PNAS RESEARCH ARTICLE | NEUROSCIENCE
PSYCHOLOGICAL AND COGNITIVE SCIENCES

Sensory integration, temporal prediction, and rule discovery reflect interdependent inference processes

Lucas Benjamin[a,b,c,1], Benjamin Morillon[a,2], and Valentin Wyart[b,c,1,2]

Affiliations are included on p. 10.

Edited by Peter Dayan, Max Planck Institute for Biological Cybernetics, Tübingen, Germany; received September 11, 2025; accepted December 17, 2025 by Editorial Board Member Peter L. Strick

Deciphering the structure of variable sensory input is key to building an accurate model of one's environment. Humans can accumulate evidence from sequences of stimuli to estimate their sensory statistics, predict the timing of upcoming stimuli, but also discover rules governing sequence generation. However, whether these three forms of inference operate independently or synergistically remains untested. Here, we report selective interactions between sensory integration, temporal prediction, and rule discovery in humans. Participants were exposed to rhythmic sequences of 10 stimuli governed or not by a latent rule—a predictable change in stimulus statistics after five stimuli—and then asked to predict the 10th stimulus from incomplete sequences. Individual differences in sensory integration timescale for rule-free sequences predicted efficient rule discovery. Conversely, discovering the latent rule shaped the timescale and format of sensory integration for rule-based sequences. Tampering with the rhythmicity of stimulus presentation impaired rule discovery without affecting sensory integration accuracy. Selective perturbations of recurrent neural networks trained in the same conditions confirmed these specific interactions. Together, these findings provide insights into the flexibility of human inferences based on variable yet predictable sensory input.

human learning | rule learning | inference | cognitive modeling | artificial neural networks

Significance

Humans routinely extract three types of information from stimulus sequences: their statistics, timing, and the rules that govern their generation. But do these three cognitive abilities interact? By examining how hundreds of individuals predict future stimuli in sequences governed by a hidden rule, we uncovered interactions between sensory integration, temporal prediction, and rule discovery. Individuals with longer integration timescales were more likely to discover the rule—and more so for rhythmic sequences. Moreover, individuals who discovered the rule integrated stimuli more accurately. For artificial neural networks, as for humans, shorter integration timescales and arrhythmic sequences hindered rule discovery. Therefore, sensory integration, temporal prediction, and rule discovery reflect interacting cognitive abilities, whose interplay should be reconsidered by existing theories.

The sensory environments we live in are variable, yet highly organized. Understanding the generative organization that governs their structure is crucial for reducing uncertainty. To do so, humans rely on multiple inference abilities, including sensory integration (1–4), temporal prediction (5–8), categorization based on arbitrary mappings of sensory features (9–11), learning statistical dependencies between stimuli (12–18), and the learning of hidden rules describing the patterns of relations between successive stimuli (19–23). In a majority of studies, these different learning abilities are assumed, either implicitly or explicitly, to rely on different cognitive and neural mechanisms. But some studies, on the opposite, suggest a single unified cognitive apparatus for those different learning processes (24). This debate lacks empirical data, and these abilities themselves have rarely been studied conjointly, resulting in a fragmented description of human learning abilities. Here, we studied three human inference capacities conjointly and empirically tested if they rely on similar or distinct and independent or interrelated cognitive processes: sensory integration, temporal prediction, and rule discovery.

So far, inference capacities have mostly been studied in isolation, targeting different research questions and using different methods. In perceptual decisions research, studies have examined how sensory information is integrated across sequences of stimuli—or across samples from the same stimulus—to reach a decision (3, 4, 25, 26). We refer to this ability as *sensory integration*. This process is known to be prone to a number of suboptimalities constraining the accuracy of perceptual decisions across individuals (27–30). Yet, it remains largely unknown whether these sensory integration suboptimalities are confined to sensory integration itself, or whether they act as a broader cognitive bottleneck impairing higher-order inferences, such as rule discovery.

To investigate this possible effect of sensory integration limitations on higher-order inferences, we designed a rule-based visual prediction task that requires both sensory integration and the discovery of a hidden rule for accurate predictions. Rule discovery is known to depend on the rule's algorithmic complexity and to vary across individuals (22, 31). However, except for a few specific studies (32), sequences used in most rule discovery tasks are typically stripped of any variability, eliminating the need to integrate sensory information across stimuli. In our rule-based prediction task, we used rhythmic isochronous sequences of 10 oriented stimuli (Fig. 1*A*) sampled from probability distributions whose mean orientations switch (i.e., rotate by 90°) after five stimuli—a latent rule that

Author contributions: L.B., B.M., and V.W. designed research; L.B. and V.W. performed research; L.B. and V.W. contributed new reagents/analytic tools; L.B., B.M., and V.W. analyzed data; and L.B., B.M., and V.W. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission. P.D. is a guest editor invited by the Editorial Board.

[1]To whom correspondence may be addressed. Email: lucas.benjamin78@gmail.com or valentin.wyart@inserm.fr.

[2]B.M. and V.W. contributed equally to this work.

This article contains supporting information online at https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2524629123/-/DCSupplemental.

Published February 6, 2026.

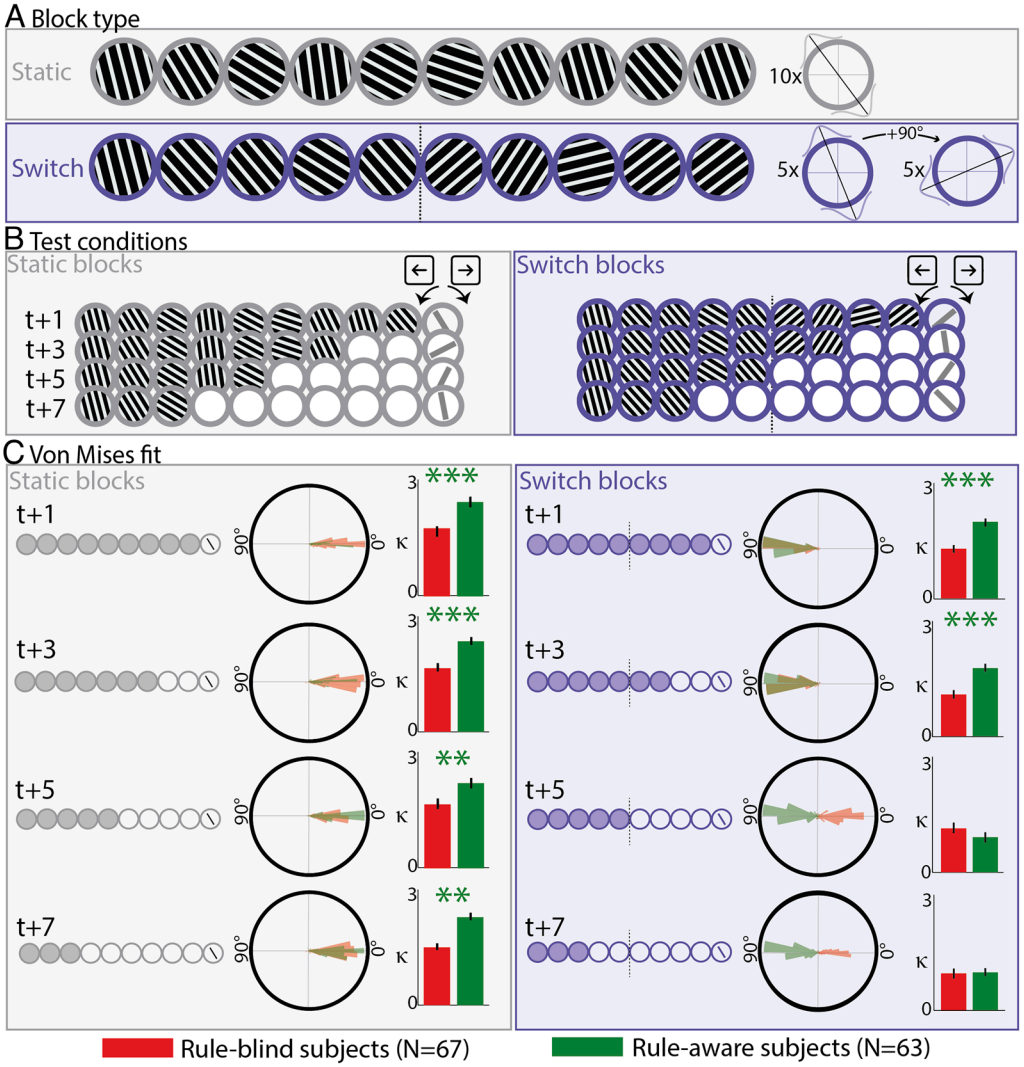

Fig. 1. (*A*) Description of the two types of blocks from Experiment 1. Each of the grating patches lasts 250 ms. In static blocks, sequences were composed of 10 items drawn from a normal distribution centered on a target orientation randomly picked for each sequence. For switch blocks, the last five stimuli of sequences were drawn from the 90° shifted distribution compared to the first five stimuli. (*B*) For each block, after the presentation of six complete sequences, participants performed different conditions, i.e., incomplete sequences of 3/5/7 or 9 first stimuli. For each, they had to predict what would have been the orientation of the last (10th) item of this particular sequence by rotating a bar. (*C*) We compared participants' predictions to the mean orientation before switch and fitted Von Mises distributions for each participant, condition, and block. Polar plots display the histogram of the location (θ) of the Von Mises distribution for each participant, while bar plots display participants' average precision (κ). Error bar represents the SE of κ values across subjects **$p_s < 0.01$ *** $p_s < 0.001$ (Bonferroni corrected across conditions).

was not described to participants. After exposing participants to full sequences, we asked them to predict the 10th stimulus from incomplete sequences of 3, 5, 7, or 9 stimuli (*switch* blocks, Fig. 1*B*). In other blocks (*static* blocks), we used sequences where all stimuli were drawn from a single distribution—without a latent rule—to investigate if suboptimalities in sensory integration measured in static blocks could impact discovery of the latent rule in switch blocks. By introducing interstimulus temporal jitter, we further investigated whether a third inference ability—temporal prediction—affected either of these two forms of inference, or their interrelation.

