Simplicity guides the discovery and use of compositionality

Valerio Rubino[1,2*], **Peter Dayan**[2,3], **and Charley M. Wu**[2,4,5,6]

[1]Center for Mind/Brain Science, University of Trento, 38068 Italy
[2]Max Planck Institute for Biological Cybernetics, Tübingen, 72076 Germany
[3]University of Tübingen, Tübingen, 72074 Germany
[4]Max Planck Institute for Human Development, Berlin 14195 Germany
[5]Centre for Cognitive Science, Institute of Psychology, Technical University of Darmstadt, Darmstadt, Germany
[6]Hessian.AI, Darmstadt, Germany
[*]valerio.rubino@unitn.it

ABSTRACT

Human cognition derives great power from compositionality, often formalized through probabilistic Languages of Thought. However, identifying useful compositional representations is computationally demanding, requiring search over vast, combinatorial hypothesis spaces. Simplicity biases are a key signature of human behaviour in such settings, however many forms of simplicity are unreasonably complex to compute. Across three maze-navigation experiments, we examined how people generate and deploy compositional hypotheses under time pressure (Exp. 1), without instruction (Exp. 2), and in statistically matched non-compositional environments (Exp. 3). Through both behavioural analyses and computational modeling, we show that behaviour was best predicted in all settings by a novel heuristic measure of fragment simplicity. Simplicity in this form avoids the need to identify globally minimal expressions, and was both robust to time pressure and generally accelerated response times. This pattern was observed at multiple scales of behaviour, for both individual actions and abstract sequences of spatial primitives (templates), with participants exhibiting a robust bias towards simpler hypotheses than warranted by the structure of the mazes. Together, our results suggest that by adopting a heuristic form of simplicity, we overcome the difficulty of generating compositional hypotheses.

Introduction

Humans excel at making sense of novel situations from limited information. Consider the challenge of navigating an unfamiliar airport in a foreign country while rushing to catch a flight. Even without prior exposure to that specific airport, people can effortlessly draw on familiar elements—security checkpoints, departure gates, duty-free shops—to generate and test plausible hypotheses about where to look. This ability hinges on being able to form and evaluate a potentially unbounded number of hypotheses from a finite set of building blocks in a compositional manner[1–3]. Thus, this capacity for compositional reasoning is commonly claimed as a potentially unique factor[4] distinguishing us from other animals[5–8] and enabling our astonishing potential for creativity and generalization[9].

According to an influential proposal[9,10], compositionality presupposes a Language of Thought (LoT), whose vocabulary and syntax enable us to generate an unbounded number of hypotheses. This emphasis on unbounded generativity echoes Wilhelm von Humboldt's famous characterization of language as "the infinite use of finite means"[11]. Yet the very openness of the LoT hypothesis implies that any event or observation is compatible with an infinite number of potential explanations. As a result, additional criteria are required for narrowing down the number of possibilities[12,13], often taking the form of simplicity biases defined over a domain's representational medium[10]. Such biases are widespread in cognition[14,15]: item memory and comparison difficulty scale with the complexity of the simplest description[16–20], and under incomplete information, people preferentially extrapolate by extending the simplest pattern consistent with the data[6,21]. However, despite their ubiquity, the origins of these simplicity biases remain poorly understood.

Early LoT models sought to formalize the generation and evaluation of simple compositional hypotheses through explicit, rule-based inference mechanisms[22,23], but were criticized as being too rigid to capture the flexibility of human cognition[24,25]. More recent approaches have developed probabilistic Language of Thought (pLoT) models[4,26], which achieve greater flexibility by maintaining probability distributions over candidate compositional hypotheses[3,25,27]. These distributions typically combine prior probabilities and likelihoods measuring the fit to observations, derived by evaluating individual hypotheses[28,29]. Simplicity biases then naturally emerge by assigning higher prior probability to simple hypotheses. The effectiveness of this probabilistic approach is evident through successful accounts of human cognition across a wide range of settings, from learning to categorize concepts[25,27,30], identifying and remembering visuospatial patterns[31–33] and recognizing violations in geometrical shapes[8] and auditory sequences[20,34].

However, while pLoT models cope with some of the limitations of the rigid early LoT proposals[35], their algorithmic implementation remains unclear, particularly with respect to feasibility under realistic resource constraints. That is, understanding compositional reasoning not only requires us to appreciate the open-ended productivity of human cognition, but also to confront the “finite means” imposed by our limited computational resources[36,37]. Whether compositional representations ultimately help or hinder reasoning under resource constraints remains an open question. On one hand, compositional representations afford more efficient compression[38–40], potentially making it cheaper to store and execute solutions with limited cognitive resources. On the other hand, compositionality results in a combinatorial explosion of hypotheses, making search exponentially more difficult with each added component or relational dependency[39,41,42].

One promising strategy for mitigating the cognitive costs of searching for compositional representations is to build and maintain libraries of reusable primitives[26,34,43,44]. Such libraries allow frequently co-occurring combinations of primitives to be cached and reused, thereby accelerating hypothesis search rather than starting from scratch each time[42]. Consistently, exposure to familiar combinations of primitives biases subsequent hypotheses toward incorporating them[45–47]. Ideally, people would maintain libraries of primitives useful for their tasks, potentially developing into domain-specific languages. However, it remains unclear what factors shape the acquisition and use of newly learned primitives, and how these mechanisms are themselves constrained by limited cognitive resources.

Goals and scope

Here, we investigate how real-world cognitive constraints and hierarchical structure shape the generation and use of compositional hypotheses, using a maze navigation task with an underlying hidden path (Fig. 1a). In Exp. 1, participants were subject to time pressure (within-subject: limited vs. unlimited) to reduce the availability of cognitive resources (Fig. 1b). In Exp. 2, participants navigated the same mazes but without prior instruction about the compositional structure of the task (Fig. 1c). Lastly, in Exp. 3, we presented participants with mazes lacking any compositional structure (but matched in conditional probabilities) to assess spontaneous inductive biases in hypothesis testing.

Across all three experiments, we find that participants' accuracy is best predicted by a novel measure inspired by Alexander and Carey[16], which we call *fragment simplicity*. This heuristic counts how often substrings of any length are repeated or mirrored in a sequence, across multiple scales (Fig. 1d). Throughout, we use *simplicity* to refer to the theoretical preference for simpler latent hypotheses, and *fragment simplicity* to refer to our operational measure of that preference. The fragments we consider are either at the level of individual actions or templates (abstract sequences of spatial primitives), yielding the derived quantities *action simplicity* and *template simplicity*. Participants exhibited a pervasive, multiscale inductive bias towards simplicity defined at the both of these levels, consistently exceeding what was warranted by the structure of the mazes. This simplicity bias was also associated with faster responses, suggesting it reflects an efficient heuristic for compositional hypothesis generation given the finite means of human cognition.

Results

Across three experiments, participants navigated a series of mazes, each with a single hidden correct path. Participants iteratively navigated the mazes by selecting an action (*left*, *right*, or *up*), and received feedback for each choice. Correct actions resulted in progressing to the next tile, while incorrect actions resulted in staying at the current location and losing a “life”. We incentivized participants to reach the goal row of each maze (Fig. 1a) while making as few mistakes as possible.

Exp. 1: Compositionality is robust to time pressure

In Exp. 1, we informed participants that the path in each maze was generated as a composition of two out of six primitives, which they were previously shown. Participants received both written instructions and an interactive tutorial. During maze navigation, each successfully completed primitive was visually highlighted, to help participants acquire and use them. We used a within-subject manipulation of time pressure (unlimited vs. limited to 20 s in each maze) to test how resource limitations affected the use of compositional representations.

Participants leverage primitives

Overall, participants achieved 74% accuracy in unlimited time and 69% in limited time rounds (Fig. 2a), as measured by the proportion of correct action choices, $p(\text{correct step})$. As expected, time pressure significantly decreased accuracy ($t(65) = 8.4$, $p < .001$, $d = 0.9$), with participants unable to finish in 46% $\pm$ 3.83 of rounds (see Fig. S2a). Using a mixed-effects logistic regression including participants as random intercepts and predicting whether each action was correct or incorrect, we found that the negative effect of time pressure (Odds Ratio, OR: 0.84 [0.81, 0.87], $p < .001$) was significant even when controlling for the influence of timed-out rounds (OR: 0.91 [0.87, 0.95], $p < .001$). Nonetheless, participants' accuracy in limited time rounds still surpassed simulations relying only on the true underlying marginal probabilities ($t(65) = 17.7$, $p < .001$, $d = 2.2$) and conditional probabilities ($t(65) = 12.8$, $p < .001$, $d = 1.6$) of actions. Thus, participants did not rely solely on the Markovian statistics of the task, but instead used the hierarchical compositional structure in both time conditions.

As an initial test for whether participants learned and exploited primitives, we contrasted *between-primitives* steps (i.e., choosing between primitive A and B at the start of a new 4-action sequence) against *within-primitive* steps (i.e., completing a primitive once it has been started). Intuitively, participants who infer and exploit the compositional structure of the maze should commit fewer errors in within- than between-primitives steps. We define a *primitive use index* as the difference between *within-primitive* and *between-primitives* ac-

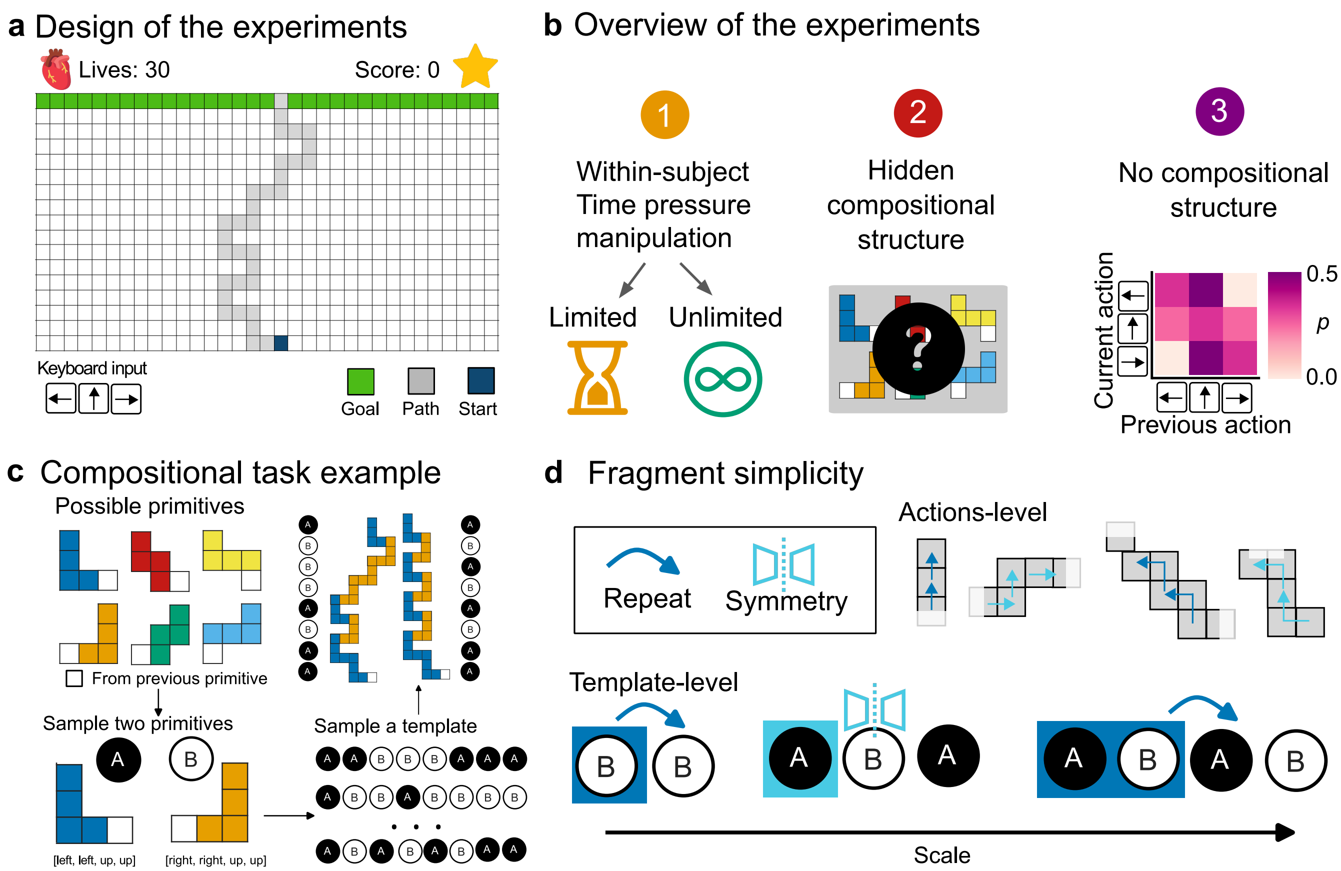

Figure 1. Task details. **a**) Task screenshot. Participants navigated from the start tile (blue) to the goal row (in green) by selecting actions (*left*, *right*, and *up*) using the corresponding arrow keys. Progress on each maze required following a hidden path (grey), which was compositionally generated from pairs of spatial primitives in Exps. 1-2 (see Fig. S1a,b). Incorrect moves resulted in the loss of a life. **b**) Exp. 1 used a within-subject manipulation of time pressure (limited vs. unlimited time). Exp. 2 omitted explicit instructions about the compositional structure. Exp. 3 used paths containing no compositional structure, but with matched Markovian statistics (see Fig. S1c,d). **c**) For Exp. 1 and Exp. 2, each path was generated by selecting two spatial primitives from a set of six (left panel) and combining them according to an abstract template (lower right panel). **d**) Fragment simplicity was defined either at the action or template level, based on the number of repetitions and symmetries across scales.

curacy (Fig. 2b). Although indirect, this measure indexes primitive use because it tracks the performance benefit on steps whose actions can be predicted by participants who know and use primitives. Primitive use was lower under time pressure ($t(65) = -2.1$, $p = .040$, $d = 0.3$), although within-primitive steps were still easier for participants even under time pressure ($t(65) = 21.4$, $p < .001$, $d = 2.1$). Although the overall effect was only marginally significant, the reduction in primitive use under time pressure became substantially stronger when we restricted the analysis to trials occurring after participants had discovered both primitives, thus providing a cleaner test of primitive use independent of primitive discovery ($t(65) = -3.6$, $p < .001$, $d = 0.5$; Fig. S2b). Applying mixed-effects logistic regression provides further confirmatory evidence, with participants achieving higher within-primitive accuracy (OR: 1.77 [1.70, 1.84], $p < .001$) with time pressure reducing this effect (OR: 0.92 [0.87, 0.98], $p = .008$; Fig. S2c).

These results demonstrate that participants learned and used primitives in both conditions, although limiting cognitive resources diminished this effect. Furthermore, they suggest that participants draw hypotheses at the higher-level of templates rather than at the action level.

