cells. Front. Neural Circuits *6*, 6.
12. Leutgeb, S., Leutgeb, J.K., Treves, A., Moser, M.-B., and Moser, E.I. (2004). Distinct ensemble codes in hippocampal areas CA3 and CA1. Science *305*, 1295–1298.
13. Leutgeb, S., Leutgeb, J.K., Barnes, C.A., Moser, E.I., McNaughton, B.L., and Moser, M.-B. (2005). Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. Science *309*, 619–623.
14. Latuske, P., Kornienko, O., Kohler, L., and Allen, K. (2018). Hippocampal remapping and its entorhinal origin. Front. Behav. Neurosci. *11*, 253.
15. Dupret, D., O'Neill, J., Pleydell-Bouverie, B., and Csicsvari, J. (2010). The reorganization and reactivation of hippocampal maps predict spatial memory performance. Nat. Neurosci. *13*, 995–1002.
16. Danielson, N.B., Zaremba, J.D., Kaifosh, P., Bowler, J., Ladow, M., and Losonczy, A. (2016). Sublayer-specific coding dynamics during spatial navigation and learning in hippocampal area CA1. Neuron *91*, 652–665.
17. Tryon, V.L., Penner, M.R., Heide, S.W., King, H.O., Larkin, J., and Mizumori, S.J.Y. (2017). Hippocampal neural activity reflects the economy of choices during goal-directed navigation. Hippocampus *27*, 743–758.
18. Lee, H., Ghim, J.W., Kim, H., Lee, D., and Jung, M. (2012). Hippocampal neural correlates for values of experienced events. J. Neurosci. *32*, 15053–15065.
19. Lee, S.H., Huh, N., Lee, J.W., Ghim, J.W., Lee, I., and Jung, M.W. (2017). Neural signals related to outcome evaluation are stronger in CA1 than CA3. Front. Neural Circuits *11*, 40.
20. Gauthier, J.L., and Tank, D.W. (2018). A dedicated population for reward coding in the hippocampus. Neuron *99*, 179–193.e7.