Our findings, obtained from nearly 600 participants across five experiments, uncovered a clear bidirectional interaction between sensory integration and rule discovery. We found that participants' timescale of sensory integration during static blocks predicted whether they would discover the latent hidden rule in switch blocks. Conversely, we found that discovering the latent rule led rule-aware participants to accurately mentally rotate presented stimuli in switch blocks before integrating them. Additional experiments controlled for general attention and working memory to confirm that the timescale of sensory integration has a unique effect on latent rule discovery. Moreover, we showed that temporal unpredictability selectively impairs rule discovery, without altering sensory integration. Finally, we used targeted manipulations of recurrent neural networks (33–35) to establish a causal link between sensory integration timescale and latent rule discovery.

Results

Experiment 1: Sensory Integration Predicts Rule Discovery. We first investigated if two key human inference processes—sensory integration and rule discovery—are interacting or independent by testing if rule discovery in switch sequences could be predicted by subject performances during static sequences. First, we identified behavioral signatures diagnostic of latent rule discovery, to classify participants accordingly. Separately for each type of block (static/switch) and each type of prediction (t+1/t+3/t+5/t+7), we fitted the distribution of tilt between each participant's response (their

predicted orientation for the 10th stimulus) and the mean orientation before the switch (*Materials and Methods*). The fitted location (θ) provided a clear signature of rule discovery (Fig. 1*C*). Indeed, for conditions $t + 5$ and $t + 7$ in switch blocks, the orientation of the 10th stimulus is shifted by an average of 90° relative to the mean of the preswitch distribution. A participant who discovered the hidden rule should have a location θ close to 90°, whereas a participant who did not should have a location θ close to zero, as they have only seen the first half of the sequence in these conditions. Fitted location for these conditions in switch blocks revealed a clear bimodal distribution around these two behaviors (Fig. 1*C*). We split participants into two groups based on their fitted location in these conditions: *rule-aware* for participants with $\theta > 45°$ (N = 63) and *rule-blind* for participants with $\theta < 45°$ (N = 67, Fig. 1*C*). The order of block presentation in the experiment had no impact on the discovery of the rule (71 participants started with a static block: 48% rule-aware; 59 started with a switch block: 49% rule-aware).

We compared these two groups in terms of other aspects of behavior. The first step was to compare their response precision κ (*Materials and Methods*). Rule-aware participants made significantly more precise predictions than rule-blind participants in static blocks (concentration κ: $p_s < 0.01$, Bonferroni corrected; Fig. 1*C*). They also made more precise $t + 1$ and $t + 3$ predictions in *switch* blocks ($p_s < 0.001$). This improvement in prediction precision in rule-aware participants was accompanied by a higher self-reported confidence than rule-blind participants (*SI Appendix*, Fig. S2).

To characterize the nature of the increased precision of rule-aware participants, we used a computational modeling approach (28, 29). We decomposed each participant's imprecisions during static blocks into two sources of internal errors: integration errors, arising from suboptimal perception and integration of successive stimuli, and response errors, reflecting suboptimal reporting of the participant's prediction (*Model 1*; *SI Appendix*, Fig. S1). Integration errors, but not response errors, were significantly smaller for rule-aware than rule-blind participants (rank-sum test, $P < 0.001$, response error $P = 0.35$; *SI Appendix*, Fig. S1*B*).

Experiment 1: Sensory Integration Timescale Predicts Rule Discovery. To better understand the specific sensory integration parameters which affect rule discovery in switch blocks, we derived a fine-grained model of participants' behavior in static blocks (*Model 2*). It comprised five sources of suboptimality (Fig. 2*A*): 1/sensory bias (γ_{sen}) distorting stimulus orientation toward diagonals (28, 36), 2/random noise of SD σ_{int} corrupting the integration process, 3/integration leak (α), capturing the timescale of sensory integration with a stronger leak corresponding to a shorter integration timescale, 4/random noise of SD σ_{rep} corrupting the reporting of participants' belief, and 5/response bias (γ_{rep}) distorting the reported orientation toward diagonals. Parameter recovery confirmed that all parameters were dissociable and interpretable (*SI Appendix*, Fig. S3).

Fitting this model to participants' behavior in static blocks revealed a selective difference in sensory integration leak between rule-aware and rule-blind participants: Rule-aware participants showed lower integration leak (i.e., longer integration timescale–α = 0.33 ± 0.05) than rule-blind participants (0.80 ± 0.08; rank-sum test $P < 0.001$ Bonferroni corrected). No other parameter significantly differed between the two groups of participants.

Experiments 2-3: Sensory Integration Timescale and Rule Discovery Are Unrelated to Working Memory. To test whether individual differences in working memory (WM) capacity (or more generally, in executive attention) were related to individual differences in sensory integration timescale, we carried out a first control experiment. In **Experiment 2,** a new group of participants

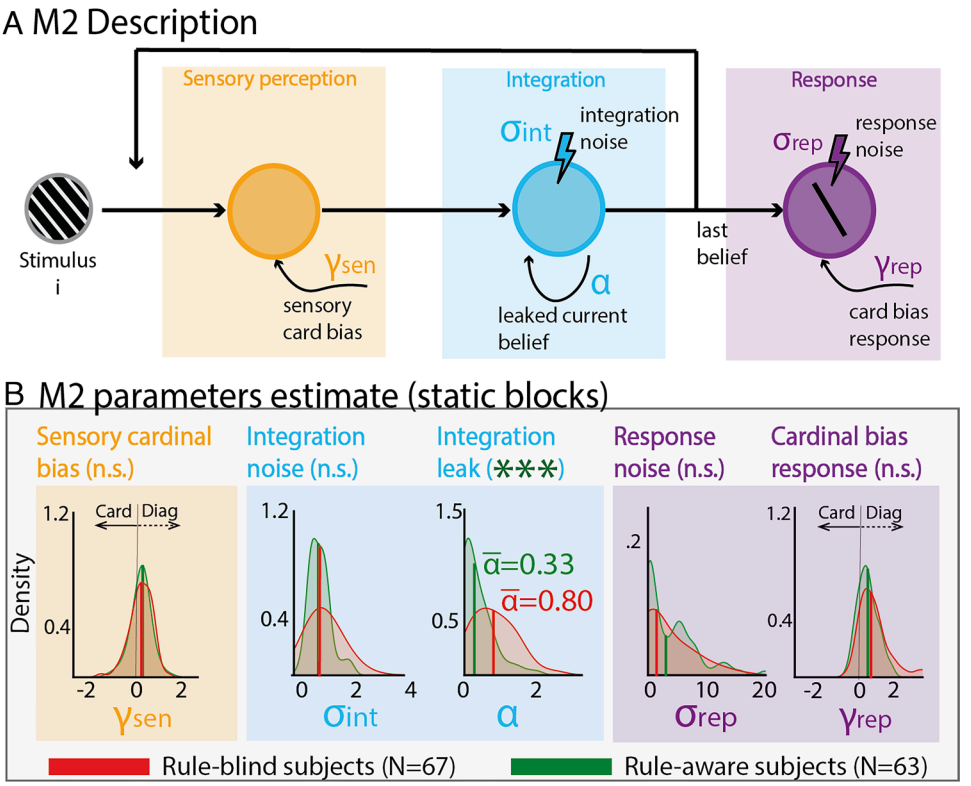

Fig. 2. (*A*) Description of Model 2 with five suboptimalities parameters (*Materials and Methods*). (*B*) Parameter fit on static blocks and group comparison between rule-aware and rule-blind participants. ***$p_s < 0.001$ (rank-sum test, Bonferroni corrected over the five parameters).

completed static blocks (without a hidden rule) alongside an n-back working memory task using the same stimuli—by asking participants to report the orientation of a single stimulus in the sequence cued at sequence offset (*Materials and Methods*). Participants' performance in the sensory integration task and in the working memory task diverged as sequence length increased: Accuracy in the n-back working memory task declined, whereas accuracy in the sensory integration task improved (ANOVA seqlen x blocktype, $P < 0.001$). This performance dissociation suggests that performance in the sensory integration task does not rely strongly on working memory capacity. In line with this dissociation, individual differences in performance in the working memory task did not correlate with individual differences in the sensory integration timescale, as estimated by the integration leak α from model M2 fitted to behavior in the sensory integration task (*SI Appendix*, Fig. S4).