Simplicity predicts accuracy

We next sought to understand how participants combine primitives by forming hypotheses about latent templates. A longstanding line of research has shown that participants' performance is sensitive to the simplicity of the latent hypothesis explaining the observations[17,20,21]. However, multiple formalizations of simplicity have been proposed. To identify the measure that best captures participants' behaviour, we ran a series of mixed-effects logistic regressions predicting accuracy from baseline regressors (time pressure, round, and step number; Fig. 2c dashed line) together with candidate complexity measures computed over maze paths at both the

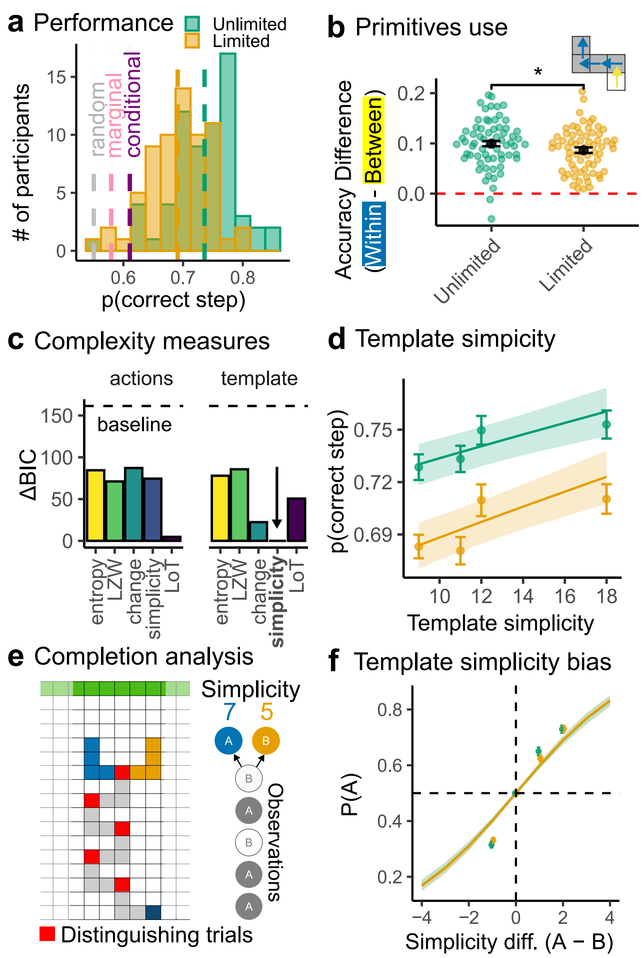

Figure 2. Experiment 1 results. **a**) Performance measured in accuracy, p(correct step), compared with simulated baselines (random actions or using ground-truth task marginal or conditional probabilities). **b**) Time pressure reduced primitive use, measured as the difference in accuracy for within-primitive versus between-primitive trials, though it remained above zero. Black dots show averages across participants and error bars denote the standard error of the mean (s.e.m.). Asterisks indicate statistical significance (*, $p < .05$; **, $< .01$; ***, $p < .001$). **c**) Comparison of complexity measures predicting participant accuracy (mixed-effects logistic regression), using Δ Bayesian Information Criterion (ΔBIC) relative to the best-fitting model (template simplicity) where lower is better. The black dotted line is the ΔBIC of a baseline model without any complexity measure but including time pressure, step, and round number. **d**) In the best-fitting model, template simplicity predicted higher accuracy, with no interaction with time pressure. Unless otherwise noted, dots show averages within quantile bins (computed across participants). **e**) Completion analysis to predict participant choices in distinguishing steps, where choices uniquely corresponded to hypotheses about the next primitive, with different template simplicity values. **f**) Participants showed a simplicity bias, preferentially selecting actions leading to completions with greater template-level fragment simplicity.

action and template levels.

In line with previous studies[20], we compared *entropy*[48] (measuring unpredictability based on ground-truth frequencies), Lempel-Ziv-Welch (*LZW*) compressibility[49] (using an algorithm that compresses strings by progressively building a dictionary of repeated patterns), *change complexity*[19] (counting changes in substrings, weighted by length) and *LoT complexity*[20] (measuring the length of the shortest expression describing the paths in a LoT they defined). Lastly, we included a novel measure of *fragment simplicity* inspired by[16], which simply counts the number of repetitions and symmetries across scales (Fig. 1d; see Methods and Fig. S2d for the correlation among measures).

The best-fitting model used *template simplicity* (i.e. fragment simplicity at the level of templates) along with the baseline regressors and their interactions (BIC: 123746, Fig. 2c; see Fig. S4e-g for details). In this winning model, accuracy increased with template simplicity (OR: 1.10 [1.07, 1.12], $p < .001$). This result suggests that participants hypothesized that maze paths were simple at the level of templates, so that they achieved higher accuracy when the correct paths matched their hypotheses. Crucially, while time pressure decreased accuracy across the board (OR: 0.84 [0.82, 0.86], $p < .001$), there was no interaction with template simplicity (OR: 1.02 [1.00, 1.05], $p = .101$). Therefore, time pressure did not change the influence of template simplicity on accuracy: performance was equivalently predicted by template simplicity in both unlimited (OR: 1.07 [1.04, 1.09], $p < .001$) and limited conditions (OR: 1.08 [1.06, 1.11], $p < .001$).

Among action-level measures, LoT complexity was the best predictor, and the second overall (BIC: 123848.3). However, adding it to the winning model did not improve the fit (BIC: 123754; $\chi^2(1) = 3.54$, $p = .060$), while adding action simplicity did (BIC: 123747; $\chi^2(1) = 9.61$, $p = .002$). In this augmented model (both action and template simplicity), accuracy increased with action simplicity (OR: 1.02 [1.01, 1.04], $p = .002$) and, again, time pressure did not attenuate the effect of template simplicity (OR: 1.03 [1.00, 1.05], $p = .084$).

Thus, the cognitive limitations imposed by time pressure generally reduced accuracy, but did not impair the predictive influence of template simplicity. This suggests that participants exploited the compositional structure of the task and favoured simpler primitive combinations to the same degree, even under time pressure. We next present a finer-grain analysis characterising how participants test hypotheses at a trial-by-trial level.

Simplicity bias in hypothesis testing

To identify a potential simplicity bias in the hypotheses participants generated about the compositional structure of the correct path, we conducted a *completion analysis* (Fig. 2e) to predict choices based on the observed structure of the maze. Because fragment simplicity best captured participants' performance, we used it to operationalize simplicity bias in the completion analysis. After observing a subtemplate such as AABAB, participants could choose either primitive A or B to continue the sequence, yielding subtemplates with simplicity values of 7 and 6, respectively. Thus, choosing primitive A over B would favour simplicity. We focused on *distinguishing steps* (see Methods), where the chosen primitive and its corresponding completion could be inferred from the participant's actions.

Using a mixed-effects logistic regression to model choices on distinguishing steps, we found that participants systematically favoured completions with higher template simplicity (OR: 1.67 [1.56, 1.78], $p < .001$; Fig. 2f). This effect did not interact with time pressure (OR: 1.01 [0.94, 1.09], $p = .829$). In the same model, we included additional regressors (see Fig. S2g) to control for lower-level predictors, such as repeating the last primitive (OR: 3.60 [3.23, 3.93], $p < .001$), and choosing primitives whose action yielded to higher action simplicity (OR: 1.14 [1.08, 1.21], $p < .001$). Furthermore, correct actions were more likely to be chosen (OR: 1.24 [1.15, 1.33], $p < .001$). Therefore, participants generated hypotheses favouring simpler completions in a way that was resilient to time pressure.

Computational modelling of template simplicity bias

Finally, we fit computational models to predict trial-by-trial choices of each participant (Fig. 3a). These models are themselves compositional, with statistical learning (S), primitives inference (P), template simplicity (T), and action simplicity (A) components (see Methods for a full description).

Statistical learning (S) learns the conditional probability of actions by updating the parameters of a Dirichlet distribution. *Primitives inference* (P) uses sequential Monte Carlo sampling (i.e. i.e., a particle filter[22,50,51]) to infer which pair of primitives (out of the total of six) was used to generate each maze: $p(\text{primitives}|\text{data}) \propto p(\text{data}|\text{primitives})p(\text{primitives})$. Briefly this is accomplished by propagating a finite set of hypotheses ($n = 100$ for generality), which are then weighted and resampled according to the likelihood of the observed feedback (see Methods for details). Combining statistical and primitives inference (SP), we use a mixture parameter γ that determines the relative contribution of P over S (weighted by $1-\gamma$). *Template simplicity* (T) biases how inferred primitives (from P) are sampled from the posterior $p(\text{primitives} \mid \text{data})$. β_T is a free parameter scaling the bonus for primitives that lead to completions with greater template simplicity. Whenever the model has to choose between primitives, it considers completions by attaching candidate primitives to the observed subtemplate. The model computes the simplicity advantage of each completion by subtracting from its simplicity the baseline simplicity of the observed subtemplate. β_T scales the simplicity advantages, modulating the preference for primitives yielding simpler completions. Similarly, we also included an *action simplicity* (A) component, which biases the distribution over actions by the bonus term β_A, before it is combined with S. Lastly, all models include a source of random errors ε, defining the probability of uniformly sampling from all valid actions. After fitting each model variant to participants' data, we used them to simulate data and verified all models were recoverable and distinguishable from each other (see Fig. S3a,b).

The full SPTA model using all four components provided the best fit for participants in both limited and unlimited time conditions (Fig. 3b; mean BIC=989.02 $\pm$ 18.00; PXP $\approx$ 1.00). Likelihood-ratio tests against simpler models also confirmed each component of the SPTA model was justified (all $p <$ .001).

We then inspected the estimated parameters of the winning SPTA model (Fig. 3c; see Fig. S3c-f for parameter recovery). Consistent with the behavioural analyses, time pressure did not modulate the template simplicity bias β_T ($t(65) = 0.1$, $p = .899$, $d = 0.02$), which remained robust across conditions. Furthermore, as with our behavioural index of primitive use (Fig. 2b), estimates of γ were significantly lower under time pressure ($t(65) = -3.7$, $p < .001$, $d = 0.3$), even while the model simultaneously accounted for an increase in random errors ε ($t(65) = 3.6$, $p < .001$, $d = 0.5$). Lastly, time pressure did not have a significant effect on action simplicity β_A ($t(65) = -2.0$, $p = .054$, $d = 0.3$).

To aid interpretation of the model parameters, we subsequently investigated their behavioural correlates (Fig. S3d,e). For each participant we computed accuracy and their primitive use index, while also running separate completion analyses to quantify their template and action simplicity biases. To better understand how model parameters influenced accuracy, we fit a mixed-effects logistic regression predicting trial-level accuracy, including random intercepts for participants and controlling for the effect of time pressure (Fig. S3d,e). We found that accuracy increased with γ (OR: 1.23 [1.21, 1.25], $p < .001$) and β_A (OR: 1.06 [1.04, 1.07], $p < .001$), while it unsurprisingly decreased with ε (OR: 0.90 [0.89, 0.91], $p < .001$). However, the effect of β_T was not significant (OR: 1.01 [0.99, 1.02], $p = .314$), consistent with the fact that the mazes were generated across a wide range of template simplicity values and thus did not incentivize this bias.

Lastly, we investigated how model parameters influenced reaction times (RTs). We fit a mixed-effects linear regression predicting trial-level RTs (Fig. S4a) and selected the interaction structure using a stepwise model comparison approach. We found that higher γ values slowed responses ($b = 0.18$, [0.16, 0.19], $p < .001$), whereas β_T sped up RTs ($b = -0.02$, [-0.03, -0.01], $p = .001$) albeit less so under time pressure ($b = 0.06$, [0.05, 0.07], $p < .001$). To better understand how β_T affected RTs, we restricted the analysis to distinguishing trials, where participants had to choose between primitives. In these trials, we found that RTs decreased with β_T ($b = -0.06$, [-0.10, -0.03], $p < .001$) and this effect was stronger when the chosen completion was simpler ($b = -0.02$, [-0.04, -0.01], $p = .004$). Thus, a stronger template simplicity bias enabled participants to generate faster hypotheses, while primitive use required longer time.

Together, these results indicate that participants favoured primitives corresponding to simpler hypotheses about the underlying templates in distinguishing steps, consistent with a

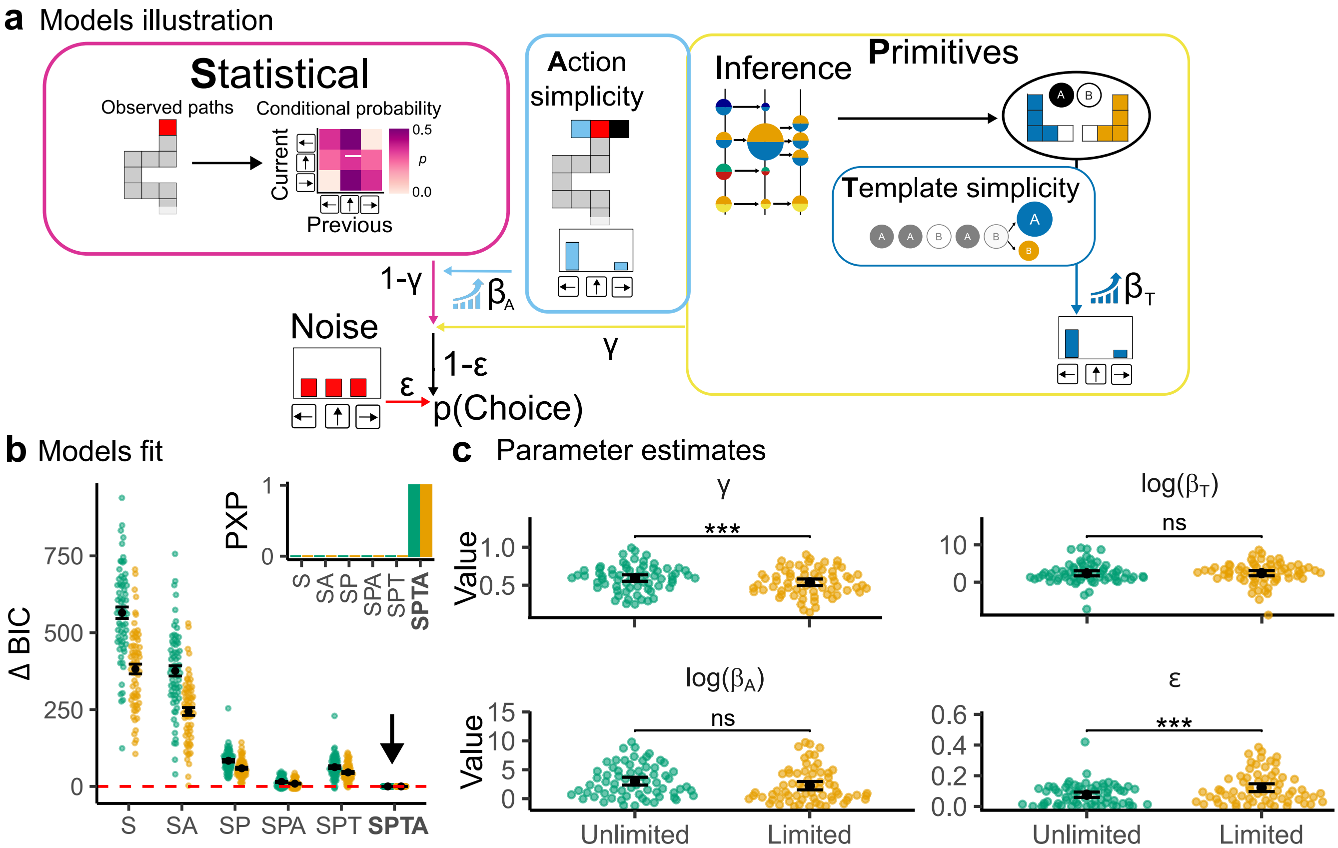

Figure 3. Computational modelling in Exp. 1. **a**) Computational models are defined compositionally: The statistical component (S) learns the conditional probability of actions; The action simplicity component (A) boosts the probability of actions based on the action simplicity of the resulting completion; The primitives acquisition inference (P) adds inference over the two candidate primitives in the current maze (using a particle filter; see Methods); The template simplicity component (T) biases primitive selection by boosting completions that result in higher template simplicity. **b**) Model comparison showing $\Delta\, BIC$ relative to the best-fitting model (SPTA). The inset shows the Protected Exceedance Probabilities (PXP). **c**) Parameter estimates for the winning SPTA model. Time pressure increased random errors (ε) and reduced primitive use (γ), but did not affect the bias towards action simplicity (β_A) nor template simplicity (β_T), both shown on a log-scale.

robust simplicity bias that was insensitive to rewards and unaffected by time pressure. In fact, participants with a stronger template bias were faster in choosing between primitives on distinguishing trials. In contrast, using primitives to choose actions required more time, and primitive use was significantly reduced under time pressure. Thus, while time pressure did not change the kind of compositional hypotheses participants generated, it diminished their effects on behaviour.