21. Singer, A.C., and Frank, L.M. (2009). Rewarded outcomes enhance reactivation of experience in the hippocampus. Neuron *64*, 910–921.
22. Ambrose, R.E., Pfeiffer, B.E., and Foster, D.J. (2016). Reverse replay of hippocampal place cells is uniquely modulated by changing reward. Neuron *91*, 1124–1136.
23. Tabuchi, E., Mulder, A.B., and Wiener, S.I. (2003). Reward value invariant place responses and reward site associated activity in hippocampal neurons of behaving rats. Hippocampus *13*, 117–132.
24. Duvelle, É., Grieves, R.M., Hok, V., Poucet, B., Arleo, A., Jeffery, K.J., and Save, E. (2019). Insensitivity of place cells to the value of spatial goals in a two-choice flexible navigation task. J. Neurosci. *39*, 2522–2541.
25. Michon, F., Sun, J.-J., Kim, C.Y., and Kloosterman, F. (2020). A dual reward-place association task to study the preferential retention of relevant memories in rats. Front. Behav. Neurosci. *14*, 69.
26. McNaughton, B.L., Barnes, C.A., and O'Keefe, J. (1983). The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. Exp. Brain Res. *52*, 41–49.
27. Buzsáki, G. (2015). Hippocampal sharp wave-ripple: a cognitive biomarker for episodic memory and planning. Hippocampus *25*, 1073–1188.
28. O'Neill, J., Senior, T.J., Allen, K., Huxter, J.R., and Csicsvari, J. (2008). Reactivation of experience-dependent cell assembly patterns in the hippocampus. Nat. Neurosci. *11*, 209–215.
29. Davidson, T.J., Kloosterman, F., and Wilson, M.A. (2009). Hippocampal replay of extended experience. Neuron *63*, 497–507.
30. Karlsson, M.P., and Frank, L.M. (2009). Awake replay of remote experiences in the hippocampus. Nat. Neurosci. *12*, 913–918.
31. Kentros, C., Hargreaves, E., Hawkins, R.D., Kandel, E.R., Shapiro, M., and Muller, R.V. (1998). Abolition of long-term stability of new hippocampal place cell maps by NMDA receptor blockade. Science *280*, 2121–2126.
32. Anderson, M.I., and Jeffery, K.J. (2003). Heterogeneous modulation of place cell firing by changes in context. J. Neurosci. *23*, 8827–8835.
33. Lever, C., Wills, T., Cacucci, F., Burgess, N., and O'Keefe, J. (2002). Long-term plasticity in hippocampal place-cell representation of environmental geometry. Nature *416*, 90–94.
34. Frank, L.M., Stanley, G.B., and Brown, E.N. (2004). Hippocampal plasticity across multiple days of exposure to novel environments. J. Neurosci. *24*, 7681–7689.
35. Karlsson, M.P., and Frank, L.M. (2008). Network dynamics underlying the formation of sparse, informative representations in the hippocampus. J. Neurosci. *28*, 14271–14281.
36. Fyhn, M., Molden, S., Hollup, S., Moser, M.B., and Moser, E. (2002). Hippocampal neurons responding to first-time dislocation of a target object. Neuron *35*, 555–566.
37. Silva, D., Feng, T., and Foster, D.J. (2015). Trajectory events across hippocampal place cells require previous experience. Nat. Neurosci. *18*, 1772–1779.
38. Diba, K., and Buzsáki, G. (2007). Forward and reverse hippocampal place-cell sequences during ripples. Nat. Neurosci. *10*, 1241–1242.
39. Carr, M.F., Jadhav, S.P., and Frank, L.M. (2011). Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval. Nat. Neurosci. *14*, 147–153.
40. Joo, H.R., and Frank, L.M. (2018). The hippocampal sharp wave-ripple in memory retrieval for immediate use and consolidation. Nat. Rev. Neurosci. *19*, 744–757.
41. Ólafsdóttir, H.F., Bush, D., and Barry, C. (2018). The role of hippocampal replay in memory and planning. Curr. Biol. *28*, R37–R50.
42. Xu, H., Baracskay, P., O'Neill, J., and Csicsvari, J. (2019). Assembly responses of hippocampal CA1 place cells predict learned behavior in goal-directed spatial tasks on the radial eight-arm maze. Neuron *101*, 119–132.e4.
43. Foster, D.J., and Wilson, M.A. (2006). Reverse replay of behavioural sequences in hippocampal place cells during the awake state. Nature *440*, 680–683.
44. Ji, D., and Wilson, M.A. (2007). Coordinated memory replay in the visual cortex and hippocampus during sleep. Nat. Neurosci. *10*, 100–107.
45. Rothschild, G., Eban, E., and Frank, L.M. (2017). A cortical-hippocampal-cortical loop of information processing during memory consolidation. Nat. Neurosci. *20*, 251–259.
46. Gomperts, S.N., Kloosterman, F., and Wilson, M.A. (2015). VTA neurons coordinate with the hippocampal reactivation of spatial experience. eLife *4*, e05360.
47. Pfeiffer, B.E., and Foster, D.J. (2013). Hippocampal place-cell sequences depict future paths to remembered goals. Nature *497*, 74–79.
48. Deng, X., Liu, D.F., Karlsson, M.P., Frank, L.M., and Eden, U.T. (2016). Rapid classification of hippocampal replay content for real-time applications. J. Neurophysiol. *116*, 2221–2235.
49. Ciliberti, D., Michon, F., and Kloosterman, F. (2018). Real-time classification of experience-related ensemble spiking patterns for closed-loop applications. eLife *7*, e36275.
50. Millman, K.J., and Aivazis, M. (2011). Python for scientists and engineers. Comput. Sci. Eng. *13*, 9–12.
51. Harris, C.R., Millman, K.J., van der Walt, S.J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N.J., et al. (2020). Array programming with NumPy. Nature *585*, 357–362.
52. Virtanen, P., Gommers, R., Oliphant, T.E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., et al.; SciPy 1.0 Contributors (2020). SciPy 1.0: fundamental algorithms for scientific computing in Python. Nat. Methods *17*, 261–272.
53. McKinney, W. (2010). Data structures for statistical computing in Python. Proceeding of the 9th Python in Science Conference (SciPy), pp. 56–61.
54. Hunter, J.D. (2007). Matplotlib: a 2D graphics environment. Comput. Sci. Eng. *9*, 90–95.
55. Perez, F., and Granger, B.E. (2007). IPython: a system for interactive scientific computing. Comput. Sci. Eng. *9*, 21–29.
56. Kloosterman, F., Layton, S.P., Chen, Z., and Wilson, M.A. (2014). Bayesian decoding using unsorted spikes in the rat hippocampus. J. Neurophysiol. *111*, 217–227.
57. Sodkomkham, D., Ciliberti, D., Wilson, M.A., Fukui, K., Moriyama, K., Numao, M., and Kloosterman, F. (2016). Kernel density compression for real-time Bayesian encoding/decoding of unsorted hippocampal spikes. Knowl. Base. Syst. *94*, 1–12.
58. Kloosterman, F., Davidson, T.J., Gomperts, S.N., Layton, S.P., Hale, G., Nguyen, D.P., and Wilson, M.A. (2009). Micro-drive array for chronic in vivo recording: drive fabrication. J. Vis. Exp. 1094.
59. Nguyen, D.P., Layton, S.P., Hale, G., Gomperts, S.N., Davidson, T.J., Kloosterman, F., and Wilson, M.A. (2009). Micro-drive array for chronic in vivo recording: tetrode assembly. J. Vis. Exp. ●●●, 7–9.