To further test whether individual differences in general attention or working memory capacity were related to individual differences in rule discovery, we carried out a second control experiment. In **Experiment 3,** an additional group of participants performed *switch* blocks (with the hidden rule) together with two working memory assessments: a first assessment targeting WM capacity and a second assessment targeting WM maintenance. Although participants' performance in the two WM assessments correlated with each other (Spearman correlation R = 0.31; $P < 0.01$), neither distinguished rule-aware from rule-blind participants in the switch blocks (rank-sum tests both $ps > 0.24$, *SI Appendix*, Fig. S4). Together, these results indicate that individual differences in working memory performance do not explain the observed effect of sensory integration timescale (reflected in the integration leak) on rule discovery.

Experiment 1: Rule Discovery Alters Sensory Integration. Analysis of the static blocks provided clear evidence for a "bottom–up" interaction between sensory integration timescale (as measured by participants' leak) and rule discovery. To explore a possible "top–down" interaction of rule discovery on human integration strategy, we turned to the analysis of participants' behavior in switch blocks from **Experiment 1**. We compared two hypotheses: Either participants that discovered the rule do not change their integration process but only add a correction offset to their response to account for the rule (Model 3), or they flexibly apply mental rotation to each of the first 5 stimuli, prior to sensory integration (Model 4).

We fitted those two models (Fig. 3*A*) on switch data together with the baseline model without any offset (M2; Fig. 2*A*). Model comparison revealed the integration offset model (M4) to be the best fit for rule-aware participants (AIC: M2 = 671, M3 = 647, M4 = 523; $P < 0.001$, Fig. 3*A*). We checked that the offset rule implemented in this model (same offset for the first five stimuli and no offset for the last five) was correct by fitting a complementary model

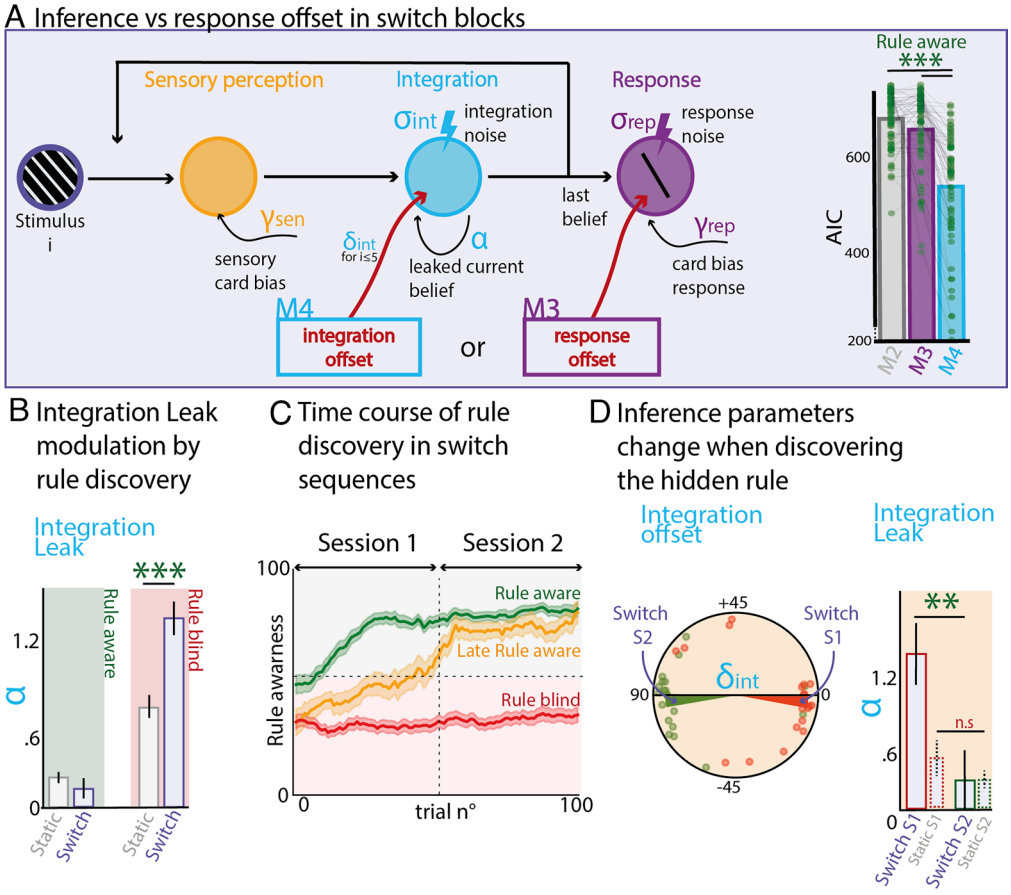

Fig. 3. (*A*, *Left*) Description of Models 3 and 4. Model 3 is similar to Model 2 (Fig. 2) with an additional offset parameter applied to the response while Model 4 has instead an offset parameter applied during the integration of the five first stimuli of the sequence. (*Right*) AIC comparison when fitted on rule-aware participants' responses to switch blocks for Models 2 ("no offset"), 3 ("response offset"), and 4 ("integration offset"). (*B*) Comparison of the integration leak within participants, between static and switch blocks, for rule-aware and rule-blind participants. ***$p_s < 0.001$ (rank-sum test). (*C*) Time-course of the rule discovery for each group of subjects. The y-axis represents the probability that a given trial was closer to the rule-aware answer compared to the rule-blind answer. (*D*) Analysis of the 20 late rule-aware participants that were rule-blind in session 1 but rule aware in session 2. Discovering the rule in session 2 modified their integration by adding a 90° offset but most importantly drastically reducing the leak, further evidencing the causal role of rule discovery on integration timescale.

with nine offsets, one for each stimulus (Model 5). This confirmed our hypothesis, showing rule-aware participants had an exact representation of the switch rule and applied a 90° offset to the first five stimuli and no offset to the last five (*SI Appendix*, Fig. S5). Discovering the rule modified participant's sensory integration process, as they integrate mentally rotated information, rather than simply accounting for the rule in their reporting process.

To further support this idea, we explored one of its predictable implications. Participants who did not discover the rule should perceive switch blocks as more volatile than static blocks, while participants who found the rule can use this knowledge to explain away perceived volatility. As perceived volatility is typically accompanied by a smaller integration timescale (37, 38), we compared the leak during the static and the switch blocks for both groups. We found an increased leak in switch compared to static blocks for rule-blind participants (α switch = 1.5 ± 0.13, static = 0.79 ± 0.08; rank-sum test $P < 0.001$), up to the point where they almost ignored preswitch stimuli, while rule-aware participants kept a similar leak, revealing the use of the rule knowledge to reduce perceived volatility of switch sequences (α switch = 0.31 ± 0.06, static = 0.32 ± 0.05; $P = 0.6$; Fig. 3*B*, interaction $P < 0.001$). Integration noise was also modulated by rule discovery, in a way that suggests that the mental rotation applied to stimuli in switch blocks by rule-aware participants is associated with an additional source of random errors (*SI Appendix*, Fig. S6).

To investigate whether discovering the rule causally influences changes in integration parameters (rotation and leak), we performed a quasi-causal analysis by focusing on a specific subgroup of participants, labeled as *late rule-aware* ($N = 20$): classified as rule-blind in the first session (on day 1) but rule-aware in the second session (on day 2). If rule discovery genuinely modifies sensory integration parameters, we expected to observe changes in parameter values in the switch block of the second session compared to the first session. Trivially, rule-aware participants had a greater integration offset but more importantly a substantial reduction of integration leak in the second session compared to the first session, only during *switch* blocks ($P < 0.01$–Fig. 3*C*).

Experiments 4-5: Temporal Prediction Selectively Affects Rule Discovery. Sensory integration and rule discovery are not the only two forms of inference humans are known to be capable of performing. Indeed, presented with stimulus sequences, humans cannot only predict *what* they expect to occur but also *when* they expect a given element to appear. To investigate if temporal prediction affects the two others, we designed a jittered version of Experiment 1 and measured the influence of temporal unpredictability on both sensory integration and rule discovery. We reasoned that if sensory integration and rule discovery reflect distinct forms of inference—rather than two aspects of a single core cognitive ability (24)— temporal unpredictability could selectively disrupt only one of these forms of inference (**Experiment 4**, see *Materials and Methods*).

We first replicated all main effects of Experiment 1 (*SI Appendix*, Fig. S6). Then, we tested the influence of temporal jitter on sensory integration by fitting Model 2 to the data from static blocks with temporally jittered sequences (Experiment 4). We found no difference in the fitted parameters compared to those from static blocks with rhythmic sequences (Experiment 1) (rank-sum test all *ps* > 0.05; Fig. 4*A*). We confirmed this absence of difference in sensory integration parameters between rhythmic and temporally jittered sequences using a within-subject design (**Experiment 5**), where the same participants performed static blocks with and without temporal stochasticity (*SI Appendix*, Fig. S7).

Despite these unchanged integration parameters, temporal stochasticity significantly reduced participants' likelihood of discovering the rule during switch blocks (–39%, $P = 0.01$, bootstrap resampling; Fig. 4*B*), revealing a specific influence of temporal prediction on rule discovery. We found that this decreased probability of discovering the latent rule with temporally jittered sequences (in switch blocks) was independent of individual differences in integration leak (in static blocks) by fitting a logistic function that relates integration leak to the probability of rule discovery in each experiment. The intercept term, but not the slope, was significantly higher for rhythmic sequences (Experiment 1) compared to temporally jittered sequences (Experiment 4, bootstrap resampling $P < 0.05$), indicating that the effect of temporal stochasticity on rule discovery is not mediated by changes in integration leak (Fig. 4*B*; see *SI Appendix*, Fig. S8 for additional control analyses).