Exp. 2: Implicit learning of compositional structure

Exp. 2 used the same set of mazes as Exp. 1, but participants navigated them with unlimited time, were not told about the compositional structure of paths or the existence of primitives, and did not receive feedback about the primitive structure upon completing each path. This design allowed us to assess participants' spontaneous inductive biases about paths, their implicit discovery of primitives, and the factors influencing both.

Despite the lack of information, participants still achieved 72% accuracy on average (Fig. 4a), which was significantly higher than simulations based solely on marginal ($t(58) = 14.5$, $p < .001$, $d = 1.9$) or conditional probabilities ($t(58) = 12.2$, $p < .001$, $d = 1.6$). Instead, this accuracy was not significantly different from the one achieved by participants in Exp. 1 ($t(123) = -0.0$, $p = .966$, $d = 0.01$). This suggests they were able to discover and use the compositional structure of the task implicitly.

Using the same mixed-effects logistic regression as in Exp. 1 (Fig. 4b), we found that participants were more accurate in within-primitive vs. between-primitives steps (OR: 1.50 [1.44, 1.57], $p < .001$), with this gap increasing with round number (OR: 1.05 [1.00, 1.09], $p = .030$). This effect was primarily driven by within-primitive steps having increased accuracy over successive rounds (OR: 1.07 [1.04, 1.09], $p < .001$), whereas between-primitives accuracy was unaffected by round number (OR: 1.02 [0.99, 1.06], $p = .235$).

These results suggest that primitive use increased over rounds and raise the question of how participants combined them into templates.

Implicit, multi-scale simplicity biases

We tested the influence of different complexity measures on accuracy using the same mixed-effects logistic regressions including baseline regressors as in Exp. 1. Here, participants received no instruction about the computational structure, and their performance was best explained by action fragment simplicity (BIC: 61833, $p < .001$; Fig. 4c-d; see Fig. S4b). Fragment template simplicity was still the best template-level predictor (BIC: 61944, $p < .001$), and including it significantly improved the action simplicity model (BIC: 61826, $\chi^2(1) = 5.00$, $p = .025$).

Indeed, a more fine-grained completion analysis revealed that participants exhibited both action- and template-level simplicity biases, defined by fragment simplicity. Using the same completion analysis with distinguishing steps as in Exp. 1, here we found that both template simplicity (OR: 1.55 [1.44, 1.66], $p < .001$; Fig. 4e) and action simplicity (OR: 1.77 [1.61, 1.95], $p < .001$; Fig. 4f) predicted participant choices. These results control for the repetition of the last primitive (OR: 3.73 [3.29, 4.23], $p < .001$) and the choice of the correct action (OR: 1.23 [1.11, 1.38], $p < .001$) as baseline regressors in the same mixed-effects regression model (BIC: 13541; see Fig. S4e for details).

We then fit separate logistic regressions for each participant's trial-by-trial choices to estimate coefficients for their action- and template-level simplicity biases. This analysis revealed that action and template simplicity biases were negatively correlated across participants ($r = -.67$, $p < .001$; Fig. 4g), suggesting a tradeoff between action and template simplicity in hypotheses generation. As expected, the template simplicity bias for each participant scaled with their primitive use index ($b = 0.56$, [0.33, 0.79], $p < 0.001$; Fig. 4h), suggesting participants who relied more on primitives also exhibited a stronger template simplicity bias.

Here, individual differences can be partly explained by path dependence based on the order of mazes participants were randomly assigned. Exposure to mazes with high ground-truth template simplicity in the first half of rounds increased the template simplicity bias in the second half ($b = 1.16$, [1.07, 1.25], $p < .001$), with the effect remaining significant after binarizing early ground-truth template simplicity through a median split ($b = 1.24$, [1.06, 1.46], $p = .008$; Fig. 4i). Furthermore, early ground-truth template simplicity increased primitive use in the second half of the task ($b = 0.02$ [0.00, 0.03], $p = .030$; Fig. 4j), while the effect of early ground-truth action simplicity was not significant ($b = 0.01$ [-0.00, 0.03], $p = .065$). These results suggest that participants who experienced mazes with greater compositional structure (higher template simplicity) were more likely to adopt and use primitives later in the task.

Together, these findings suggest that participants initially exhibited a bias toward action-level simplicity, which was reoriented toward template-level simplicity, once they discovered and began using primitives.

Computational modelling of implicit simplicity biases

We fit the previously described models (Fig. 3a) on data from Exp. 2 (see Fig. S5a,b for model recovery). The full model SPTA again provided the best fit to participants' choices (mean BIC: 1049.50 ± 33.27; PXP ≈ 1.00; Fig. 4k), with likelihood-ratio tests confirming all components were necessary (all $p < .001$).

After confirming all SPTA parameters were reliably recoverable (Fig. S5c–f), we inspected the estimated parameters and computed the correlations with their corresponding behavioural indices (Fig. S6g). As in Exp. 1, we found that β_A positively correlated with accuracy ($r = .59$, $p < .001$), suggesting an action simplicity bias might be incentivized by the compositional structure of paths.

To further investigate how model parameters influenced performance, we fit a mixed-effects logistic regression (Fig. S7a). We found that accuracy decreased with ε (OR: 0.92 [0.90, 0.95], $p < .001$), but increased with γ (OR: 1.29 [1.25, 1.33], $p < .001$) and β_A (OR: 1.04 [1.01, 1.08], $p = .008$). However, the effect of β_T was not significant (OR: 1.00 [0.98, 1.03], $p = .847$). This dissociation reflects the fact that the compositional structure of the mazes promoted a bias towards action simplicity, but not template simplicity (Fig. S7b). While the templates spanned multiple levels of template simplicity, the repeated use of the same two primitives in each maze increased the ground-truth action simplicity of paths – making an action simplicity bias facilitate higher accuracy.

We next examined how model parameters modulated RTs in distinguishing trials. As in Exp. 1, we found that the template simplicity of chosen completions reduced RTs ($b = -0.07$, [-0.09, -0.06], $p < .001$), and this effect was stronger for participants with higher β_T ($b = -0.03$, [-0.05, 0.01], $p = .007$). When looking at all trials, RTs slowed down with β_A ($b = 0.17$, [0.04, 0.31], $p = .015$) and sped up with ε ($b = -0.19$, [-0.30, -0.08], $p < .001$), while the effects of β_T ($b = 0.07$, [-0.06, 0.20], $p = .282$) and γ were not significant ($b = 0.01$, [-0.10, 0.11], $p = .886$).

We then compared parameters of the SPTA model from Exp. 1 and Exp. 2 (Fig. 4l). To enable a direct comparison with Exp. 2, we restricted the analysis to parameter estimates from the unlimited time condition in Exp. 1. In Exp. 2, participants were not informed about the compositional structure of the task, nor were they explicitly signalled a primitive was completed, as in Exp. 1. In line with this difference, γ was lower in Exp. 2 ($t(123) = -3.3$, $p = .001$, $d = 0.6$). In contrast, neither β_T, β_A, nor ε differed reliably across experiments (β_T: $t(123) = -0.4$, $p = .716$, $d = 0.07$; β_A: $t(123) = -0.7$, $p = .483$, $d = 0.1$, ε: $t(123) = -1.4$, $p = .161$, $d = 0.3$; Fig. S5l).

Together, these results indicate that participants implicitly discovered and exploited primitives, even in the absence of explicit instruction. Withholding instructions about primitives made action simplicity the main driver of accuracy, but

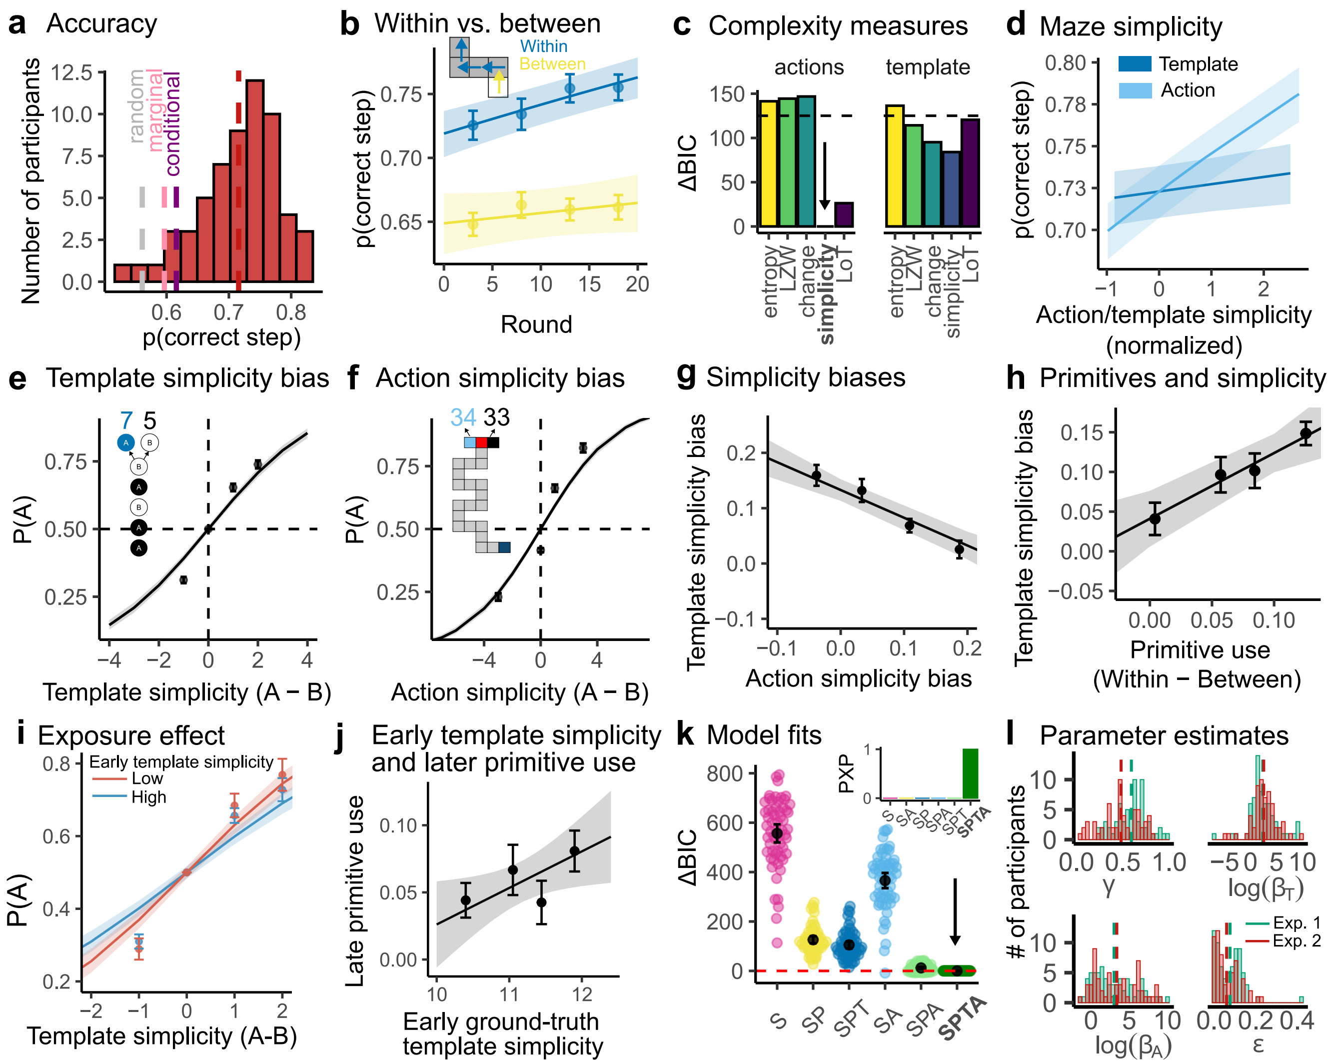

Figure 4. Experiment 2 results. **a**) Accuracy of participants compared to simulations relying only on action statistics. **b**) Accuracy on within-primitive steps increased over rounds, while performance on between-primitives steps did not, indicating that participants learned to exploit primitives. **c**) Comparison of complexity measures in predicting participant accuracy (mixed-effects logistic regression), illustrated using ΔBIC relative to the best-fitting model (action simplicity), where lower is better. The black dotted line is the BIC of a baseline model without any complexity measure, but including step and round number. **d**) The best-fitting model (action simplicity) was improved by including the best template complexity predictor (template simplicity). **e-f**) Completion analysis. The probability of choosing an action increased with the simplicity of its completion, computed both at the template-level (**e**) and action-level (**f**). **g**) Action and template simplicity biases were negatively correlated across participants. **h**) Template simplicity bias increased with primitive use across participants. **i**) Ground-truth template simplicity in the first half of rounds predicted greater template simplicity bias in the second half. **j**) Ground-truth template simplicity in the first half of rounds predicted greater primitive use in the second half. **k**) Model comparison showing the $\Delta\,BIC$ relative to the winning model (SPTA), with the inset showing PXP. **l**) Comparison of SPTA parameter estimates in Exp. 1 and 2 (with log-transformed β_T and β_A).

template simplicity still influenced participants' choices. Furthermore, participants who relied more on primitives also had a stronger bias towards hypotheses with greater template simplicity. Conversely, exposure to simpler templates was associated with both stronger template simplicity biases and greater reliance on primitives. Although withholding explicit information about primitives reduced their use, it did not reliably attenuate simplicity biases, consistent with the view that such biases arise spontaneously. Although the simplicity bias did not increase reward, it made responses faster, suggesting that simplicity bias is a useful heuristic for hypotheses generation.

Exp. 3: Simplicity biases in the absence of compositional structure

In Exp. 1 and 2, we found evidence participants generated compositional hypotheses that reused higher-level primitives, even when cognitive resources were limited or when the compositional structure was not explicitly introduced. In both cases, participants exhibited spontaneous biases towards hypotheses with greater template and action simplicity. This result suggests people are characterized by simplicity biases operating at multiple scales of learning.