Please cite this article in press as: Michon et al., Single-trial dynamics of hippocampal spatial representations are modulated by reward value, Current Biology (2021), https://doi.org/10.1016/j.cub.2021.07.058

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Deposited data		
Processed data and analysis code to recreate all main and supplemental figures.	Open Science Framework	https://osf.io/smzby
Experimental models: organisms/strains		
Rat: Long Evans	Janvier, France	Cat# RjOrl:LE
Software and algorithms		
Scientific Python	Millman and Aivazis,[50] Harris et al.,[51] Virtanen et al.,[52] McKinney,[53] Hunter,[54] and Perez and Granger[55]	https://numpy.org https://www.scipy.org https://pandas.pydata.org
Neural decoding without spike sorting	Kloosterman et al.[56] and Sodkomkham et al.[57]	N/A
Other		
Digilynx SX acquisition system with HS-36 analog headstage and Cheetah software	Neuralynx	https://www.neuralynx.com/
12 μm polyimide-insulated nickel-chrome tetrode wire	Sandvik	Cat# NI055820

RESOURCE AVAILABILITY

Lead contact

Further information and requests for resources should be directed to and will be fulfilled by the lead contact, Fabian Kloosterman (fabian.kloosterman@nerf.be).

Materials availability

This study did not generate new unique reagents.

Data and code availability

The data and analysis routines used in this study are available via the Open Science Framework: https://osf.io/smzby. This will enable the recreation of all main and supplemental figures. Additional code is also available at http://bitbucket.org/kloostermannerflab.

EXPERIMENTAL MODEL AND SUBJECT DETAILS

A total of 6 male Long-Evans rats were used for this study. All experiments were carried out following protocols approved by the KU Leuven animal ethics committee (P119/2015) in accordance with the European Council Directive (2016/63/EU). Animals with implanted electrodes were housed separately in individually ventilated cages (IVC) with *ad libitum* access to water and controlled intake of food pellets. Body weight and general health status were checked daily by the researchers and animal care personnel. The results presented here derive from the analysis of a subset of data collected for a previously published study.[6]

METHOD DETAILS

Behavioral procedure

All rats were trained in a dual reward-place association task.[25] Rats performed the task on an elevated maze that was split into two environments located at the right and left side of the experimental room. The two environments were separated by a divider and connected to a common home platform via 30 cm long tracks. Each environment consisted of a choice platform with connections for up

Please cite this article in press as: Michon et al., Single-trial dynamics of hippocampal spatial representations are modulated by reward value, Current Biology (2021), https://doi.org/10.1016/j.cub.2021.07.058

to 6 radially emanating arms. Each arm ended in a reward platform. The goal in the dual reward-place task is for the animal to learn and remember which of the 6 locations in each environment is associated with either large (9 pellets) or small (1 pellet) reward.