RNN Simulations: Shorter Sensory Integration Timescales Prevent Rule Discovery. Previous analyses showed that participants' integration timescale (measured by integration leak) was specifically associated with their capacity to discover the hidden rule in switch blocks. Detailed analyses of late rule-aware subjects (Exp 1) and jittered experiments (Exp 4 & 5) gave insights regarding the causal nature of the effect of sensory integration timescale on rule discovery. To further strengthen these findings, we turned to recurrent neural networks (RNNs) to test directly for a causal effect of sensory integration timescale on latent rule discovery.

In practice, we trained RNNs to predict each upcoming stimulus in the sequence. This approach allowed us to measure the network's surprise at each stimulus in the sequence (*Materials and Methods*). Based on these surprise estimates, we defined a *Rule Learning Index (RLI)* as the network's "surprise" to the switch (stimulus 6), relative to the average surprise across all other stimuli. This index, unavailable from human behavior, captures the latent dynamics of the ability of the network to correctly predict the switch rule and thus provides a quantitative metric of rule discovery in RNNs. In particular, it allows for precise dissection of the learning of sensory integration and rule discovery. Throughout the training, we observed a rapid decline in surprise for all stimuli except the switch, corresponding to the learning of pure sensory integration. This was followed by a later decrease of the surprise evoked by the 6th stimulus, indicating rule discovery (Fig. 5*B*). This dynamic is revealed in a sharp initial rise in the RLI (sensory integration learning), followed by a slower decrease (rule learning) until convergence.

To causally explore the role of each suboptimality, we sought to selectively modulate each of them in our networks, and observe their effect on RLI. For that, we aimed at introducing five selective perturbations that trigger the five suboptimalities observed in human behavior (Fig. 2). We simulated the behavior of selectively perturbed RNNs and fitted Model 2 to the RNN behavior during static blocks—similarly to what was done in humans to estimate their suboptimality parameters. We showed that distorting the network input with cardinal bias triggered a sensory cardinal suboptimality in RNN behavior. Similarly, distorting RNN response with cardinal bias or random noise introduced respectively cardinal bias response and response noise in their behavior. Most importantly, adding noisy recurrent connections triggered integration noise (σ_{int}) and sparsifying the hidden layer specifically triggered leak (α) in RNN behavior (Fig. 5*C*).

We used the *RLI* to assess how each suboptimality causally affects the learning trajectory in switch sequences and showed that integration noise and integration leak modified RLI

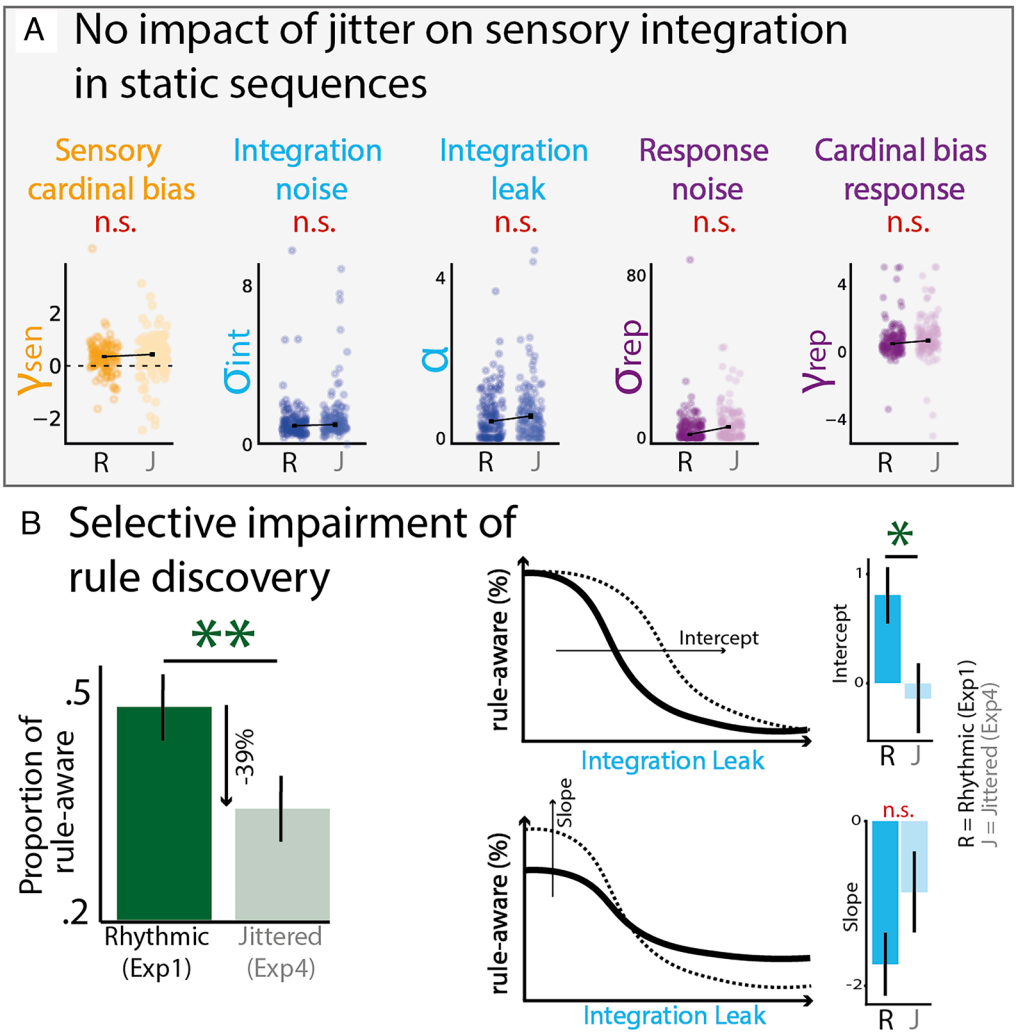

Fig. 4. (*A*) Comparing integration parameters during static sequence for rhythmic (Experiment 1) and jittered (Experiment 4) presentation revealed no impact of temporal stochasticity on integration. (*B*) However, jitter during the presentation significantly reduced the percentage of participants that found the hidden rule in switch sequences ($P = 0.01$). Fitting sigmoid between leak parameters during static sequences and rule-awareness probability gave a different intercept ($P < 0.05$) but a similar slope, revealing independent impact of jitter and leak on rule discovery.

dynamics. More precisely, while adding integration noise (σ_{int}) via noisier recurrent connections globally impaired both sensory integration and rule discovery, increasing integration leak (α) specifically delayed discovery of the switch rule without altering sensory integration (Fig. 5*C*). We considered learning timecourse as a complementary metric to closely match humans' analyses (Fig. 3*C*) and replicated our main effects (*SI Appendix*, Fig. S11). We also confirmed that this relationship was not limited to the particular case of a single switch in the middle of the sequence by replicating this analysis using another rule with sequences of the form AABBAABBAA (*SI Appendix*, Fig. S9).

To support the selective interaction between temporal predictability and rule discovery observed in Experiments 4-5, we trained a new set of RNNs to make temporally resolved predictions and trained them with rhythmic or temporally jittered stimulus sequences, using the exact same parameters used in Experiment 4. A Rule Learning Index (RLI) adapted to this temporally resolved condition showed that, like humans, the training of RNNs to perform sensory integration was unaffected by temporal stochasticity, whereas the latency at which the same RNNs discovered of the latent rule was strongly delayed by temporal stochasticity (Fig. 5*D*, see details in *SI Appendix*, Fig. S10).

Discussion

To make sense of sequences of stimuli, humans rely on multiple inferences often studied independently, leaving unknown how they interact to form a comprehensive learning system. A few recent studies have begun to address this issue, by investigating the impact of temporal prediction on statistical learning (39), or by asking human participants to classify sequences of stimuli drawn either from deterministic rules or from probabilistic distributions (40). The present study aims at understanding how sensory integration, temporal prediction, and rule discovery interact to form a flexible inference system that is able to make sense of sensory events unfolding in time. Our results revealed a strong two-way interplay between sensory integration and rule discovery: 1) a bottleneck exerted by the integration timescale (leak) of sensory integration on the probability of further discovering the hidden rule, and 2) the application of the hidden rule to incoming stimuli prior to their integration, as well as the use of the rule knowledge to adapt integration timescale. Our findings further reveal a selective impact of temporal prediction on rule discovery, leaving sensory integration unaffected.