However, action simplicity might reflect an adaptation to the compositional structure of the mazes, due to the repetition of actions occurring in primitives. Therefore, an action simplicity bias may reflect an environmental adaptation to the true underlying compositional structure of the mazes. To rule out this possibility, we designed a final experiment devoid of any compositional structure. To do so, we generated paths by randomly sampling and concatenating actions (*left*, *right* and *up*) to match the marginal and conditional probabilities from the set of mazes used in Exp. 1 and 2. Thus, Exp. 3 was matched in zeroth- and first-order action statistics, but lacked compositional structure.

In this case, participants achieved 58% accuracy on average (Fig. 5a), which was below the accuracy of baseline simulations using the true marginal ($t(57) = -2.6$, $p = .012$, $d = 0.3$) or conditional probability ($t(57) = -6.2$, $p < .001$, $d = 0.8$), but still better than random chance ($t(57) = 2.7$, $p = .009$, $d = 0.4$).

To assess which complexity measure best predicted behaviour, we compared mixed-effects logistic regressions using only action-level measures since there were no primitives to compute templates (see Fig. S5a–c for details). Here, the best-fitting model used action simplicity (BIC: 83952; Fig. 5b), which predicted higher performance (based on random variation across generated paths; OR: 1.06 [1.04, 1.08], $p < .001$; Fig. 5c). This model also included the effect of step number, which increased performance (OR: 1.06 [1.05, 1.08], $p < .001$) and had a negative interaction with action simplicity (OR: 0.98 [0.97, 1.00], $p = .020$). This interaction may suggest that participants partially reduced their simplicity bias after noticing it did not consistently improve performance within a round.

We confirmed these findings by examining step-by-step choices with a variant of the completion analysis applied at the level of actions (Fig. 5d left; see Fig. S5d for details). These results show that choices were biased by action simplicity (OR: 2.33 [2.25, 2.42], $p < .001$; Fig. 5d right), even when controlling for repetition of the last action (OR: 2.61, [2.48, 2.76], $p < .001$) and choice of the correct action (OR: 1.71 [1.62, 1.81], $p < .001$). Additionally, this model also revealed that action simplicity interacted with various factors, with its effects reducing over successive steps (OR: 0.88 [0.86, 0.90], $p < .001$) and slightly increasing with round number (OR: 1.01 [1.01, 1.01], $p < .001$). Thus, participants reduced their simplicity bias within each maze upon observing that it did not predict correct choices. However, exposure to more rounds of mazes without compositional structure did not override the bias, suggesting participants continued to expect mazes with simple action sequences. Nonetheless, early exposure to mazes high in action simplicity increased the subsequent action simplicity bias ($b = 1.03$, [1.01, 1.05], $p = .010$), even when early action simplicity was binarized ($b = 1.14$, [1.09, 1.19], $p < .001$; Fig. 5e). Thus, the action simplicity bias was at least in part sensitive to the ground-truth action simplicity.

Finally, a model comparison (Fig. 5f left) revealed that incorporating a simplicity bias SA (mean BIC: 1885.89 ± 8.20) significantly improved fit over the baseline statistical model S (mean BIC: 1937.75 ± 8.60; $\chi^2(58) = 3412.4$, $p < .001$; PXP$\approx$ 1.00; see Fig. S9a for parameter estimates). Again, we used simulations to ensure models (Fig. S9b,c) and their parameters (Fig. S9d,e) were recoverable.

We then compared β_A across the three experiments (Fig. 5g), and found that in Exp. 3 it was significantly lower than in Exp. 1 (two-sample t-test: $t(122) = 5.6$, $p < .001$, $d = 1.0$) and in Exp. 2 ($t(115) = 5.9$, $p < .001$, $d = 1.1$). In contrast with the previous experiments, β_A and action simplicity biases were not correlated with accuracy across participants (respectively: $\rho = -0.13$, $p = .342$; $\rho = -0.17$, $p = .209$; Fig. 5h; see Fig. S9f for partial correlations). As converging evidence, a mixed-effects logistic regression also did not find any effect of β_A on accuracy (OR: 1.00 [0.94, 1.05], $p = .936$; Fig. S9g). Therefore, the action simplicity bias neither helped nor hindered performance.

Model parameters did not significantly influence RTs on their own (Fig. S9h). However, β_A interacted significantly with action simplicity ($b = -0.02$, [-0.02, -0.01], $p < .001$; Fig. S9i), such that participants with higher β_A were faster when choosing actions leading to higher action simplicity. In contrast to previous experiments, here β_A sped up responses. This is because here, β_A only indexed the action simplicity bias, whereas in previous experiments β_A was correlated with primitive use (Fig. S6g) and correct choices (Fig. S7b).

Altogether, these findings indicate that the completion analysis and the SA model yield closely aligned evidence for an action simplicity bias. Importantly, the lack of significant correlation between β_A and accuracy indicates that action simplicity bias was not driven by task incentives in the maze paths. In sum, these results show that simplicity biases are integral to hypothesis generation, even in the absence of compositional structure and rewards incentivizing it.

Discussion

Compositional representations are a hallmark of human cognition[4], conferring important advantages such as systematic generalization[52–54] and efficient compression[20,34]. Across many domains, people display a robust inductive bias for simpler hypotheses[14,15], which has been argued to reflect a search for simple programs or language-like expressions defined over domain-specific primitives[10,25]. Here, we investigated the factors influencing this tendency towards simplicity in the

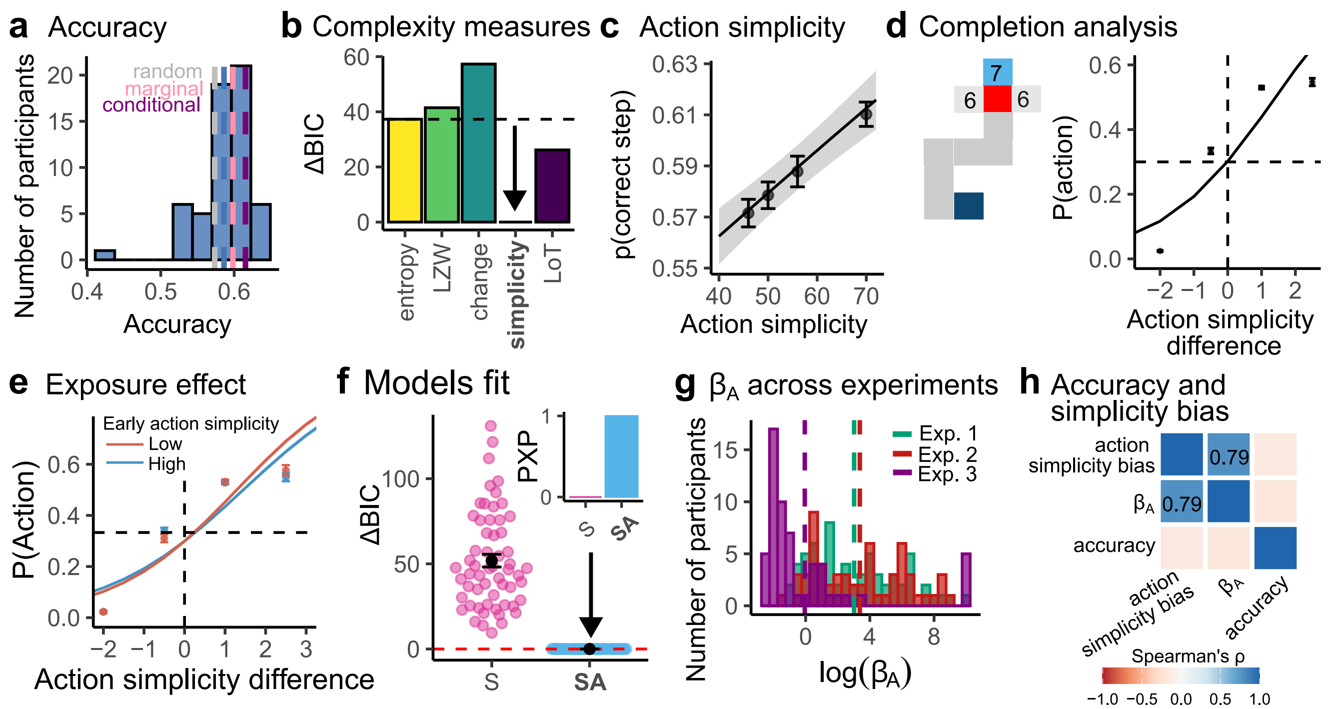

Figure 5. Experiment 3 results. **a**) Participants' accuracy compared to simulations relying only on action statistics. **b**) Model comparison (ΔBIC) for logistic regressions predicting accuracy using different complexity measures (action simplicity only due to lack of primitives). **c**) In the best-fitting regression, accuracy increased with action simplicity. **d**) Completion analysis: the probability of choosing an action increased with the simplicity of the resulting subpath, even when controlling for baseline regressors (e.g., marginal and conditional probabilities; left panel). **e**) Higher ground-truth action simplicity in the first half of the rounds predicted a stronger action simplicity in the second. **f**) Model comparison (ΔBIC) showing the SA model outperforms the other. **g**) Comparison of β_A, which control the action simplicity bias, across the three experiments. **h**) The fitted β_A correlates with the behavioural action simplicity bias but not with accuracy across participants.

generation of compositional hypotheses, and its interplay with the acquisition of primitives.

Across three experiments, participants navigated mazes with hidden paths. In the first two, paths were generated by combining spatial primitives of four actions (e.g., left, left, up, up) into higher-level templates (e.g., AABB...A; Fig. 1). In Exp. 1, we used a within-subject time pressure manipulation to test whether reduced cognitive resources influenced the discovery and use of compositional structure, in a setting where primitives, which consist of multiple actions, need to be learned (Fig. 2). Participants achieved higher accuracy when executing primitives than in choosing between them (as initial evidence of primitive use), and generally performed better on mazes with greater ground-truth template simplicity. To provide more fine-grained evidence, we used a completion analysis on distinguishing trials to show that participants' choices also favoured primitives corresponding to greater template simplicity. Time pressure did not reduce this template simplicity bias, but it did reduce primitive use. Crucially, computational modelling (Fig. 3) confirmed that these effects were not reducible to increased randomness, and persisted even when controlling for reduced decision precision under time pressure. Accordingly, our model-derived template simplicity bias was associated with faster responses, whereas primitive use slowed them down. Thus, converging evidence suggests that time pressure did not change the kind of compositional hypotheses participants generated, particularly when choosing between primitives. However, time pressure did diminish the contribution of simplicity biases to behaviour in favour of simpler, non-compositonal forms of statistical learning.

In Exp. 2, a different set of participants navigated the same mazes as in Exp. 1, but were not informed about the underlying spatial primitives, allowing us to examine their implicit acquisition (Fig. 4). We again found evidence for primitive use. Accuracy regressions and the completion analyses revealed participants exhibited both action- and template simplicity biases, which were negatively correlated. This raises the possibility that participants had a spontaneous bias towards action simplicity which was then reoriented at the level of templates due to the discovery of primitives. Individual differences between participants revealed that exposure to simpler templates in the first half of the experiment predicted both greater subsequent primitive use and template simplicity bias. However, our maze paths did not allow us to confirm the action simplicity bias was not induced by their underlying structure.

Finally, Exp. 3 tested whether the action simplicity bias was merely induced by the structure of the task (Fig. 5). Participants were presented with mazes devoid of any compositional structure, with paths generated to match the conditional probabilities of actions in Exps. 1–2. Strikingly, we still observed that action simplicity predicted accuracy on each maze and participants' trial-by-trial choices. Although this bias showed partial adaptation to the ground-truth action simplicity of paths in the first half of the mazes, it remained robust and corresponded to faster RTs, suggesting that the simplicity bias was not incentivized by the task, but still sped up hypothesis generation.

Together, our results shed light on the mechanisms by which compositional hypotheses are generated and selected. Commonly, the use of compositional representations is assumed to depend on some form of Bayesian inference, whereby prior knowledge and observations are integrated to evaluate hypotheses about the best underlying LoT expression or program[25,27]. In these accounts, hypotheses are favored according to their minimal description length (MDL;[14], i.e. the length of their shortest descriptions), introducing a bias toward simplicity. For instance, in visual and auditory domains, participants expect primitives to combine according to abstract templates whose learnability is determined by the complexity of their description in a LoT (e.g.[20]). Yet, generating hypotheses to search for the simplest hypothesis consistent with the data requires probing a combinatorial and potentially infinite hypothesis space[55], posing important computational challenges. While a simplicity bias could in principle help cope with these limitations by favouring compact representations, identifying them is itself computationally demanding. Rather than exhaustively searching for these globally minimal descriptions, people may rely on biased search heuristics that naturally favor compact representations.

In line with this possibility, we found that accuracy was best predicted by a measure we call *fragment simplicity*. Instead of presupposing that participants identify a single globally consistent expression minimizing MDL, our measure, which was inspired by past work by Alexander and Carey[16], counts the number of repetitions and symmetries in fragments across multiple scales. While remaining straightforward, this approach quantifies the amount of structure that can be reused to represent a sequence, even if it is not described by a unified, well-defined expression.

A key difference between ours and prior studies is that the latter repeatedly exposed participants to the same sequence, potentially enabling the search for a best-fitting LoT expression[20,31,56]. In contrast, we asked participants to predict the next step of novel partial sequences, making it more difficult for them to converge on a single, globally optimal hypothesis. Thus, in the context of our experiments, simplicity biases might not arise from an exhaustive, computationally costly search for the simplest hypothesis. Rather, they might instead reflect the use of efficient heuristics to generate compositional hypotheses, by creatively reusing fragments of them. Consistent with this interpretation, we found that across the three experiments, simplicity biases were associated with faster RTs and were not disrupted by resource limitations imposed through time pressure in Exp. 1.

A well-known strategy to alleviate the computational burden of search is to coarse-grain the hypothesis space by caching useful abstractions[34,42]. Indeed, previous studies have shown that increasing the frequency of certain combinations of primitives increases the likelihood that participants will use those combinations to parse new, ambiguous stimuli[45–47]. While these findings might suggest participants integrate frequent combinations into their primitives repertoire, they might also reflect local, iterative adjustments to prior hypotheses[55,57,58].

To measure the extent participants genuinely acquired and used novel primitives, we relied on both a behavioural index (i.e., increased accuracy on within- vs. between-primitives actions) and the parameter γ in our computational models (which weighted the use of primitives over mere statistical learning). Both of these measures revealed that time pressure significantly decreased primitive use, with the model also controlling for the increased random errors. Furthermore, higher γ was associated with longer RTs, suggesting that primitive use required additional time and cognitive resources. Consistently, previous studies have shown that higher-level sequence processing is not automatic[59,60], raising the possibility that resource limitations may hinder the use of newly learned primitives. Therefore, caching new primitives may help cope with limited resources only after these primitives have been well practiced.

Furthermore, our behavioural primitive use index enabled us to investigate the factors influencing the emergence of primitives. We found that participants who experienced mazes with higher template simplicity in the first half of Exp. 2 were more likely to use primitives in the second half. This suggests that the adoption of novel primitives is not only determined by the associative strength of their components[61,62]. One possibility is that participants are more likely to choose primitives if they enable them to generate simpler hypotheses. Thus, preferences for simpler hypotheses may promote the acquisition of new primitives. This highlights the potential interplay between two types of reuse, whereby the reuse of hypothesis fragments drives the incorporation of stable primitives into a library.