The dual reward-place association task is a repeated acquisition task in which rats need to learn and remember different associations every day. Each daily session is subdivided in an instruction phase and a test phase that were separated by a 2h delay. During the instruction phase, only the rewarded target arm is physically present in each environment. Across 5 trial blocks, rats were allowed to sequentially visit the target arm in the two environments to collect and consume reward. The rats were next placed in an enclosure located in the experimental room for the 2h delay. Following the delay, the rats were tested for their memory of the daily reward-location association separately for both environments in the presence of three additional distractor arms. In each session, task parameters were varied pseudo-randomly, i.e., the location of the target and distractor arms, the large/small reward assignment to left/right environment and the order in which environments were presented to the animal during instruction and test phase.

Electrophysiological recordings

A custom-designed 3D-printed micro-drive array,[58,59] carrying up to 24 tetrodes and 3 stimulation electrodes was surgically attached to the rat skull. During surgery, the array was positioned above the cortical surface through craniotomies located above the dorsal hippocampus for the recording electrodes (center coordinates: 4 mm posterior to Bregma, 2.5 mm right from the midline) and above the ventral hippocampal commissure for the stimulation electrodes (center coordinates: 1.3 mm posterior to Bregma, 0.9 mm right from the midline). Note that the stimulation electrodes were not used as part of this study. Following 1 week of post-operative recovery, the electrodes were lowered toward the pyramidal cell layer of hippocampal area CA1 over the course of 2-3 weeks.

During experiment, electrophysiological recordings were performed using a 128-channel data acquisition system (Digilynx SX, HS-36 analog headstage and Cheetah software; Neuralynx, Bozeman, MO). Wide-band (0.1-6000 Hz) signals and waveform snippets of online detected spikes in the band-pass filtered signal (600-6000 Hz) were sampled at 32 kHz. The position of the rats in the maze was tracked and captured at 50 Hz using an overhead video camera and colored LEDs mounted on the headstage.

QUANTIFICATION AND STATISTICAL ANALYSIS

Data Analysis

Analysis of neural and behavioral data was performed using Python and its scientific extension modules,[50] augmented with custom Python and C++ toolboxes. In particular, we used numpy[51] for numerical computations, scipy,[52] pandas,[53] statsmodels, scikit-learn and scikit-posthocs for statistical analyses, matplotlib[54] and seaborn for visualization, and Jupyter/IPython[55] for interactive analyses.

Behavior

The position of the rats was tracked using an overhead video camera. Running speed was computed as the magnitude of the Gaussian (bandwidth 0.5 s) smoothed gradient vector of position. In the instruction trials, the average running speed to (from) the reward platforms was computed over the full journey between leaving home (reward platform) and arriving at the reward platform (home).

Detection of sharp-wave ripple (SWR) events

The local field potentials from 1-3 tetrodes were downsampled from 32 kHz to 4 kHz and filtered in the ripple frequency band (140-225 Hz). The ripple envelope was computed as the absolute value of the Hilbert-transformed filtered ripple signal, averaged across the selected tetrodes and smoothed with a Gaussian kernel (bandwidth 15 ms). Slow trends in the ripple envelope were removed using a moving median filter (window length 3 s). Finally, start and end times of ripple events were detected when the detrended ripple envelope exceeded a low threshold of $\mu + 0.5 \times \sigma$ and the maximum envelope exceeded a high threshold of $\mu + 8 \times \sigma$. Here, μ and σ represent the mean and standard deviation of the detrended ripple envelope. Ripple events that were separated by less than 20 ms were merged into a single event, and events with a duration shorter than 40 ms were excluded.