First, Experiment 1 showed that participants' integration timescale [as measured by integration leak (28, 30, 41, 42)], measured

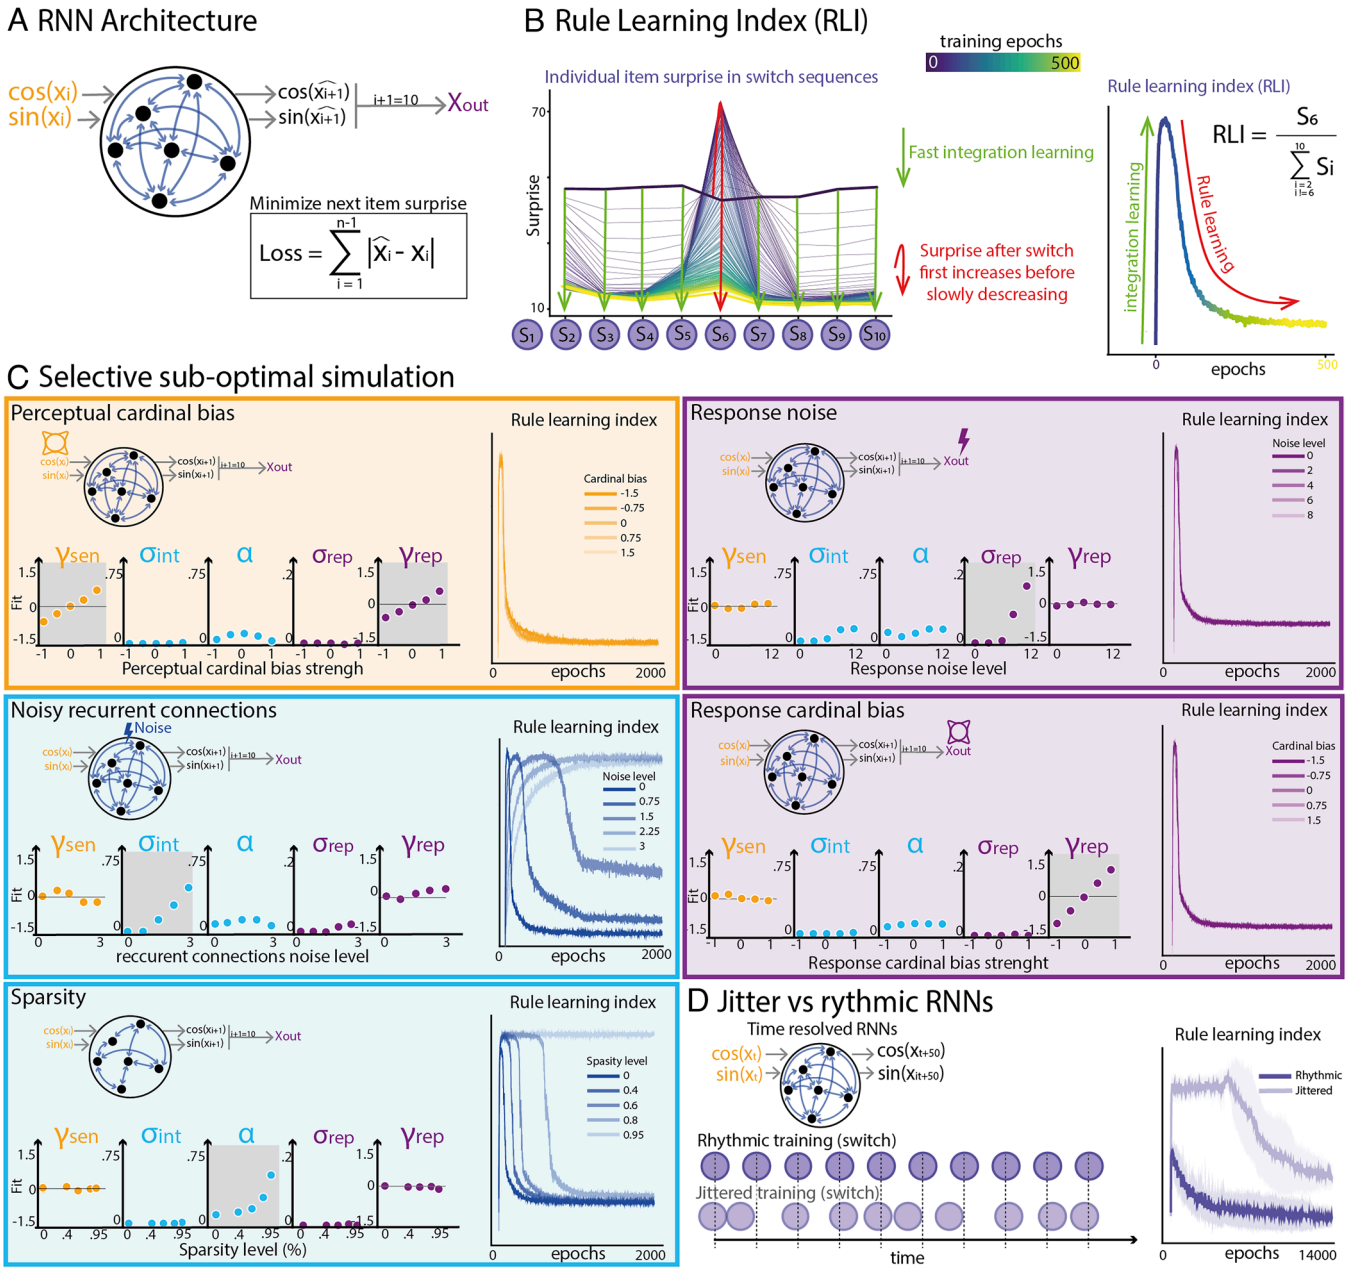

Fig. 5. (*A*) Architecture of the recurrent neural network (RNN). The RNN receives as input the current stimulus orientation (x_i) and is trained to predict the next orientation (x_{i+1}). The difference between the predicted and actual orientation defines the surprise. (*B*, *Left*) Surprise for each stimulus (2–10) in the switch blocks as a function of training epochs (color-coded from purple to yellow). Green arrows highlight the rapid decrease in surprise for stimuli 2/3/4/5 and 7/8/9/10 corresponding to the learning of sensory integration predictions. The red arrow highlights the slower decrease in surprise following the switch, corresponding to the discovery of the rule. (*Right*) This learning dynamic is summarized by the Rule Learning Index (RLI), defined as the ratio of surprises related to rule learning (stimulus 6) and sensory integration (stimuli 2-to-5 and 7-to-10). (*C*) Five variants of the RNN designed to implement each of the five suboptimalities defined in Model 2, with five levels of suboptimalities (from low to high) for each. 1. Perceptual cardinal bias applied to each orientation xi. 2. Integration noise modeled via noisy computation in the recurrent layer. 3. Integration leak obtained by sparsifying to the recurrent layer. 4&5. Response-related parameters were applied at the output level (Xout). Bottom graphs on each panel. After training, we fitted Model 2 to each RNN's predictions on static blocks to estimate how each manipulation impacted each of the parameters. Right graph on each panel. For each level of suboptimality, we estimated the RLI to evaluate how it affected the learning dynamic. (*D*) Comparaison of RLI for rhythmic vs jittered sequences.

in rule-free sequences, predicts the likelihood of discovering a rule in other sequences (Fig. 2*B*). Note that the hidden rule in switch sequences triggered not a gradual but a brutal 90° shift in mean stimulus orientation between the 5th and 6th stimuli, which therefore did not theoretically require a long integration timescale to be detected. Participants who did not find the hidden rule in switch blocks showed moderate, not extreme integration leak in static blocks. As integration timescale likely reflects one's prior belief about environmental volatility, participants who expect a highly volatile environment lower their integration timescale—making them less likely to detect the latent rule. On the other hand, participants who expect a low environmental volatility rely on a larger integration timescale and can look for a latent rule explaining away the change in statistics from the middle of sequences. Importantly, we verified that this effect of sensory integration timescale on rule discovery is not attributable to differences in general attention or working memory and was robust to which block type was first experienced by subjects.

To assess the possible causal role of integration timescale in participants rule-blindness, we applied selective causal perturbations

to RNNs trained to perform the same task as human participants. This effort is part of a growing body of research that leverages these flexible artificial systems to understand cognition, not by fitting their free parameters (connection weights) to behavior (43, 44), but by analyzing the structural and functional configurations which, through constrained optimization, result in the same suboptimalities as biological brains (45–47). Sparsifying recurrent connections produced a graded increase in integration leak during static blocks without affecting any other suboptimality. In contrast, adding computational noise in recurrent units (29, 45) increased integration noise independently of leak. In switch blocks, these manipulations revealed a clear distinction: Increasing sparsity impaired rule discovery without disrupting the training dynamics of sensory integration. In contrast, adding computation noise in the hidden layer (resulting in increased integration noise) impaired both the training of sensory integration and the latency of rule discovery. This distinction was captured in RNN surprise dynamics during training: Networks first reduced surprise to all stimuli except the 6th (indicating sensory integration), then gradually to the 6th (indicating rule discovery). Only integration leak delayed the latter without altering the former. The link between leak and delayed rule discovery generalized to sequences with a different hidden rule structure (with 90° shifts every two stimuli), supporting the generality of this causal relationship. Although the number of trials required to train humans and "blank-slate" RNNs on our task differs by orders of magnitude, the shared impact of sensory integration timescale on rule discovery similarly constrains humans and RNNs.

While integration timescale was reliably and consistently tied to poorer rule discovery abilities both in humans and RNNs, our results also show that discovering the hidden rule modified how human participants integrate stimuli in switch blocks. In humans, discovering the rule causally triggered the mental rotation of individual stimuli prior to integrating them, in order to predict the final (10th) stimulus. This finding directly contradicts dominant theories and computational models of sensory integration, which typically postulate that integration occurs within the "native sensory space" defined by stimulus features (e.g., orientation) (48–52). If sensory integration was encapsulated in this way, rule discovery would result in a post hoc shift applied to the integrated result. This is not what we observed: Participants' behavior was best explained by a model in which stimuli were mentally rotated according to the hidden rule during integration. Such abstract mental manipulations of stimuli have already been described in rule learning tasks, notably in the context of temporal rule in stimulus sequences, such as the reordering (21, 23), mapping (53), or compression (54). These findings are much sparser in the context of noisy sensory integration. Our findings extend previous results showing dissociated neural correlates of the sensory processing of individual stimuli and the integration of their associated category evidence (9, 11, 55).