Collectively, our findings dovetail with recent work showing that the search for programs can be made more efficient and aligned to that of humans by using heuristics to transform previously inferred or learned ones[39]. One such heuristic is reuse, whereby new programs are generated by recombining components of previously learned ones[34,40,63]. For instance, a recent study asked participants to write programs to make a robot navigate an arena by reusing on-the-fly fragments of previous ones[40], finding that participants reused program subroutines. However, in their case, this reuse heuristic might have been influenced by the incentive structure of their task, which explicitly rewarded participants for using fewer primitives. Here, in the absence of explicit incentives for reuse, our robust simplicity bias further suggests there may be an intrinsic tendency to reuse fragments of previously considered hypotheses. Moreover, the fact that fragment simplicity mechanistically also operates over mirrored symmetries of past fragments suggests a tendency towards reuse is not only restricted to exact repetition, but may also allow for reuse with transformation. Thus, compositional hypothesis generation under cognitive constraints may rely on a broader form of creative reuse, in which fragments from previous hypotheses are flexibly adapted to construct new ones.

Limitations and future directions

Our results suggest that simplicity biases do not necessarily stem from an exhaustive, computationally costly search and evaluation of hypotheses, but instead reflect efficient heuristics for fragment reuse in hypothesis generation (Fig. 1d). However, in the present work we modeled simplicity through a scoring function that takes a sequence as input and returns a simplicity score. Future research should develop process-level models of hypothesis generation to uncover the mechanisms underlying this bias, for instance by explicitly sampling fragments from previous hypotheses and transforming them to generate novel hypotheses.

Although we tested simplicity bias in a more realistic setting where primitives consist of multiple actions that participants must learn, we did not model the exact mechanism by which participants acquired them. Future work could study the mechanisms that stitch together individual actions into high-level primitives by examining regular periodicity in RTs, which are a common signature of chunking[64]. Furthermore, although we found that time pressure reduced the use of primitives, our results do not rule out the possibility that primitives are learned in a more robust format over longer timescales, making their execution to more automatic and cheap. For instance, hippocampal replay has been proposed as a mechanism for building up and consolidate a library of primitives for later use[65,66]. Neuroimaging methods such as MEG or fMRI could provide a window into neural replay[67] and thus reveal its role in consolidating and combining primitives.

More generally, neuroimaging methods could help to understand the origin of the simplicity bias. Our current results suggest that the simplicity bias is an efficient heuristic for hypothesis construction, rather than simply a scoring metric. However, our behavioural data cannot distinguish between whether the bias arises during hypothesis generation or hypothesis selection, or both. Future work is thus required to better understand these mechanisms.

Conclusion

The use of compositional representations has long been linked to a preference for simple hypotheses[10,21]. This pervasive simplicity bias in human cognition[14,15,17] is often seen as reflecting a scoring process over candidate hypotheses to identify those best fitting observations. Yet this view exposes a tension: simple hypotheses afford compression but demand a computationally costly search. Rather than the "finite means"[11] of human cognition acting as an obstacle to compositional reasoning, our results suggest that resource limitations do not impair the operation of heuristics such as simplicity, which reduce computational cost, compress representations, and perhaps, enable flexible and efficient (albeit sometimes suboptimal) compositional generalization.

Methods

Participants and design

All participants were recruited through Prolific. The studies were approved by the Ethics in Psychological Research Commission of the University of Tübingen (Wu_2021/0124/213) and informed consent was obtained from all subjects.

Exp. 1 used a within-subject design manipulating time pressure. Half of the rounds had *unlimited time*, with participants free to complete the maze at their own pace. The other half of the rounds were *limited time* and participants had only 20 seconds to reach the goal. If the time expired before they reached the goal row, the round ended and they received no points. However, if they were able to complete it in time, their performance bonus on that round would be doubled. This was to ensure participants were motivated to engage with the limited time rounds, even though they were more demanding and less likely to be completed. For Exp. 1, we recruited 66 participants (33 female; M_{age}=38.00; *SD*=10.95) and paid them £3.75 for participation, plus a performance contingent bonus of up to £3.75; they spent 29.0 $\pm$ 10.68 minutes on the task and earned £4.71 $\pm$ 0.59 in total on average.

For Exp. 2, we recruited 59 participants (30 female; M_{age}=37.32; *SD*=13.23) and compensated them with £2.80 for taking part in the experiment and a performance contingent bonus of up to £2.80. Participants spent 15.83 $\pm$ 0.72 minutes on the task and earned £4.32 $\pm$ 0.05 in total.

For Exp. 3, we recruited 60 participants (M_{age}=37.2; *SD*=12.02), who received a base fee of £2.80 and, in addition, a performance contingent bonus of up to £2.80. Participants spent 9.76 $\pm$ 14 minutes on the task and earned £4.32 $\pm$ 0.05 in total.

Materials

We defined six spatial primitives, each consisting of four actions (either *left*, *right*, or *up*, hereafter *l*, *r*, and *u* respectively; Fig. 1c). Furthermore, we randomly generated 10 templates consisting of binary algebraic patterns (e.g. ABABABAB). For each template, we drew two pairs of primitives, generating 20 paths in total. Participants experienced each path in time-unlimited and time-limited conditions (40 rounds in total). We randomly sampled possible pairs of primitives to achieve the same frequency for mirrored primitives (e.g. *lulu* and *ruru*; Fig. S1a) and to make each primitive appear in combination with at least two others (Fig. S1b). Also, the sample of sequences balanced the marginal (Fig. S1c) and conditional

(Fig. S1d) probabilities of l and r actions. Furthermore, we also constrained each path to avoid backtracking and tight loops (e.g. *luru*), to make paths more visually distinguishable.

Participants navigated mazes constructed on a 17 row by 33 column grid world. Each maze had a single hidden solution, or path (Fig. 1a). For Exp. 1 and 2, we generated each path by combining two primitives, according to an abstract template (Fig. 1c). We defined three pairs of symmetric primitives (six in total). Each primitive consisted of four actions (either *left*, *right* or *up*). We chose a limited set of primitives to facilitate their acquisition and to ensure that each primitive could be combined with at least two others, preventing participants from inferring the second from the first one within each maze.

The templates were binary algebraic patterns (e.g. *ABABABAB*, where A and B denote the first and second primitive of each path, respectively). Thus, each template defined a sequence of eight primitives, resulting in paths of 32 actions. We constructed ten such templates, approximately balancing the number of occurrences of A and B and spanning a range of compositional complexity. For each template, we chose two pairs of primitives, yielding 20 paths in total. Paths were used only if they (i) reached the upper goal row and (ii) did that only on the last step, (iii) remained within the grid boundaries, and (iv) contained no loops. Furthermore, we filtered the paths to balance the marginal and conditional probability of *left* and *right* actions. In Experiment 1, each sequence was presented twice: once under unlimited time and once under time pressure.

For Exp. 3, we randomly generated valid paths (using the same criteria as for Exp. 1 and 2) and sampled 20 of them. We selected paths to match the marginal and conditional probabilities of actions from Exp. 1 and 2. Additionally, we only included paths that overlapped on no more than 22 steps (70% of actions), to ensure sufficient variability across mazes.

Procedure

After providing informed consent, participants were shown instructions, and completed an interactive tutorial and comprehension check before starting the experiment. The task consisted of 40 rounds for Exp. 1, and 20 for Exp. 2 and 3. In Exp. 1 we first taught participants about the generative process underlying path generation in a tutorial.

Each round began from the centre tile on the bottom row, and the goal was to reach the top row coloured in green (Fig. 1a). On each step, participants chose an action using the left, right and up arrow keys to take a step. The experiment ignored *invalid* choices, defined as steps outside the boundaries of the grid or "backtracking" (i.e., selecting left after a correct right move, and vice versa). We also excluded invalid actions from models.

After each valid choice, participants received feedback, with successful steps marked in blue and incorrect steps displayed in red. Incorrect actions resulted in the loss of a life and did not make the participant move from the current tile. Participants began each round with 20 lives in Exp. 1, and 30 lives in Exp. 2 and 3. They were incentivized to complete each maze with as many lives left as possible, as this determined their bonus. In Exp. 1, participants also completed time-limited rounds, with 20 seconds to reach the goal row. Points earned in these rounds were doubled, to motivate participants to engage seriously with the more difficult condition. After completing all rounds, participants provided demographic information and described the strategies they used to solve the task.

Complexity measures

We introduce a complexity measure called *fragment simplicity*, inspired by the work by Alexander and Carey[16] (Fig. 1d). The *fragment simplicity* of a string is the number of repetitions or subsymmetries occurring in all its substrings of lengths from 2 through n, where n is the length of the string. The simplicity of a maze can be defined at both the template and the action level. In line with previous studies (e.g.[20]), we compare simplicity with a number of other complexity measures. We estimate the entropy[48] of a template T by computing the probability of A and B within T and then summing $\sum_{i=1}^{n} -\log P(T(i))$, where $T(i)$ denotes the primitive at step i in the template; action-level entropy is computed analogously for individual paths, based on the probabilities of the actions they contain. Lempel-Ziv-Welch (LZW) complexity[49] measures the compressibility of templates and paths with the LZW algorithm. Change complexity[19] counts changes between consecutive substrings, weighted by the size of each. Lastly, LoT complexity[20] measures the complexity of a template or a path by the length of a minimal LoT expression describing it.

Behavioural analyses

We used mixed-effects logistic regressions to predict the outcome of each step (i.e., whether feedback was correct or incorrect). Candidate predictors, such as complexity measures, were evaluated by comparing them to a baseline model. Improvements in model fit were quantified using likelihood ratio tests (χ^2) on nested models.

For the completion analyses, we used logistic regressions to predict action selection based on competing strategies. We only considered the first attempt participants made for each step. We used a fixed-effects logistic regression because the randomized labeling of response options eliminated stable participant-level intercept differences, causing mixed-effects models with random intercepts to yield singular fits. For Exp. 1 and 2, for each distinguishing step we only considered the two actions corresponding to two different primitives. For Exp. 3, we predicted the action participants chose out of three possible alternatives; to take into account the path structure within a logistic regression, we defined the simplicity regressor for each action as the average difference between its completion simplicity and that of the two alternatives. All regressors were mean-centred and scaled to improve convergence.

Model-based analyses

We fit computational models predicting choices of participants in all steps (Fig. 3a). We compare a baseline model relying solely on the learned statistics of the task with models equipped with simplicity biases and understanding of the compositional structure underlying paths. The models induce probability distributions over actions, which are used to evaluate the negative log-likelihood (nLL) of each participant's choice. The baseline statistical model (S) learns conditional probabilities $p(\text{action} \mid \text{prevAction})$ of actions from their observed frequencies. The frequencies are initialized with a uniform Dirichlet prior (implemented with a pseudocount of 1) and updated after each observation. We tested another variant of the model including a parameter governing the contribution of conditional relative to marginal probabilities, but we found that participants almost unanimously only relied on the former and that this additional free parameter did not significantly improve the fit.

We also use a *compositional* model (SP), which chooses actions based on sampling hypotheses about the two primitives of each maze. The model uses a particle filter, where each particle ($n = 100$ for generality) is a possible pair of primitives. These hypotheses are resampled on each between-primitives step; the feedback is used to remove invalidated hypotheses. On each step, we compute a distribution over actions based on the primitives hypotheses and the current step number. This primitives-informed distribution over actions π_P is combined with π_S and weighted by the free parameter γ, which controls the contribution of primitives:

$$\pi_{SP} = (1-\gamma)\pi_S + \gamma\pi_P \tag{1}$$

The SP model learns the probability that two primitives co-occur within a maze. As for conditional probability of actions, these probabilities are initialized with a uniform Dirichlet prior (implemented via a pseudocount of 1) and updated whenever a pair of primitives is unambiguously observed. The SP model then samples hypothesized primitive pairs from the resulting posterior distribution, subject to consistency with previous observations.

The template simplicity model (SPT) biases the distribution over primitives, boosting the probability of primitives corresponding to simpler completions. Here, we compute simplicity for all primitives sim_T by subtracting the simplicity of the observed subtemplate from the simplicity of the candidate completion (obtained by concatenating the observed subtemplate to each candidate primitive). Here, the free parameter β_T determines the strength of the bias for template simplicity:

$$\pi_P \propto \pi_P + \beta_T \text{sim}_T \tag{2}$$

The resulting scores were renormalized to sum to 1, yielding a valid probability distribution over primitives.

We also equipped models with an action simplicity bias (SA, SPA, SPTA). The distribution over actions π_S is boosted by the action-level completion simplicity, with the free parameter β_A determining the strength of the bias for action simplicity.

$$\pi_S \propto \pi_S + \beta_A \text{sim}_A \tag{3}$$

Again, the resulting scores were renormalized to sum to 1, yielding a valid probability distribution over actions.

In SPA and SPT models, this distribution is then mixed with the distribution π_P weighted by the primitive weight parameter γ. We fitted parameters to minimize nLL, computing a maximum likelihood estimate for each model, and then compared them based on BIC. Nested models are compared with a log-likelihood ratio test (χ^2).

Acknowledgements

We are grateful for feedback from Philipp Schwartenbeck on early versions of the experiment. We thank Mani Hamidi for help running some of the analyses on a computing cluster. Preliminary analyses were previously presented at the 2023 meeting of The Cognitive Science Society[68], and at the 2023 Conference on Cognitive Computational Neuroscience[69].

Author contributions statement

V.R., P.D., and C.M.W. conceived the experiments, V.R. and C.M.W. conducted the experiments, V.R. and C.M.W. analysed the results. V.R. and C.M.W. wrote the first draft of the manuscript with feedback from P.D. All authors reviewed and approved the final manuscript.

Additional information

Data and code are publicly available at https://github.com/vrubino/mazeExperiment-public/ The authors declare no competing interests.

References

1. Boole, G. *An investigation of the laws of thought: On which are founded the mathematical theories of logic and probabilities* (Walton and Maberly, London, UK, 1854).
2. Simon, H. A. *The Sciences of the Artificial* (MIT Press, 1969).
3. Lake, B. M., Ullman, T. D., Tenenbaum, J. B. & Gershman, S. J. Building machines that learn and think like people. *Behav. brain sciences* **40** (2017).
4. Dehaene, S., Al Roumi, F., Lakretz, Y., Planton, S. & Sablé-Meyer, M. Symbols and mental programs: a hypothesis about human singularity. *Trends Cogn. Sci.* (2022).
5. Chomsky, N. *Cartesian linguistics: A chapter in the history of rationalist thought* (Harper & Row, 1966).