Place cell analysis

Spike waveforms were automatically clustered into putative single units using Kilosort2 (https://github.com/MouseLand/Kilosort) and manually curated (using software package Phy, https://github.com/cortex-lab/phy) to identify well-isolated and stable single unit activity clusters. Units with a waveform peak-to-trough duration less than 0.5 ms or more than 1% contamination in a 2 ms refractory period were excluded from analysis.

For each unit, we computed the spatial tuning curve as the average firing rate in 1 cm bins along the trajectories from home to reward platform (outbound) and back (inbound). Only time windows in which the run speed exceeded 5 cm/s and no ripple events occurred were included to compute the tuning curves. Separate tuning curves were computed for the trajectories to/from the left and right environments. All tuning curves were smoothed with a Gaussian kernel (bandwidth 10 cm).

Based on the spatial tuning curves, place fields were defined as contiguous locations where the firing rate exceeded a low threshold of 0.1 Hz and the peak firing rate exceeded a threshold of 1 Hz. We used this relatively low threshold of 1 Hz to make sure to include fields that appear or disappear over the five instruction trials. The above procedure leads also to the inclusion of fields

Please cite this article in press as: Michon et al., Single-trial dynamics of hippocampal spatial representations are modulated by reward value, Current Biology (2021), https://doi.org/10.1016/j.cub.2021.07.058

with weakly modulated firing rates across all trials. Therefore, to exclude such fields, we set a minimum in-field firing rate of 5 Hz in at least one of the five trials. Finally, narrow place fields (less than 10 cm distance between the outer most spikes in the field) were excluded from the analysis.

In a total of 37 separate sessions (median 7.0 sessions/animal, range 2-8 sessions/animal), we recorded from 758 place cells (median 18.0 cells/session, range 6-49 cells/session). Place cells had a median of 2.0 place fields (percentage of clusters with 1 field: 45.8%, 2 fields: 28.2%, 3 or more fields: 26.0%). Since the home location in the maze overlapped between run trajectories to/from the reward locations, we excluded place fields in home from further analyses.

Speed-corrected in-field firing rates

To remove the relation between run speed and place cell firing rates, we computed a speed-corrected in-field firing rate per instruction trial. First, we model the rate-speed relation by performing kernel regression on all fields across all animals and sessions. Kernel regression is a non-parametric method that tries to discover a non-linear relation between two variables. We used a Gaussian kernel and automatic bandwidth selection based on least-squares cross-validation. In the model, we only included trials and fields with non-zero rates. Next, we computed the corrected in-field firing rate as:

$$r_{ij} = r_{ij} \times \frac{\bar{r}}{\hat{r}_{ij}}$$

where r_{ij} is the rate for field i in trial j, $\bar{r}$ is the mean rate across all fields and trials and $\hat{r}_{ij}$ is the speed-predicted rate for field i and trial j. We used a multiplicative correction to make sure that the corrected rates remained positive. After correction, no relation between speed and in-field firing rate remained (Figure S1).

Population vector correlation

For each instruction trial, a vector of speed-corrected in-field firing rates was constructed for all place fields or subset of place fields. Trial-to-trial similarity of the population activity was calculated as the Spearman correlation coefficient between the corresponding in-field rate vectors. The population vector correlation was computed at a per-session level and subsequently averaged across sessions.

Emerging and vanishing place fields

Emerging and vanishing place fields were defined as those fields with zero in-field firing rate in the first or last trial(s). To test if a sequence of zero-rate trials occurred by chance (assuming the occurrence of zero-rate trials is randomly distributed across all trials), the fractions of fields with n zero-rate trials on the trials $[t, \ldots, t+n-1]$ was computed for $n \in [1, \ldots, 4]$ and $t \in [1, \ldots, 5]$ with $t + n - 1 \leq 5$. The fractions were expressed relative to the expected fraction of fields with n zero-rate trials computed as the probability of drawing n trials out of 5. A value above 1 indicates that it is more likely than chance that a sequence of n zero-rate trials occur from trial t.

Bayesian neural decoding

Per recording session, an encoding model that relates hippocampal spiking activity to the animal's position was constructed from the data acquired in the instruction phase of the task. Only spikes emitted during run epochs (run speed > 10 cm/s) with a minimum spike amplitude of 60 μV were incorporated into the model. Tetrodes with a mean spiking rate during run epoch below 0.1 Hz were excluded.