Moreover, the discovery of the latent rule also protected rule-aware participants from the typical decrease in integration timescale that often accompanies increased volatility (37, 38). Indeed, despite a much larger sensory volatility in switch blocks compared to static blocks, the knowledge of the rule allowed rule-aware participants to explain away this sensory volatility and therefore keep their integration timescale as long as in static blocks. By contrast, rule-blind participants perceived switch blocks as more volatile, and therefore—and adaptively so—reduced drastically their integration timescale reflected by a much larger integration leak. These findings further show that rule-blind participants were not insensitive to the difference between static and switch blocks. Rather, they interpreted the difference between blocks as an increased sensory volatility in switch blocks, instead of a latent rule.

Our study is not the first to investigate the impact of hidden rule discovery on decision-making. However, the hidden rules used in previous studies typically did not interact with the processing of presented stimuli. For example, the hidden rule used in Schuck et al. (56) afforded bypassing entirely the processing of the pattern of dots to focus solely on their color to decide. This bypass prevented testing for the interplay between stimulus processing and rule discovery. By contrast, in our task, a leak in sensory integration directly prevented participants from discovering the hidden rule, an objective limitation in constructing accurate predictions of upcoming stimuli. RNN simulations showed that this effect was not limited to the specific AAAAABBBBB rule used in our task but also generalized to other rules like AABBAABBAA, even though future research in humans should investigate the set of rules over which this effect generalizes. The rule we have used—with a switch in the middle of a sequence of stimuli—corresponds to the "reversal learning" paradigm that has been extensively used

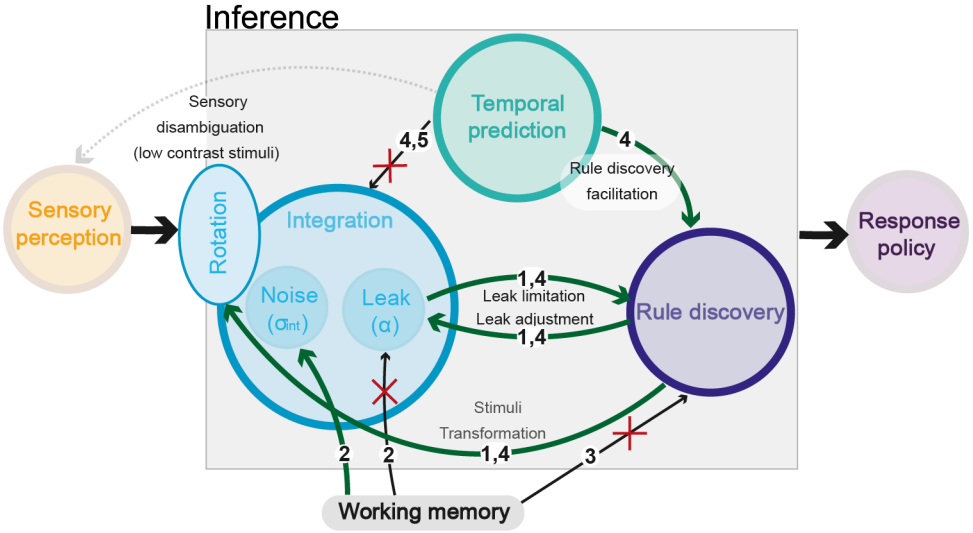

Fig. 6. Visual summary of our findings across all experiments: Green arrows represent demonstrated links between cognitive abilities and crossed black arrows where no such links were found. Numbers correspond to experiments testing each arrow. The gray arrow between temporal prediction and sensory perception corresponds to a known effect of sensory disambiguation, not depicted here.

to study adaptive learning in humans and animals (57–61). We have adapted this paradigm so that the switch followed a hidden rule, which could either be perceived as predictable (if the hidden rule has been discovered) or as randomness in stimulus sequences (if the hidden rule has not been discovered).

RNN perturbations showed that integration leak selectively emerged only for extremely sparse networks, resulting in a substantial decrease in the number of computations performed. Indeed, RNNs with only 5% of active connections showed an integration leak comparable to human participants, without any other suboptimality. The integration leak measured in human participants could therefore partly stem from a cost–benefit optimization, which trades the benefits of efficient sensory integration and rule discovery against the cost of more dense computations. Such cost–benefit arbitration has been postulated to explain other suboptimalities in decision-making such as integration noise, which could arise from the flexibility of associative circuits that can map arbitrary functions of sensory features (here, the integration of stimulus orientations with rule-predicted shifts) onto actions (29, 45). Indeed, the attentional selection and pooling of sensory signals according to a current context (62, 63) by low-rank associative neural circuits has been theorized to be critical for decision accuracy (64).

Finally, we found that temporal predictability [i.e., predicting the timing rather than the content of a perception (65)], was selectively interacting with rule discovery, while having no direct impact on sensory integration. This dissociation suggests that the precise timing at which the latent switch occurred mattered for triggering latent rule discovery, possibly because temporal jitter triggered stronger surprise for all stimuli of the sequence, therefore reducing the relative importance of the surprise after the switch (6th element) compared to the surprise induced by other stimuli of the sequence. By contrast, sensory integration of high-contrast stimuli seems unaffected by temporal jitter—consistent with earlier work showing its robustness to temporal parameters of sequences (41). Together, our findings emphasize that human inferential abilities result from multiple processes with complex interdependencies (Fig. 6). It further highlights the need to conjointly study these several inference abilities to obtain a clear picture of how humans are able to build a rich mental representation of uncertain environments.

Materials and Methods

Participants. All participants for all experiments were UK-based adults (>18 yo), recruited online via the Prolific platform, and gave informed consent. The study was approved by Inserm Ethical Review Committee (IRB00003888 on 13/11/2018). *Experiment 1*: 130 participants performed two sessions of the task on consecutive days (52 female, mean age = 41.5). They were paid £8 for the first and £9 for the second session. *Experiment 2*: 100 participants performed the task (53 female, mean age = 42.3) and performed a single session retributed 7£. *Experiment 3*: 102 participants were recruited online (49 female, mean age = 41.6) and performed a single session retributed £7. *Experiment 4*: 123 participants performed two sessions of the task on consecutive days (72 female, mean age = 41.1). They were paid £8 for the first and £9 for the second session. *Experiment 5*: 126 participants performed two sessions of the task on consecutive days (57 female, mean age = 38.7). They were paid £8 for the first and £9 for the second session.

Experimental Paradigm and Stimuli. All experiments were coded in JavaScript using jspsych toolbox (66).

Experiments 1 and 4 consists of two types of blocks: In static blocks, each stimulus of the sequence is a grating patch whose orientation is drawn from a normal distribution ($\sigma = 15°$). In *switch* blocks, we added a 90° rotation from the middle of the sequences. Importantly, the means of the distributions were randomly chosen for each trial so that the abstract +90° rule must be learned. Each sequence was sequentially presented at 2 Hz, rhythmically in Experiment 1, or with random ISI (uniform [0, 500] ms) for Experiments 4 and 5. The block type was indicated by the color of the circle surrounding the stimuli. Block order was random across participants. In each block, after the presentation of six full training sequences, participants were presented with 20 incomplete sequences, stopping after 3, 5, 7, or 9 stimuli (five of each), and had to predict the orientation of the last (10th) stimulus by rotating a bar, and estimate their confidence level. Four full sequences were randomly placed between the incomplete test sequence to consolidate participant learning. After each answer, they saw colored feedback indicating how far they are from the real orientation. They performed a total of 400 test trials over the two sessions.

Experiment 5 had an identical procedure to Experiments 1 and 4 but with only static trials, half of the blocks being rhythmic, and the other half jittered.

Experiment 2 had participants perform static blocks, identical to Experiment 1, followed by memory blocks in which they were presented with sequences of stimuli and asked on each trial to report the orientation of one of the stimuli in the sequence (selected randomly on each trial). In both conditions, sequences could contain 3, 5, 7, or 9 oriented stimuli. Orientations were drawn uniformly in working memory blocks so that the orientation of a stimulus could not be predicted (by integration) from the orientation of another stimulus better than chance.

Experiment 3 had participants perform sensory integration blocks with the same latent rule as in Experiment 1 (Switch blocks) and two types of working memory blocks measuring their ability to memorize the orientation of stimuli at positions 1 to 5 in the sequence (working memory capacity block) or their ability to maintain the orientation of a cued stimulus at positions 1 to 5 in the sequence (working memory maintenance block).

Preregistration of Experiment 1. The procedure, as well as the primary outcome of this analysis (bidirectional interactions between sensory integration and rule discovery), has been preregistered prior to data collection. The relevant details can be found here: https://aspredicted.org/wy43v.pdf.

Von Mises Distributions Fit. We performed model-free analysis by fitting von Mises distributions of the difference between participants' answers and the average orientation of the five first stimuli (only three first for t + 7 conditions) of the sequence for each condition separately. Von Mises distributions have two parameters: location (θ) which indicates the average direction of the distribution, and the concentration (κ), which reflects variability of the response.