6. Kumar, S., Dasgupta, I., Daw, N. D., Cohen, J. D. & Griffiths, T. L. Disentangling abstraction from statistical pattern matching in human and machine learning. *PLOS Comput. Biol.* **19**, e1011316 (2023).
7. Morgan, T. J. & Feldman, M. W. Human culture is uniquely open-ended rather than uniquely cumulative. *Nat. Hum. Behav.* **9**, 28–42 (2025).
8. Sablé-Meyer, M., Ellis, K., Tenenbaum, J. & Dehaene, S. A language of thought for the mental representation of geometric shapes. *Cogn. Psychol.* **139**, 101527, DOI: 10.1016/j.cogpsych.2022.101527 (2022).
9. Fodor, J. A. & Pylyshyn, Z. W. Connectionism and cognitive architecture: A critical analysis. *Cognition* **28**, 3–71, DOI: 10.1016/0010-0277(88)90031-5 (1988).
10. Fodor, J. *The Language of Thought* (Harvard University Press, 1975).
11. Von Humboldt, W. *Über die Verschiedenheit des menschlichen Sprachbaues und ihren Einfluss auf die geistige Entwickelung des Menschengeschlechts* (Dümmler, Berlin, 1836).
12. Goodman, N. *Fact, Fiction, and Forecast* (Harvard University Press, 1955).
13. Goodman, N. Safety, strength, simplicity. *Philos. Sci.* **28**, 150–151, DOI: 10.1086/287795 (1961).
14. Chater, N. & Vitányi, P. Simplicity: a unifying principle in cognitive science? *Trends cognitive sciences* **7**, 19–22 (2003).
15. Feldman, J. The simplicity principle in perception and cognition. *Wiley Interdiscip. Rev. Cogn. Sci.* **7**, 330–340 (2016).
16. Alexander, C. & Carey, S. Subsymmetries. *Percept. & Psychophys.* **4**, 73–77 (1968).
17. Feldman, J. Minimization of boolean complexity in human concept learning. *Nature* **407**, 630–633 (2000).
18. Grünwald, P. D. *The minimum description length principle* (MIT press, 2007).
19. Aksentijevic, A., Mihailovic, A. & T. Mihailovic, D. Time for change: Implementation of aksentijevic-gibson complexity in psychology. *Symmetry* **12**, 948, DOI: 10.3390/sym12060948 (2020).
20. Planton, S. *et al.* A theory of memory for binary sequences: Evidence for a mental compression algorithm in humans. *PLoS computational biology* **17**, e1008598 (2021).
21. Simon, H. A. & Kotovsky, K. Human acquisition of concepts for sequential patterns. *Psychol. review* **70**, 534 (1963).
22. Wu, C. M., Meder, B. & Schulz, E. Unifying principles of generalization: past, present, and future. *Annu. Rev. Psychol.* **76**, 275–302, DOI: 10.1146/annurev-psych-021524-110810 (2025).
23. Piantadosi, S. T. The computational origin of representation. *Minds machines* **31**, 1–58 (2021).
24. Smolensky, P. *On the Proper Treatment of Connectionism*, 229–254 (MIT Press, 1993).
25. Piantadosi, S. T. & Jacobs, R. A. Four problems solved by the probabilistic language of thought. *Curr. Dir. Psychol. Sci.* **25**, 54–59, DOI: 10.1177/0963721415609581 (2016).
26. Ellis, K. *et al.* Dreamcoder: growing generalizable, interpretable knowledge with wake–sleep bayesian program learning. *Philos. Transactions Royal Soc. A* **381**, 20220050 (2023).
27. Goodman, N. D., Tenenbaum, J. B., Feldman, J. & Griffiths, T. L. A rational analysis of rule-based concept learning. *Cogn. science* **32**, 108–154 (2008).
28. Tenenbaum, J. B. & Griffiths, T. L. Generalization, similarity, and bayesian inference. *Behav. brain sciences* **24**, 629–640 (2001).
29. Kemp, C. & Tenenbaum, J. B. The discovery of structural form. *Proc. Natl. Acad. Sci.* **105**, 10687–10692 (2008).
30. Griffiths, T. L., Chater, N., Kemp, C., Perfors, A. & Tenenbaum, J. B. Probabilistic models of cognition: exploring representations and inductive biases. *Trends Cogn. Sci.* **14**, 357–364, DOI: 10.1016/j.tics.2010.05.004 (2010).
31. Amalric, M. *et al.* The language of geometry: Fast comprehension of geometrical primitives and rules in human adults and preschoolers. *PLoS computational biology* **13**, e1005273 (2017).
32. Zhou, Y. & Lake, B. M. Flexible compositional learning of structured visual concepts. *arXiv preprint arXiv:2105.09848* (2021).
33. Mills, T., Tenenbaum, J. & Cheyette, S. Human spatiotemporal pattern learning as probabilistic program synthesis. *Adv. Neural Inf. Process. Syst.* **36**, 54354–54365 (2023).
34. Zhou, H., Nagy, D. G. & Wu, C. M. Harmonizing program induction with rate-distortion theory. *arXiv preprint arXiv:2405.05294* (2024).
35. Quilty-Dunn, J., Porot, N. & Mandelbaum, E. The best game in town: The reemergence of the language-of-thought hypothesis across the cognitive sciences. *Behav. Brain Sci.* **46**, e261 (2023).
36. Bhui, R., Lai, L. & Gershman, S. J. Resource-rational decision making. *Curr. Opin. Behav. Sci.* **41**, 15–21 (2021).
37. Lieder, F. & Griffiths, T. L. Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources. *Behav. brain sciences* **43** (2020).
38. Mathy, F., Friedman, O. & Gauvrit, N. Can compression take place in working memory without a central

contribution of long-term memory? *Mem. & Cogn.* **52**, 1726–1736, DOI: 10.3758/s13421-023-01474-8 (2023).

39. Rule, J. S. *et al.* Symbolic metaprogram search improves learning efficiency and explains rule learning in humans. *Nat. Commun.* **15**, DOI: 10.1038/s41467-024-50966-x (2024).

40. Correa, C. G. *et al.* Exploring the hierarchical structure of human plans via program generation. *Cognition* **255**, 105990 (2025).

41. Ullman, T. D., Goodman, N. D. & Tenenbaum, J. B. Theory learning as stochastic search in the language of thought. *Cogn. Dev.* **27**, 455–480, DOI: 10.1016/j.cogdev.2012.07.005 (2012).

42. Bramley, N. R., Zhao, B., Quillien, T. & Lucas, C. G. Local search and the evolution of world models. *Top. Cogn. Sci.* (2023).

43. Ellis, K., Morales, L., Sablé-Meyer, M., Solar-Lezama, A. & Tenenbaum, J. Learning libraries of subroutines for neurally–guided bayesian program induction. *Adv. Neural Inf. Process. Syst.* **31** (2018).

44. Bowers, M. *et al.* Top-down synthesis for library learning. *Proc. ACM on Program. Lang.* **7**, 1182–1213 (2023).

45. Cheyette, S. & Piantadosi, S. Knowledge transfer in a probabilistic language of thought. In *CogSci* (2017).

46. Tano, P., Romano, S., Sigman, M., Salles, A. & Figueira, S. Towards a more flexible language of thought: Bayesian grammar updates after each concept exposure. *Phys. Rev. E* **101**, DOI: 10.1103/physreve.101.042128 (2020).

47. Zhao, B., Lucas, C. G. & Bramley, N. R. A model of conceptual bootstrapping in human cognition. *Nat. Hum. Behav.* **8**, 125–136 (2024).

48. Shannon, C. E. A mathematical theory of communication. *The Bell system technical journal* **27**, 379–423 (1948).

49. Welch, T. A. A technique for high-performance data compression. *Computer* **17**, 8–19 (1984).

50. Speekenbrink, M. A tutorial on particle filters. *J. Math. Psychol.* **73**, 140–152 (2016).

51. Doucet, A. & Johansen, A. M. A tutorial on particle filtering and smoothing: Fifteen years later. *Handb. nonlinear filtering* **12**, 3 (2009).

52. Hupkes, D., Dankers, V., Mul, M. & Bruni, E. Compositionality decomposed: How do neural networks generalise? *J. Artif. Intell. Res.* **67**, 757–795 (2020).

53. Lake, B. M. Compositional generalization through meta sequence-to-sequence learning. *Adv. neural information processing systems* **32** (2019).

54. Dekker, R. B., Otto, F. & Summerfield, C. Curriculum learning for human compositional generalization. *Proc. Natl. Acad. Sci.* **119**, e2205582119 (2022).

55. Fränken, J.-P., Theodoropoulos, N. C. & Bramley, N. R. Algorithms of adaptation in inductive inference. *Cogn. Psychol.* **137**, 101506 (2022).

56. Al Roumi, F., Marti, S., Wang, L., Amalric, M. & Dehaene, S. Mental compression of spatial sequences in human working memory using numerical and geometrical primitives. *Neuron* **109**, 2627–2639.e4, DOI: 10.1016/j.neuron.2021.06.009 (2021).

57. Bramley, N. R., Dayan, P., Griffiths, T. L. & Lagnado, D. A. Formalizing neurath's ship: Approximate algorithms for online causal learning. *Psychol. review* **124**, 301 (2017).

58. Dasgupta, I., Schulz, E. & Gershman, S. J. Where do hypotheses come from? *Cogn. psychology* **96**, 1–25 (2017).

59. Bekinschtein, T. A. *et al.* Neural signature of the conscious processing of auditory regularities. *Proc. Natl. Acad. Sci.* **106**, 1672–1677 (2009).

60. Strauss, M. *et al.* Disruption of hierarchical predictive coding during sleep. *Proc. Natl. Acad. Sci.* **112**, E1353–E1362 (2015).

61. Wu, S., Éltető, N., Dasgupta, I. & Schulz, E. Chunking as a rational solution to the speed–accuracy trade-off in a serial reaction time task. *Sci. Reports* **13**, DOI: 10.1038/s41598-023-31500-3 (2023).

62. Éltető, N., Nemeth, D., Janacsek, K. & Dayan, P. Tracking human skill learning with a hierarchical bayesian sequence model. *PLOS Comput. Biol.* **18**, e1009866, DOI: 10.1371/journal.pcbi.1009866 (2022).

63. Nagy, D. G., Shen, T., Zhou, H., Wu, C. M. & Dayan, P. Analogy making as amortised model construction. In *Reinforcement Learning Conference: Finding the Frame Workshop*, DOI: 10.48550/arXiv.2507.16511 (2025).

64. Koch, I. & Hoffmann, J. Patterns, chunks, and hierarchies in serial reaction-time tasks. *Psychol. Res. Psychol. Forschung* **63**, 22–35, DOI: 10.1007/pl00008165 (2000).

65. Schwartenbeck, P. *et al.* Generative replay underlies compositional inference in the hippocampal-prefrontal circuit. *Cell* **186**, 4885–4897.e14, DOI: 10.1016/j.cell.2023.09.004 (2023).

66. Kurth-Nelson, Z. *et al.* Replay and compositional computation. *Neuron* (2023).

67. Schuck, N. W. & Niv, Y. Sequential replay of nonspatial task states in the human hippocampus. *Science* **364**, eaaw5181 (2019).

68. Rubino, V., Hamidi, M., Dayan, P. & Wu, C. M. Compositionality under time pressure. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, vol. 45 (2023).

69. Rubino, V., Dayan, P. & Wu, C. Biases towards compositionally simpler hypotheses are robust and unaffected by learning. In *Conference on Cognitive Computational Neuroscience (CCN 2023)*, 817–820 (2023).

Supplementary Information for:

Simplicity guides the discovery and use of compositionality

Valerio Rubino, Peter Dayan, Charley M. Wu

Exp. 1

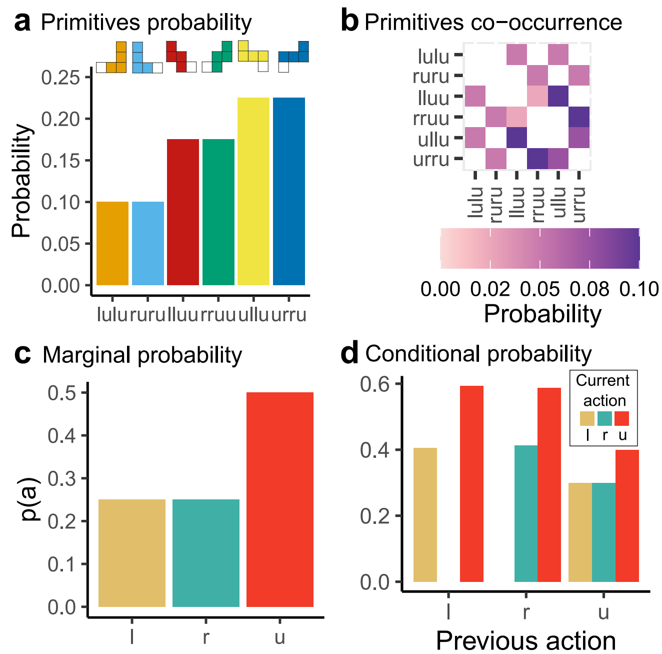

Figure S1. Task details. (**a**) Each of the 20 paths used in Exp. 1 and 2 were constructed by combining pairs of spatial primitives, selected from a set of six: *lulu*, *ruru*, *lluu*, *rruu*, *ullu*, and *urru*, where *l*, *r*, and *u* denote left, right, and up actions, respectively. We ensured that mirrored pairs, such as *lulu* and *ruru*, occurred equally often. (**b**) Each primitive appeared in combination with at least two others to emphasize the combinatorial structure of the task. (**c**) The final set of paths was selected to symmetrically balance the marginal probability of individual actions and (**d**) the conditional probability of actions given the previous action.

Logistic mixed-effects regressions of accuracy on complexity

We used logistic mixed-effects regressions to identify the complexity measure that accounted best for participants' accuracy. We considered entropy, LZW complexity, change complexity, LOT complexity and simplicity. We used these measures to compute the complexity of mazes, both at the template- and at the action-level. Because the measures were moderately correlated, we tested them in separate regression models to avoid multicollinearity (Figure S2c-d).

First, we fitted a logistic mixed-effects regression, specifying the fixed effect of time pressure and a random effect for participant (BIC = 125735). Adding the effect of round number significantly improved the fit (BIC = 125742; Likelihood ratio test against the model with only time pressure: $\chi^2(1) = 4.80$, $p = .028$). Adding the interaction between the two terms did not improve the fit (BIC = 125753 $\chi^2(1) = 0.01$, $p = .925$). Therefore, we used a baseline model including the negative effects of time pressure (OR: 0.80 [0.78, 0.82], $p = .001$) and round number (OR: 0.98 [0.97, 1.00], $p = .028$) and added complexity measures: actions entropy (BIC = 124250; $\chi^2(1) = 1.00$, $p = .336$), templates entropy (BIC = 124245; $\chi^2(1) = 5.15$, $p = .023$), actions LZW (BIC = 124238; $\chi^2(1) = 12.26$, $p < .001$), templates LZW (BIC = 124246; $\chi^2(1) = 4.80$, $p = .029$), actions change complexity (BIC = 124249; $\chi^2(1) = 1.11$, $p = .291$), templates change complexity (BIC = 124189; $\chi^2(1) = 61.16$, $p < .001$), actions LoT Complexity (BIC = 124168; $\chi^2(1) = 82.66$, $p < .001$), templates LoT Complexity (BIC = 124207; $\chi^2(1) = 39.10$, $p < .001$), actions simplicity (BIC = 124235; $\chi^2(1) = 15.78$, $p < .001$), templates simplicity (BIC = 124173; $\chi^2(1) = 76.63$, $p < .001$).