We used a decoding approach that directly relates spike amplitude features to position without prior spike sorting.[56] Under the assumption that all spikes on a tetrode occur conditionally independent of past spikes and that the firing rate is determined by position in the maze, the hippocampal activity on a single tetrode can be modeled as a marked temporal Poisson process that is fully characterized by the rate function $\lambda(a,x)$, where a represents the vector of spike amplitudes and x represent position in the maze. The likelihood of observing a set of spikes with an amplitude of $a_{1:n}$ in time interval Δ for a given position is then expressed as:[56]

$$P(a_{1:n}|x) = \Delta^n \left[\prod_{i=1}^{n} \lambda(a_i, x) \right] e^{\Delta \lambda(x)} \qquad \text{(Equation 1)}$$

And the joint likelihood across K tetrodes is obtained by product of the single tetrode likelihoods:

$$P(a^{1:K}|x) = \prod_{k=1}^{K} P(a_{1:n_k}|x)$$

.

We used a compressed kernel density estimator to evaluate the rate function $\lambda(a,x)$ and the marginal rate function $\lambda(x)$ from their component spike count and position occupancy probability distributions.[57] The bandwidth of the Gaussian kernel was set to 30 μV for spike amplitude and 5 cm for position. Mahalanobis distance threshold for compression was set to 1.0 for an acceptable trade-off between decoding accuracy and computation time.

We used two different encoding models to decode spatial information in hippocampal replay events. The trial-average encoding model is a single model that incorporates the place field activity from all instruction trials. The single-trial model is composed of a

Please cite this article in press as: Michon et al., Single-trial dynamics of hippocampal spatial representations are modulated by reward value, Current Biology (2021), https://doi.org/10.1016/j.cub.2021.07.058

separate model for each of the trial blocks which incorporates only the place field activity within that trial block. For each candidate replay event in a trial block, decoding of the spatial information is performed using the corresponding trial block encoding model.

To perform neural decoding for the spiking activity recorded on K tetrodes in time window Δ and estimate the posterior probability distribution over position (sampled at a regular grid with 4 cm spacing) from the likelihood we resort to Bayes' rule:

$$P(x|a^{1:K}) = \frac{P(a^{1:K}|x)P(x)}{P(a^{1:K})} \quad \text{(Equation 2)}$$

where a uniform prior $P(x)$ is used.

To evaluate the performance of the decoder, we used a five-fold cross-validation procedure in which four out of five instruction trial blocks were used to build the encoding model and the remaining instruction trial block was used for decoding the animal's position on the maze in Δ = 100 ms time bins. The decoding error was defined as the distance along the track between the estimated and real position. For each session, we evaluated the decoding error distribution separately for left and right environments, as well as for the complete maze. Cross-validation was only performed with the trial-average encoding model. Only sessions with a good decoding performance during the instruction phase (75 percentile of decoding error distribution is below 30 cm) were selected for subsequent analysis of hippocampal replay (29 sessions from 6 animals).

Replay analysis

A smoothed multi-unit activity (mua) rate histogram (5 ms bin size, Gaussian kernel with 15 ms bandwidth) was computed from all unsorted spikes recorded from the hippocampus with a peak amplitude larger than 60 μA. The rate was detrended using a moving median filter (window length 3 s). Transient bursts in the detrended mua were defined using a double threshold procedure, where the upper threshold $\mu + 4\times\sigma$ determines if a burst occurred and the lower threshold $\mu + 0.5\times\sigma$ determines burst start and end time. Bursts that were separated by less than 20 ms were merged, and bursts with a duration shorter than 80 ms were excluded. Candidate replay events were defined as mua bursts that overlapped with a ripple event and which occurred while the animal was immobile (run speed < 5 cm/s).