Model Design and Fit. We performed model-based analyses with several models:

Model 1: See details in *SI Appendix*, Fig. S1.

Model 2: This second model comprises 5 parameters. Sensory cardinal bias (γ_{sen}), integration noise (σ_{int}), integration leak (α), response noise (σ_{rep}), and response cardinal bias (γ_{rep}). σ_{int} and σ_{rep} were modeled like in Model 1. For the sensory cardinal bias, each orientation of the sequence was modified through the function $cb(x, \gamma_{sen}) = x + \frac{\sum_n \sin(4nx) I_n(\gamma_{sen})/n}{2I_0(\gamma_{sen})}$, with x the orientation, γ_{sen} the bias strength and I_n the Bessel function of order n. For the response cardinal bias, the same function was applied on the response $cb(resp, \gamma_{rep})$. The leak was modeled with an exponential decrease affecting the weight attributed to each stimulus in the participant response (38, 55, 67)–resp = $\sum_i e^{-i\alpha} x_i$. Integration noise was drawn from a normal distribution of variance σ_{int} and added to each element before summing. Response noise was drawn from a normal distribution of variance σ_{rep} and added to the response. Models 1 and 2 were fitted to static blocks to estimate participants' parameters when performing a sensory integration task.

Model 3: Model 3 is equivalent to Model 2, but we added a parameter modeling a possible fixed offset that participants could apply to all answers. resp = resp + offset.

Model 4: Model 4 is equivalent to Model 2 but with a fixed offset applied to the five first orientations of each sequence. –resp = $\sum_i e^{-i\alpha}(x_i + o_i)$ with o_i = offset for $i \leq 5$ and 0 for $i > 5$.

Model 5: Variation of Model 4 presented in *SI Appendix*, where each of the nine stimuli of the sequence can have a different offset. –resp = $\sum_i e^{-i\alpha}(x_i + o_i)$ with o_i different for each i value.

Like in previous studies (37, 68), we fitted each model using BADS (69) toolbox in MATLAB with 10,000 bootstrap resamples to minimize the negative

log-likelihood and find the best fitting parameters. Models 1 and 2 were fitted on data from static blocks only while Models 3–5 were fitted on data from switch blocks. We assessed for statistical significance between parameters models using the rank-sum test given the non-normality some parameter distributions and applied Bonferroni correction. We validated our model and the fitting procedure using recovery and simulations (*SI Appendix*, Fig. S3).

RNN Simulations. We first designed a simple RNN composed of a 64 × 64 recurrent layer. The input and output orientations are coded by a two-dimensional array with cos and sin. Importantly, the loss of the RNN was chosen to mimic human-like computation. Indeed, when given an input, the networks predict the orientation of the next stimulus. The absolute value of the difference between the prediction and the actual orientation is called the surprise, and the loss of the sequence is the sum of the surprises in the sequence. That way, networks are trained with a human-like strategy.

We then implemented each suboptimalities in the RNN. The sensory cardinal bias was introduced by applying a cardinal bias for each stimulus of the sequence before the input (see Model design for the implementation of this bias). We used five cardinal bias values: [−1, −0.5, 0, 0.5, 1]. To implement the leak, we made the recurrent layer sparse by fixing a given proportion of the connections to zero. This proportion was fixed to 0, 40%, 60%, 80%, or 95%. To implement integration noise, we added random noise in the computation of the recurrent layer activity from a centered normal distribution with a std of 0, 0.75, 1.5, 2.25, or 3. The two response parameters did not modify the RNN per se, but were applied on Xout: either random noise or cardinal bias. To test if those modifications were specifically affecting only one parameter of the sensory learning mechanism, we trained 20 randomly initialized RNNs on 2,000 epochs of 32 complete static sequences, for each parameter value.

We then simulated the behavior after training on an occurrence of our experiment (for comparison purposes between the RNNs, we always used the same 20 task occurrences that correspond to the task presented to the first 20 participants of Experiment 1). We could therefore fit the simulation results of the static sequences with Model 2 and plot the impact of the manipulation on the fitted parameters, averaged across the 20 RNNs (Fig. 5*C*). We explored the impact of these modifications by plotting a summary metric we called *Rule Learning Index* (RLI) capturing the dynamic of the integration and rule learning through training. This metric was computed as the ratio of the surprise elicited by the 6th stimulus and the average surprise elicited by all other stimuli (Fig. 4*B*). For methods on RNNs adapted to time-resolved inputs, see *SI Appendix*, Fig. S10.

Data, Materials, and Software Availability. All data and analysis scripts are publicly available at https://osf.io/5xjwe (70).

ACKNOWLEDGMENTS. This work is supported by a postdoctoral grant Espoir de la recherche from the Fondation pour la Recherche Médicale (FRM) awarded to L.B. (SPF202309017516), a grant from the French National Research Agency (MONODEC, ANR-23-CE37-0028) awarded to V.W., as well as a European grant (ERC, SPEEDY, ERC-CoG-101043344) awarded to B.M.

Author affiliations: [a]Institut National de la Santé et de la Recherche Médicale, Institut de Neurosciences des Systèmes, Aix-Marseille Université, Marseille 13010, France; [b]Institut National de la Santé et de la Recherche Médicale, Laboratoire de Neurosciences Cognitives et Computationnelles, Paris 75005, France; and [c]Département d'Études Cognitives, École Normale Supérieure–Université Paris Sciences et Lettres, Paris 75005, France