The models significantly improving the baseline model were enriched with an interaction term between the baseline regressors and the complexity measure: templates entropy (BIC = 124244; $\chi^2(3) = 36.38$, $p < .001$), actions LZW (BIC = 124243;

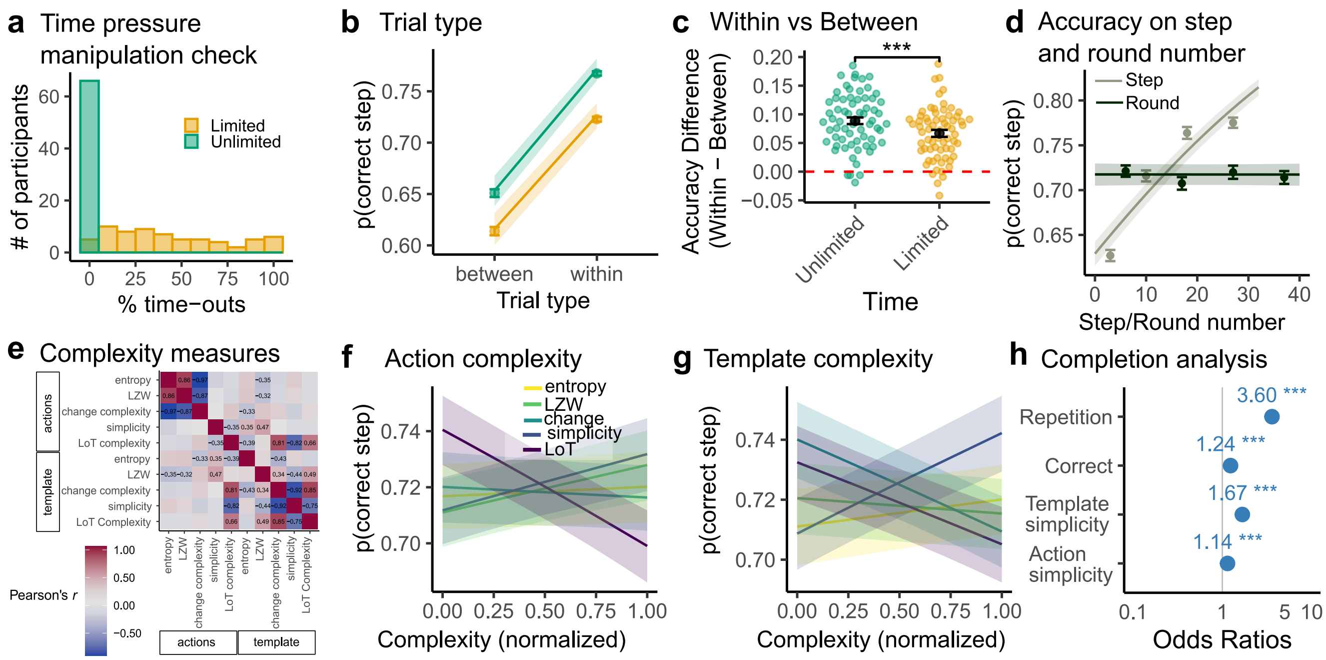

Figure S2. **Exp. 1 supplementary results**. (**a**) Proportion of rounds in which participants ran out of time. (**b**) Primitive use in steps where both primitives were already discovered. (**c**) Logistic regression predicting accuracy by step type and time pressure. (**d**) Baseline logistic regression predicting accuracy as a function of round and step number. (**e**) Pearson's r correlations among path and template complexity (only coefficients for significant correlations at $p < .05$ are reported). (**f-g**) Effects of maze complexity measures, defined either at the level of actions (**f**) or templates (**g**) (**h**) Completion analysis: main effects from a logistic regression predicting the probability of choosing an action.

$\chi^2(3) = 40.93$, $p < .001$), templates LZW (BIC = 124201; $\chi^2(3) = 79.47$, $p < .001$), templates change complexity (BIC = 123836; $\chi^2(3) = 388.13$, $p < .001$), actions LoT Complexity (BIC = 123843; $\chi^2(3) = 359.92$, $p < .001$), templates LoT Complexity (BIC = 124056; $\chi^2(3) = 186.27$, $p < .001$), actions simplicity (BIC = 124215; $\chi^2(3) = 54.88$, $p < .001$), templates simplicity (BIC = 123746; $\chi^2(3) = 463.13$, $p < .001$). For simplicity, we present BICs of interaction models in Figure 2c.

Overall, the model including template simplicity best predicted accuracy (BIC = 123746). In this model, accuracy decreased with round number (OR: 0.97, [0.95, 0.98], $p < .001$, perhaps due to a lack of mazes order randomization) and under time pressure (OR: 0.84 [0.82,0.86], $p < .001$) and it increased with both step number (OR: 1.32 [1.30, 1.34], $p < .001$) and template simplicity (OR: 1.10 [1.07, 1.12], $p < .001$). The effect of template simplicity increased with step number (OR: 1.16 [1.14, 1.17], $p < .001$) and decreased with round number (OR: 0.95 [0.94, 0.96], $p < .001$). Importantly, it did not interact with time pressure (OR: 1.02 [1.00, 1.05], $p = .101$). The significant interaction between template simplicity and round number likely resulted from the lack of counterbalancing in maze presentation order across participants, which we achieved with Exp. 2.

Step-by-step completion analysis

For the completion analysis, we focused on distinguishing steps. In a distinguishing step, each of two possible actions corresponded to a different primitive: for a maze with *ullu* and *urru* as primitives, after observing the partial path *ulluu*-, participants could choose to continue choosing *l* or *r*, reflecting their choice of the former or the latter primitive. We only used between-primitives steps and excluded steps where participants chose actions inconsistent with either primitive (in the example, participants chose *u*). Furthermore, we restricted the analysis to steps where the second primitive had been already discovered by participants (i.e., subtemplates containing at least an instance of the second primitive, i.e. *B*) to remove the effect of uncertainty on the other primitive. We only considered the first choice for each step. Each distinguishing step contributed two rows to the dataset, corresponding to the two admissible actions. The response variable indicated which action was chosen. We modeled participants' choices using a logistic regression. To eliminate any arbitrary labeling asymmetries, at each distinguishing step the two admissible actions (corresponding to primitives) were randomly assigned to response labels "A" and "B". The dependent variable therefore indicated whether the action labeled "A" was selected. Under this randomization scheme, the unconditional probability of choosing A is 0.5. Consequently, any systematic deviation from 0.5 reflects sensitivity to the predictors rather than stable side preferences.

Because labeling was randomized independently at each step, participants could not exhibit baseline tendencies toward either response label. Accordingly, including a participant-level random intercept resulted in singular fits. We therefore report fixed-effects logistic regressions.

As covariates, we included the marginal and conditional probability of each possible action as regressors. In addition, we categorized actions based on their associated primitives and identified whether each action was consistent with repeating the previous primitive, and the extent its associated primitive increased subtemplate simplicity. Lastly, we also included the contribution of each possible action to the subpath simplicity. For all predictors, we used the difference between the two alternatives within each distinguishing step. We report the estimated fixed effects in Exp. 1 in Figure S2h.

Model parameters and behavioural correlates

We examined Spearman correlations between behavioural measures and model parameters, separately by time pressure conditions.

In unlimited time rounds (Figure S3d), template simplicity bias was positively associated with γ ($\rho = 0.36$, $p < .01$), whereas action simplicity bias was not ($\rho = -0.01$, $p = .943$). Primitive use and accuracy were strongly correlated with γ ($\rho = 0.80$ and $\rho = 0.87$, respectively; both $p < .001$). Template simplicity bias was modestly correlated with β_T ($\rho = 0.25$, $p < .05$), while action simplicity bias ($\rho = -0.20$, $p = .105$), primitive use ($\rho = -0.12$, $p = .318$), and accuracy ($\rho = -0.04$, $p = .777$) were not. In contrast, primitive use and accuracy were positively associated with β_A ($\rho = 0.45$ and $\rho = 0.50$, respectively; both $p < .001$), whereas neither simplicity bias measure was. Finally, ε was not significantly correlated with any behavioural measure (all $|\rho| \leq 0.20$, all $p \geq .105$).

In limited time rounds (Figure S3e), again, primitive use and accuracy were positively associated with γ ($\rho = 0.52$ and $\rho = 0.80$, respectively; both $p < .001$), whereas neither template simplicity bias ($\rho = 0.16$, $p = .190$) nor action simplicity bias ($\rho = -0.12$, $p = .318$) was significantly related to γ. Neither simplicity bias measure, primitive use, nor accuracy was significantly associated with β_T (all $|\rho| \leq 0.07$, all $p \geq .582$). For β_A, accuracy was positively correlated ($\rho = 0.48$, $p < .001$), whereas primitive use ($\rho = 0.18$, $p = .143$) and both simplicity bias measures (all $|\rho| \leq 0.13$, all $p \geq .309$) were not. Finally, ε was not significantly correlated with any behavioural measure (all $|\rho| \leq 0.11$, all $p \geq .361$).

We fit a mixed-effects logistic regression predicting trial-level accuracy, including random intercepts for participants (Figure S3f). Adding interactions between model parameters and time pressure did not improve model fit ($\beta_T \times$ time pressure: $\chi^2(1) = 1.55$, $p = .213$; $\beta_A \times$ time pressure: $\chi^2(1) = 0.89$, $p = .347$; $\varepsilon \times$ time pressure: $\chi^2(1) = 0.31$, $p = .581$; $\gamma \times$ time pressure: $\chi^2(1) = 1.72$, $p = .199$).

We fit a mixed-effects linear regression predicting log-transformed RTs, including random intercepts for participants (Figure S4a). RTs increased with γ ($b = 0.18$, [0.16, 0.19], $p < .001$) and β_A ($b = 0.08$, [0.07, 0.09], $p < .001$), and decreased with β_T ($b = -0.02$, [-0.03, -0.01], $p = .001$), ε ($b = -0.13$, [-0.15, -0.12], $p < .001$), and limited time pressure ($b = -0.09$, [-0.10, -0.08], $p < .001$). The effect of γ was stronger under limited time pressure ($b = 0.11$, [0.10, 0.12], $p < .001$), as was the effect of β_T ($b = 0.06$, [0.05, 0.07], $p < .001$). In contrast, the effect of β_A was attenuated under time pressure ($b = -0.06$, [-0.07, -0.04], $p < .001$), and the negative effect of ε was reduced ($b = 0.03$, [0.02, 0.04], $p < .001$).

RTs in distinguishing steps

We used the same distinguishing trials from the completion analysis to investigate the influence of template and action simplicity biases on RTs.

We first fit a linear mixed-effects linear regression predicting log-transformed RTs on template simplicity of each completion, β_T of each participant and time pressure. We also included the interaction of β_T and completion template simplicity, hypothesizing that the effect of β_T would emerge only in steps where the simpler hypotheses is selected. Including the interaction of time pressure with β_T and template simplicity significantly improved the fit (respectively: $\chi^(1) = 25.37$, $p < .001$; $\chi^2(1) = 5.29$, $p = .021$). While including both interactions instead of one improved the fit ($\chi^2(1) = 4.58$, $p = .032$), the threeway interaction between the three predictions did not ($\chi^2(1) = 0.09$, $p = .0.769$). In the best model (Figure S4b, left), RTs decrease with template simplicity ($b = -0.08$, [-0.10, -0.06], $p < .001$), β_T ($b = -0.06$, [-0.10, -0.03], $p < .001$) and time pressure ($b = -0.37$, [-0.40, -0.34], $p < .001$). We also found a negative interaction between template simplicity and β_T ($b = -0.02$, [-0.04, -0.01], $p = .004$). In contrast, we found positive interactions between time pressure and template simplicity ($b = 0.03$, [0.00,0.06], $p = .032$)and between time pressure and β_T ($b = 0.09$, $[0.05, 0.13]$, $p < .001$).

We conducted an analogous analysis at the action level, predicting log-transformed RTs from action simplicity, participants' β_A, and time pressure. Including the interaction between action simplicity and β_A significantly improved model fit ($\chi^2(1) = 6.61$, $p = .010$), whereas adding the interaction between β_A and time pressure did not ($\chi^2(1) = 0.01$, $p = .935$). The interaction between action simplicity and time pressure did not yield a significant improvement in fit ($\chi^2(1) = 3.71$, $p = .054$). In the best-fitting model (Figure S4b, right), RTs decreased with action simplicity ($b = -0.06$, $[-0.07, -0.04]$, $p < .001$) and time pressure ($b = -0.36$, $[-0.39, -0.33]$, $p < .001$), but increased with β_A ($b = 0.05$, $[0.02, 0.07]$, $p < .001$). Critically, the