Candidate replay events were split into Δ = 10 ms time bins and spiking activity (amplitude > 60 μV) in each bin was used to perform decoding as described above. Next, separately for left and right target arms, weighted isotonic regression was performed on the maximum-a-posteriori (MAP) position estimates with posterior probabilities as weights.[6] A goodness-of-fit score was defined that combined both the posterior probabilities and the R^2 of the regression: $score = \frac{1}{T}\sum_{t=1}^{T} P_{MAP,t}\times R^2$, where T is the number of time bins in the event, and $P_{MAP,t}$ is the posterior probability associated with the MAP estimate in time bin t. The regression was performed twice to fit both a monotonically increasing and a monotonically decreasing trajectory to the MAP estimates, and only the best fitting trajectory with the highest score was retained.

For each event, the goodness-of-fit score is compared to the distribution of scores constructed from 500 pseudo-random events in which each posterior was randomly drawn from the complete set of candidate replay events.[29] Only candidate replay events with a Monte-Carlo p value < 0.05 were considered to contain significant trajectory replay.

For each significant replay trajectory, the run direction was decoded using the same Bayesian neural decoding approach as used for position. For each replay event, a direction bias was computed separately for the small and large reward environment as the mean difference in posterior probability between the inbound and outbound run directions across time bins. For each event, the direction bias was compared to the distributions of biases computed from 500 pseudo-random events obtained by randomly assigning each posterior probability of that same event to either the large or small reward environment. Only candidate replay events with a Monte-Carlo p value < 0.05 were considered to be significantly biased for one run direction. Replay events were classified as either 'forward' or 'reverse', depending on whether the decoded direction matched the direction of the replay trajectory.

Statistics

To compare the distributions of trial-to-trial firing rate changes, we used the two-sample Kolmogorov-Smirnov test for equal distributions with Holm-Sidak p value correction.

To compare per-session correlation coefficients between trial blocks, we first used the Friedman chi-square omnibus test, followed by a Conover posthoc test of all pairwise combinations that uses Holm-Sidak p value correction method.

To compare the number of replay events between successive pairs of trial blocks and between reward conditions, we used the Wilcoxon signed-rank test with Holm-Sidak p value correction to account for the multiple tests.

To compare the contribution of place cells to SWRs and replay events between pairs of successive trials blocks and between rewards conditions, we used the Mann-Whitney U test followed by a Holm-Sidak p value correction to account for the multiple tests.

To test the distribution of trial blocks at which individual place fields appear or disappear, we performed a chi-square test for the null hypothesis of a uniform distribution.

We use p* to indicate p values that were corrected for multiple tests.

Please cite this article in press as: Michon et al., Single-trial dynamics of hippocampal spatial representations are modulated by reward value, Current Biology (2021), https://doi.org/10.1016/j.cub.2021.07.058

Curve fitting

The relation between trial-to-trial correlation, rate or $|\Delta_{rate}|$ and the sequence number of the trial blocks (or trial block pairs) was fitted with a sigmoid growth function:

$$y(k) = \beta_1 + \frac{\beta_2}{1 + e^{-\lambda k}}$$

where, β_1, β_2 and λ are the parameters to be fitted, and $k \in [0, 1, ..., n-1]$ is the sequence number for the trial blocks or trial block pairs (i.e., $k = 0$ for $r_{1\rightarrow2}$, $k = 1$ for $r_{2\rightarrow3}$, etc.). Fitting was performed using the mean squared error and L2 regularization as the cost function. We performed global minimization of the cost function using the stochastic differential evolution approach. Ten-fold cross-validation was used to estimate the hyperparameter for the L2 regularization. The 95% confidence intervals of the fitted parameters were computed by bootstrapping (500 samples).

The curve fit was summarized by two derived parameters that express the total change from the first to the last trial block or trial block pairs $(\Delta = y(n-1) - y(0))$, and the relative change between the first two trial blocks or trial block pairs $\left(rel\Delta = \frac{y(1)-y(0)}{y(2)-y(0)} = \frac{1}{\mathrm{sech}(\lambda)+1}\right)$.

To compare the derived parameters between small and large reward conditions, we computed the bootstrap distribution of differences and calculated the two-sided p value for this distribution to be different from zero.