1. R. Bogacz, Optimal decision-making theories: Linking neurobiology with behaviour. *Trends Cogn. Sci.* **11**, 118–125 (2007).
2. B. Morillon, C. E. Schroeder, V. Wyart, Motor contributions to the temporal precision of auditory attention. *Nat. Commun.* **5**, 5255 (2014).
3. C. Summerfield, T. Egner, Expectation (and attention) in visual cognition. *Trends Cogn. Sci.* **13**, 403–409 (2009).
4. A. Wald, J. Wolfowitz, Optimum character of the sequential probability ratio test. *Ann. Math. Stat.* **19**, 326–339 (1948).
5. P. Janata, S. T. Tomic, J. M. Haberman, Sensorimotor coupling in music and the psychology of the groove. *J. Exp. Psychol. Gen.* **141**, 54–75 (2012).
6. A. C. Nobre, F. van Ede, Anticipated moments: Temporal structure in attention. *Nat. Rev. Neurosci.* **19**, 34–48 (2018).
7. I. Wollman, B. Morillon, Organizational principles of multidimensional predictions in human auditory attention. *Sci. Rep.* **8**, 13466 (2018).
8. A. Zalta, S. Petkoski, B. Morillon, Natural rhythms of periodic temporal attention. *Nat. Commun.* **11**, 1051 (2020).
9. S. Cheadle *et al.*, Adaptive gain control during human perceptual choice. *Neuron* **81**, 1429–1441 (2014).
10. D. J. Freedman, J. A. Assad, Neuronal mechanisms of visual categorization: An abstract view on decision making. *Annu. Rev. Neurosci.* **39**, 129–147 (2016).
11. V. Wyart, V. de Gardelle, J. Scholl, C. Summerfield, Rhythmic fluctuations in evidence accumulation during decision making in the human brain. *Neuron* **76**, 847–858 (2012).
12. L. Benjamin, M. Sablé-Meyer, A. Fló, G. Dehaene-Lambertz, F. A. Roumi, Long-horizon associative learning explains human sensitivity to statistical and network structures in auditory sequences. *J. Neurosci.* **44**, e1369232024 (2024).
13. L. Benjamin, A. Fló, F. Al Roumi, G. Dehaene-Lambertz, Humans parsimoniously represent auditory sequences by pruning and completing the underlying network structure. *Elife* **12**, e86430 (2023).
14. L. Benjamin *et al.*, Tracking transitional probabilities and segmenting auditory sequences are dissociable processes in adults and neonates. *Dev. Sci.* **26**, e13300 (2023).
15. S. Henin *et al.*, Learning hierarchical sequence representations across human cortex and hippocampus. *Sci. Adv.* **7**, 1–13 (2021).
16. J. R. Saffran, R. N. Aslin, E. L. Newport, Statistical learning by 8-month-old infants. *Science* **274**, 1926–1928 (1996).
17. A. C. Schapiro, N. B. Turk-Browne, K. A. Norman, M. M. Botvinick, Statistical learning of temporal community structure in the hippocampus. *Hippocampus* **26**, 3–8 (2016).
18. A. C. Schapiro, T. T. Rogers, N. I. Cordova, N. B. Turk-Browne, M. M. Botvinick, Neural representations of events arise from temporal community structure. *Nat. Neurosci.* **16**, 486–492 (2013).
19. F. Al Roumi, S. Marti, L. Wang, M. Amalric, S. Dehaene, Mental compression of spatial sequences in human working memory using numerical and geometrical primitives. *Neuron* **109**, 2627–2639.e4 (2021).
20. F. Al Roumi, S. Planton, L. Wang, S. Dehaene, Brain-imaging evidence for compression of binary sound sequences in human memory. *Elife* **12**, e84376 (2023).
21. P. Albouy, Z. E. Martinez-Moreno, R. S. Hoyer, R. J. Zatorre, S. Baillet, Supramodality of neural entrainment: Rhythmic visual stimulation causally enhances auditory working memory performance. *Sci. Adv.* **8**, eabj9782 (2022).
22. J. Quilty-Dunn, N. Porot, E. Mandelbaum, The best game in town: The re-emergence of the language of thought hypothesis across the cognitive sciences. *Behav. Brain Sci.* **46**, e261 (2022), 10.1017/S0140525X22002849.
23. Y. Xie *et al.*, Geometry of sequence working memory in macaque prefrontal cortex. *Science* **375**, 632–639 (2022).
24. J. Fiser, G. Lengyel, Statistical learning in vision. *Annu. Rev. Vis. Sci.* **8**, 265–290 (2022).
25. W. J. Ma, K. Kording, D. Goldreich, *Bayesian Models of Perception and Action: An Introduction* (The MIT Press, London, England, 2023).
26. R. Ratcliff, G. McKoon, The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Comput.* **20**, 873–922 (2008).
27. J. M. Beck, W. J. Ma, X. Pitkow, P. E. Latham, A. Pouget, Not noisy, just wrong: The role of suboptimal inference in behavioral variability. *Neuron* **74**, 30–39 (2012).
28. J. Drugowitsch, V. Wyart, A.-D. Devauchelle, E. Koechlin, Computational precision of mental inference as critical source of human choice suboptimality. *Neuron* **92**, 1398–1411 (2016).
29. C. Findling, V. Wyart, Computation noise in human learning and decision-making: Origin, impact, function. *Curr. Opin. Behav. Sci.* **38**, 124–132 (2021).
30. D. Rahnev, R. N. Denison, Suboptimality in perceptual decision making. *Behav. Brain Sci.* **41**, e223 (2018).
31. S. Dehaene, F. Al Roumi, Y. Lakretz, S. Planton, M. Sablé-Meyer, Symbols and mental programs: A hypothesis about human singularity. *Trends Cogn. Sci.* **26**, 751–766 (2022).
32. S. Wu, M. Thalmann, E. Schulz, Two types of motifs enhance human recall and generalization of long sequences. *Commun. Psychol.* **3**, 3 (2025).
33. M. Kleinman, C. Chandrasekaran, J. Kao, A mechanistic multi-area recurrent network model of decision-making. *Adv. Neural Information Processing Syst.* **34**, 23152–23165 (2021).
34. T. Kurikawa, Different timescales of neural activities introduce different representations of task-relevant information. bioRxiv [Preprint] (2024). https://doi.org/10.1101/2024.07.23.604720 (Accessed 23 July 2024).
35. B. Lin, D. Bouneffouf, G. Cecchi, Predicting human decision making in psychological tasks with recurrent neural networks. *PLoS One* **17**, e0267907 (2022).
36. A. R. Girshick, M. S. Landy, E. P. Simoncelli, Cardinal rules: Visual orientation perception reflects knowledge of environmental statistics. *Nat. Neurosci.* **14**, 926–932 (2011).
37. J. K. Lee, M. Rouault, V. Wyart, Adaptive tuning of human learning and choice variability to unexpected uncertainty. *Sci. Adv.* **9**, eadd0501 (2023).
38. O. Ossmy *et al.*, The timescale of perceptual evidence integration can be adapted to the environment. *Curr. Biol.* **23**, 981–986 (2013).
39. I. Gómez Varela, J. Orpella, D. Poeppel, P. Ripolles, M. F. Assaneo, Syllabic rhythm and prior linguistic knowledge interact with individual differences to modulate phonological statistical learning. *Cognition* **245**, 105737 (2024).
40. M. Maheu, F. Meyniel, S. Dehaene, Rational arbitration between statistics and rules in human sequence processing. *Nat Hum Behav* **6**, 1087–1103 (2022).
41. M. L. Waskom, R. Kiani, Decision making through integration of sensory evidence at prolonged timescales. *Curr. Biol.* 28, 3850–3856.e9 (2018).
42. V. Wyart, E. Koechlin, Choice variability and suboptimality in uncertain environments. *Curr. Opin. Behav. Sci.* **11**, 109–115 (2016).
43. M. K. Eckstein, C. Summerfield, N. D. Daw, K. J. Miller, Predictive and interpretable: Combining artificial neural networks and classic cognitive models to understand human learning and decision making. bioRxiv [Preprint] (2023). https://doi.org/10.1101/2023.05.17.541226 (Accessed 17 May 2023).
44. J. C. Peterson, D. D. Bourgin, M. Agrawal, D. Reichman, T. L. Griffiths, Using large-scale experiments and machine learning to discover theories of human decision-making. *Science* **372**, 1209–1214 (2021).
45. C. Findling, V. Wyart, Computation noise promotes zero-shot adaptation to uncertainty during decision-making in artificial neural networks. *Sci. Adv.* **10**, eadl3931 (2024).
46. T. Flesch, K. Juechems, T. Dumbalska, A. Saxe, C. Summerfield, Orthogonal representations for robust context-dependent task performance in brains and neural networks. *Neuron* **110**, 1258–1270.e11 (2022).

47. M. Molano-Mazón *et al.*, Recurrent networks endowed with structural priors explain suboptimal animal behavior. *Curr. Biol.* **33**, 622–638.e7 (2023).
48. S. Deneve, P. E. Latham, A. Pouget, Efficient computation and cue integration with noisy population codes. *Nat. Neurosci.* **4**, 826–831 (2001).
49. B. Ermentrout, Neural networks as spatio-temporal pattern-forming systems. *Rep. Prog. Phys.* **61**, 353 (1998).
50. J. M. Esnaola-Acebes, A. Roxin, K. Wimmer, Flexible integration of continuous sensory evidence in perceptual estimation tasks. *Proc. Natl. Acad. Sci. U.S.A.* **119**, e2214441119 (2022).
51. M. Jazayeri, J. A. Movshon, Optimal representation of sensory information by neural populations. *Nat. Neurosci.* **9**, 690–696 (2006).
52. A. Pouget, P. Dayan, R. S. Zemel, Inference and computation with population codes. *Annu. Rev. Neurosci.* **26**, 381–410 (2003).
53. M. M. Garvert, R. J. Dolan, T. E. J. Behrens, A map of abstract relational knowledge in the human hippocampal–entorhinal cortex. *Elife* **6**, 1–20 (2017).
54. S. Planton *et al.*, A theory of memory for binary sequences: Evidence for a mental compression algorithm in humans. *PLoS Comput. Biol.* **17**, e1008598 (2021), 10.1371/journal.pcbi.1008598.
55. V. Wyart, N. E. Myers, C. Summerfield, Neural mechanisms of human perceptual choice under focused and divided attention. *J. Neurosci.* **35**, 3485–3498 (2015).
56. N. W. Schuck *et al.*, Medial prefrontal cortex predicts internally driven strategy shifts. *Neuron* **86**, 331–340 (2015).
57. D. A. Hamilton, J. L. Brigman, Behavioral flexibility in rats and mice: Contributions of distinct frontocortical regions. *Genes Brain Behav.* **14**, 4–21 (2015).
58. V. D. Costa, V. L. Tran, J. Turchi, B. B. Averbeck, Reversal learning and dopamine: A bayesian perspective. *J. Neurosci.* **35**, 2407–2416 (2015).
59. L. Clark, R. Cools, T. W. Robbins, The neuropsychology of ventral prefrontal cortex: Decision-making and reversal learning. *Brain Cogn.* **55**, 41–53 (2004).
60. C. M. Glaze, J. W. Kable, J. I. Gold, Normative evidence accumulation in unpredictable environments. *Elife* **4**, e08825 (2015).
61. A. Izquierdo, J. L. Brigman, A. K. Radke, P. H. Rudebeck, A. Holmes, The neural basis of reversal learning: An updated perspective. *Neuroscience* **345**, 12–26 (2017).
62. V. Mante, D. Sussillo, K. V. Shenoy, W. T. Newsome, Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* **503**, 78–84 (2013).
63. F. Pestilli, M. Carrasco, D. J. Heeger, J. L. Gardner, Attentional enhancement via selection and pooling of early sensory responses in human visual cortex. *Neuron* **72**, 832–846 (2011).
64. A. Dubreuil, A. Valente, M. Beiran, F. Mastrogiuseppe, S. Ostojic, The role of population structure in computations through neural dynamics. *Nat. Neurosci.* **25**, 783–794 (2022).
65. B. Morillon, S. Baillet, Motor origin of temporal predictions in auditory attention. *Proc. Natl. Acad. Sci. U.S.A.* **114**, E8913–E8921 (2017).
66. J. R. de Leeuw, jsPsych: A JavaScript library for creating behavioral experiments in a web browser. *Behav. Res.* **47**, 1–12 (2015).
67. M. Usher, J. L. McClelland, The time course of perceptual choice: The leaky, competing accumulator model. *Psychol. Rev.* **108**, 550–592 (2001).
68. J. Drevet, J. Drugowitsch, V. Wyart, Efficient stabilization of imprecise statistical inference through conditional belief updating. *Nat. Hum. Behav.* **6**, 1691–1704 (2022).
69. L. Acerbi, W. J. Ma, Practical bayesian optimization for model fitting with bayesian adaptive direct search. arXiv [Preprint] (2017). https://doi.org/10.48550/arXiv.1705.04405 (Accessed 2 November 2017).
70. L. Benjamin, B. Morillon, V. Wyart Data from "Sensory *integration*, temporal prediction and rule discovery reflect interdependent inference processes". OSF. https://osf.io/5xjwe. Deposited 17 November 2025.