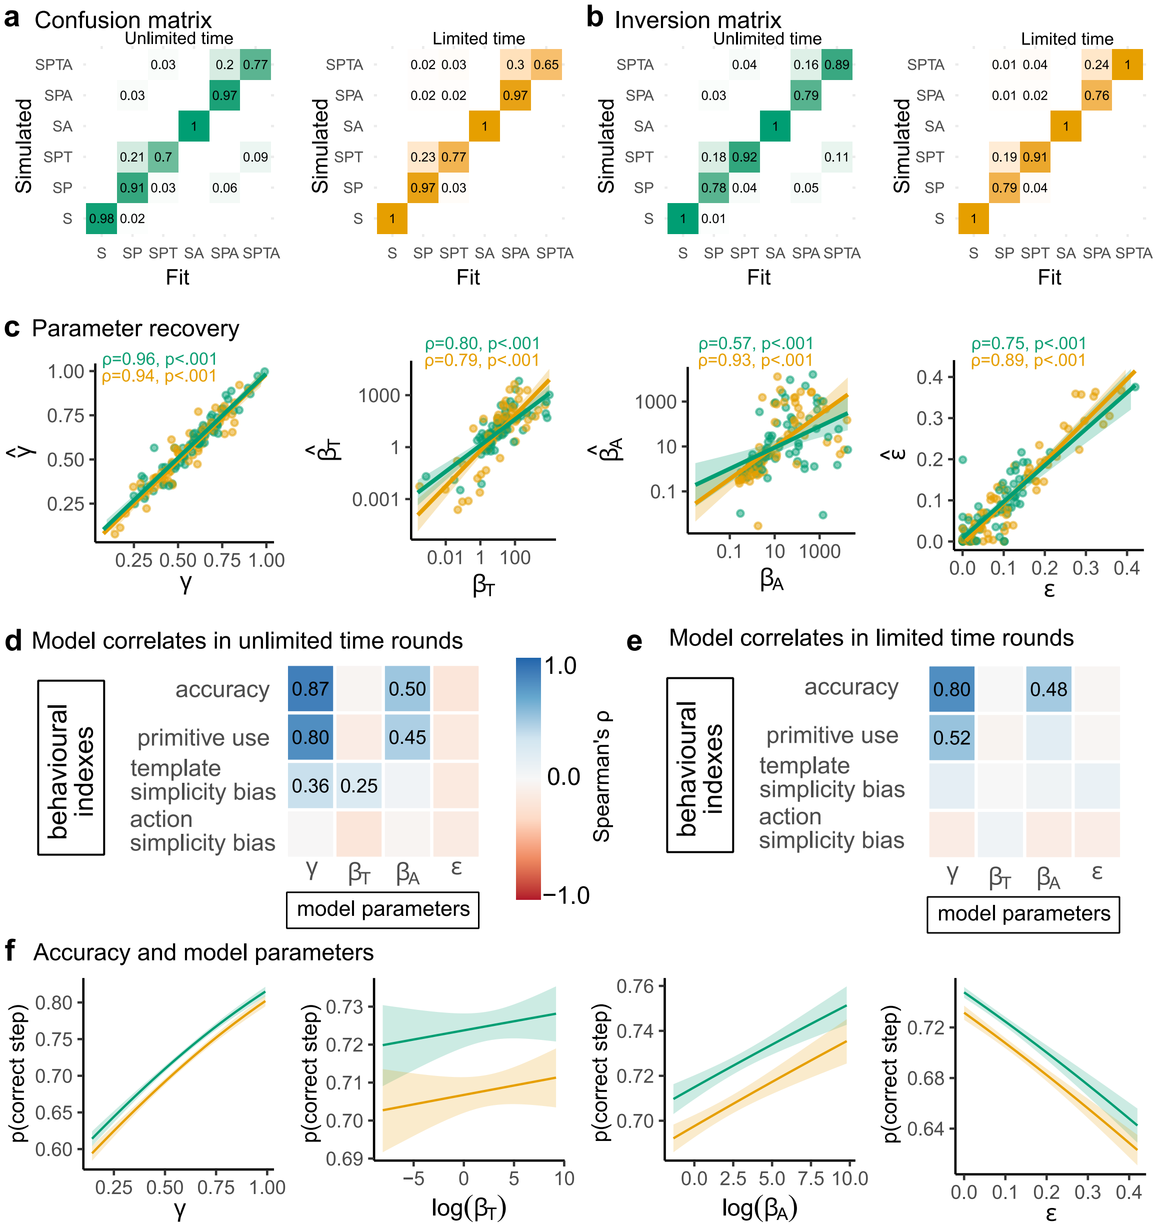

Figure S3. **Exp. 1 supplementary modelling results**. (**a**) Confusion matrices showing model recovery, (fitted model|simulated model), for unlimited (left) and limited (right) time rounds. Data were simulated from each model using participant-specific fitted parameters, and all candidate models were refit to the simulated data. The best-fitting model was selected by lowest BIC. Each cell indicates the fraction of simulated datasets for which a given fitted model was selected. Rows denote the simulated (true) model and columns denote the fitted model. Numerical values are shown only for non-zero entries. The strong diagonal structure indicates good model recoverability. (**b**) Inversion matrices showing model inference, p(simulated model|fitted model), for unlimited (left) and limited (right) time rounds. Using the same simulations and BIC-based model selection procedure as in (**a**), each cell indicates the probability that a given generating model produced the data, conditional on a particular model being selected as best fitting. Again, the strong diagonal structure rules out that winning models arise from data generated by other models. (**c**) Parameter recovery for the SPTA model. For each participant, we used the estimated SPTA parameters fit to their empirical data to generate simulated datasets, and then refit the SPTA model to these simulations. Panels show recovered versus generating parameters for γ, β_T and β_A (both shown on a logarithmic scale) and ε. Recovered and generating parameters showed strong monotonic relationships, indicating good parameter recoverability. (**d,e**). Spearman's ρ correlation coefficients between behavioural indices and model parameters, with coefficients only for significant correlations correlation ($p < .05$), in unlimited (**d**) and limited time (**e**) conditions. (**f** Mixed-effects logistic regressions regressing accuracy on SPTA parameters).

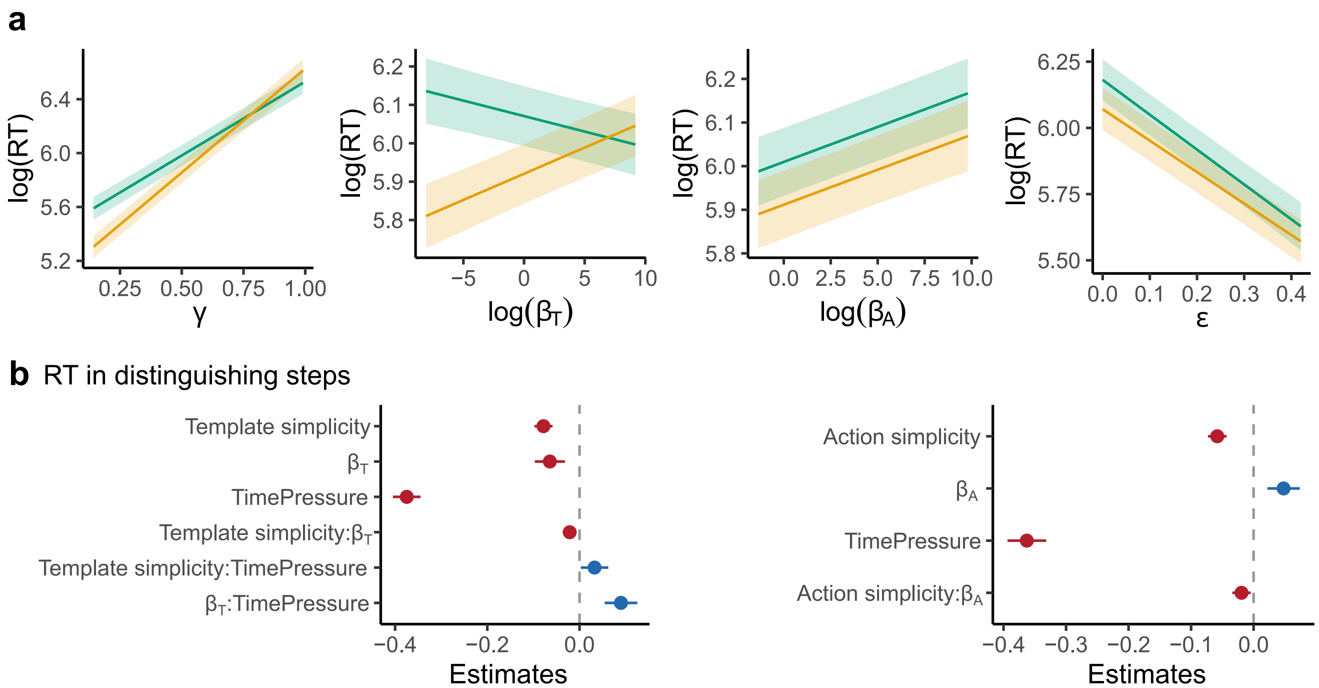

Figure S4. **Exp. 1 model parameters and RTs**. (**a**) Mixed-effects linear regressions predicting RTs on all trials on the basis of time pressure and model parameters. (**b**) Effects of β_T (left) and β_A (right) on RTs in distinguishing trials (where participants chose between primitives) and their interactions with completion simplicity.

negative interaction between action simplicity and β_A ($b = -0.02$, $[-0.03, -0.00]$, $p = .010$) indicates that the RT advantage for simpler actions was stronger for participants with higher β_A.

Exp. 2

Logistic mixed-effects regressions of accuracy on complexity

We again fit a mixed-effects logistic regression to predict accuracy. In the baseline model (BIC = 59377, $\chi^2(1) = 212.84$, $p < .001$), accuracy increased with step (OR: 1.26 [1.23, 1.28], $p < .001$) and round number (OR: 1.05 [1.03, 1.07], $p < .001$). We compared the same complexity measures used in Exp. 1: actions entropy (BIC = 61965; $\chi^2(1) = 4.24$, $p = .039$), templates entropy (BIC = 61960; $\chi^2(1) = 9.86$, $p = .002$), actions LZW (BIC = 61968; $\chi^2(1) = 1.60$, $p = .206$), templates LZW (BIC = 61943; $\chi^2(1) = 26.99$, $p < .001$), actions change complexity (BIC = 61970; $\chi^2(1) = 0.00$, $p = .958$), templates change complexity (BIC = 61967; $\chi^2(1) = 2.89$, $p = .089$), actions LoT Complexity (BIC = 61933; $\chi^2(1) = 36.40$, $p < .001$), templates LoT Complexity (BIC = 61968; $\chi^2(1) = 0.96$, $p = .328$), actions simplicity (BIC = 61856; $\chi^2(1) = 113.40$, $p < .001$), templates simplicity (BIC = 61968; $\chi^2(1) = 1.48$, $p = .22$). We added an interaction term to the models that improved the baseline one: actions entropy (BIC = 61975; $\chi^2(1) = 1.16$, $p = .28$), templates entropy (BIC = 61970; $\chi^2(1) = 0.54$, $p = .460$), templates LZW (BIC = 61947; $\chi^2(1) = 5.59$, $p = .018$), actions LoT Complexity (BIC = 61860; $\chi^2(1) = 84.11$, $p < .001$), actions simplicity (BIC = 61833; $\chi^2(1) = 33.45$, $p < .001$).

In the winning model, accuracy increases with round number (OR: 1.05 [1.03, 1.07], $p < .001$), step number (OR: 1.30 [1.27, 1.32], $p < .001$) and actions simplicity (OR: 1.12 [1.10,1.14], $p < .001$). Actions simplicity interacts significantly with step (OR: 1.06 [1.04, 1.08], $p < .001$) but not with round (OR: 0.99 [0.98, 1.01], $p = .586$) number.

In the completion analysis, we modeled participants' choices in distinguishing steps with a logistic regression (BIC: 13541). We found that the probability of choosing an action increased if it was consistent with the previous primitive (OR: 2.13, [1.92, 2.37], $p < .001$) and with its marginal (OR: 2.10 [1.32, 3.35], $p = .002$) and conditional probability (OR: 4.76 [2.59, 8.75], $p < .001$) and with the actions (OR: 2.23 [2.06, 2.41], $p < .001$) and template simplicity (OR: 1.41 [1.34, 1.49], $p < .001$) of its corresponding completion.

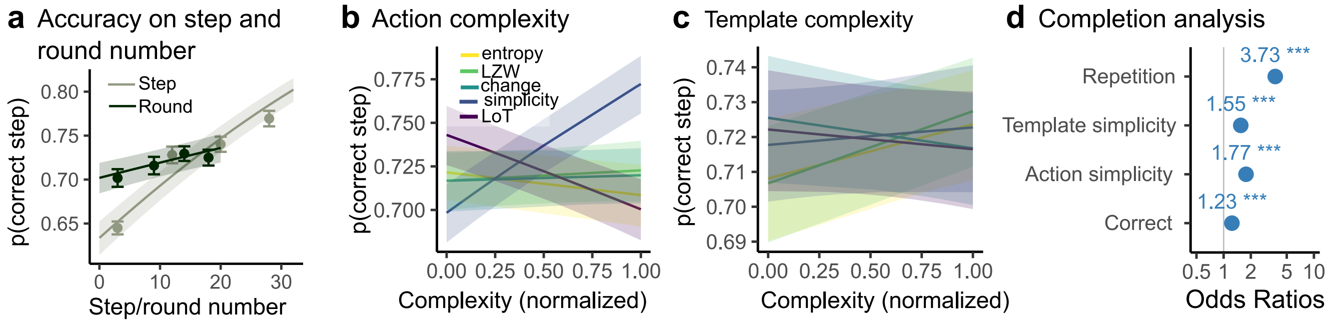

Figure S5. **Exp. 2 supplementary results**. (**a**) Mixed-effects logistic regression predicting accuracy based on round number. (**b-c**) Effects of maze complexity measures, defined either at the level of actions (**b**) or templates (**c**). (**d**) Completion analysis. Main effects from a logistic regression predicting the probability of choosing an action.

Model parameters and behavioural correlates

We examined Spearman correlations between behavioural measures and model parameters across all rounds (Figure S6g). Template simplicity bias was positively associated with γ ($\rho = 0.41$, $p < .01$), whereas action simplicity bias was not significantly related to γ ($\rho = -0.17$, $p = .188$). Primitive use and accuracy were strongly positively correlated with γ ($\rho = 0.82$ and $\rho = 0.93$, respectively; both $p < .001$). Neither template nor action simplicity bias was significantly associated with β_T (template: $\rho = 0.12$, $p = .371$; action: $\rho = -0.10$, $p = .473$). Primitive use ($\rho = -0.15$, $p = .269$) and accuracy ($\rho = 0.04$, $p = .760$) were also not significantly correlated with β_T. In contrast, primitive use and accuracy were positively associated with β_A ($\rho = 0.42$ and $\rho = 0.60$, respectively; both $p < .001$). Neither template simplicity bias ($\rho = 0.22$, $p = .089$) nor action simplicity bias ($\rho = -0.02$, $p = .890$) was significantly related to β_A. Finally, ε was not significantly correlated with any behavioural measure (action simplicity bias: $\rho = -0.09$, $p = .520$; template simplicity bias: $\rho = -0.08$, $p = .528$; primitive use: $\rho = -0.03$, $p = .834$; accuracy: $\rho = -0.13$, $p = .345$).

We fit a mixed-effects linear regression predicting log-transformed RTs, including random intercepts for participants (Figure S3c). RTs were not significantly associated with γ ($b = 0.07$, [-0.06, 0.20], $p = .277$) or β_T ($b = 0.01$, [-0.10, 0.11], $p = .886$). In contrast, RTs decreased with ε ($b = -0.19$, [-0.30, -0.08], $p < .001$) and increased with β_A ($b = 0.17$, [0.04, 0.31], $p = .012$).

RTs in distinguishing steps

As for Exp. 1, we investigated RTs of choices in distinguishing steps (Figure S7c). In the template-level analysis, RTs decreased with template simplicity ($b = -0.07$, [-0.09, -0.06], $p < .001$), whereas β_T showed no main effect ($b = 0.01$, [-0.12, 0.14], $p = .850$). Crucially, we observed a negative interaction between template simplicity and β_T ($b = -0.03$, [-0.05, -0.01], $p = .007$), indicating that the RT advantage for simpler templates was stronger for participants with higher β_T. In the action-level analysis, RTs decreased with action simplicity ($b = -0.09$, [-0.11, -0.07], $p < .001$), and increased with β_A ($b = 0.19$, [0.07, 0.31], $p = .002$). Importantly, we observed a negative interaction between action simplicity and β_A ($b = -0.04$, [-0.06, -0.02], $p < .001$), indicating that the RT advantage for simpler actions was stronger for participants with higher β_A.

Simplicity and accuracy

To understand how the structure of our mazes influenced action and template simplicity biases, we simulated the performance of i) an agent that always chose actions maximizing action simplicity and ii) an agent that chose primitives maximizing template simplicity (Fig. S7b). We then controlled for random baseline accuracy by, respectively, i) subtracting $1/3$ because there are three possible actions, and ii) subtracting $1/2$ because there are two possible primitives. Thus, this simulated performance (after controlling for random chance) indicates the normative advantage of action- and template-level simplicity biases.

These results revealed that both simplicity biases lead to accuracy significantly above chance (action: $t(19) = 16.1$, $p < .001$, $d = 3.6$; template: $t(19) = 3.1$, $p = .006$, $d = 0.7$) , and that an action simplicity bias contributed more to improving accuracy than a template simplicity bias ($t(19) = 7.1$, $p < .001$, $d = 2.4$). Furthermore, when restricting the analysis to steps where the second primitive had already been revealed, the benefit of template simplicity bias was not significantly different from 0 ($t(19) = 3.1$, $p = .006$, $d = 0.7$). Similarly, in the completion analysis, we used only steps where both primitives have been already discovered to avoid inflating the template simplicity bias. Therefore, while the action simplicity bias could be in part induced by the structure of the maze, the template simplicity bias we found cannot be only due to task incentives.

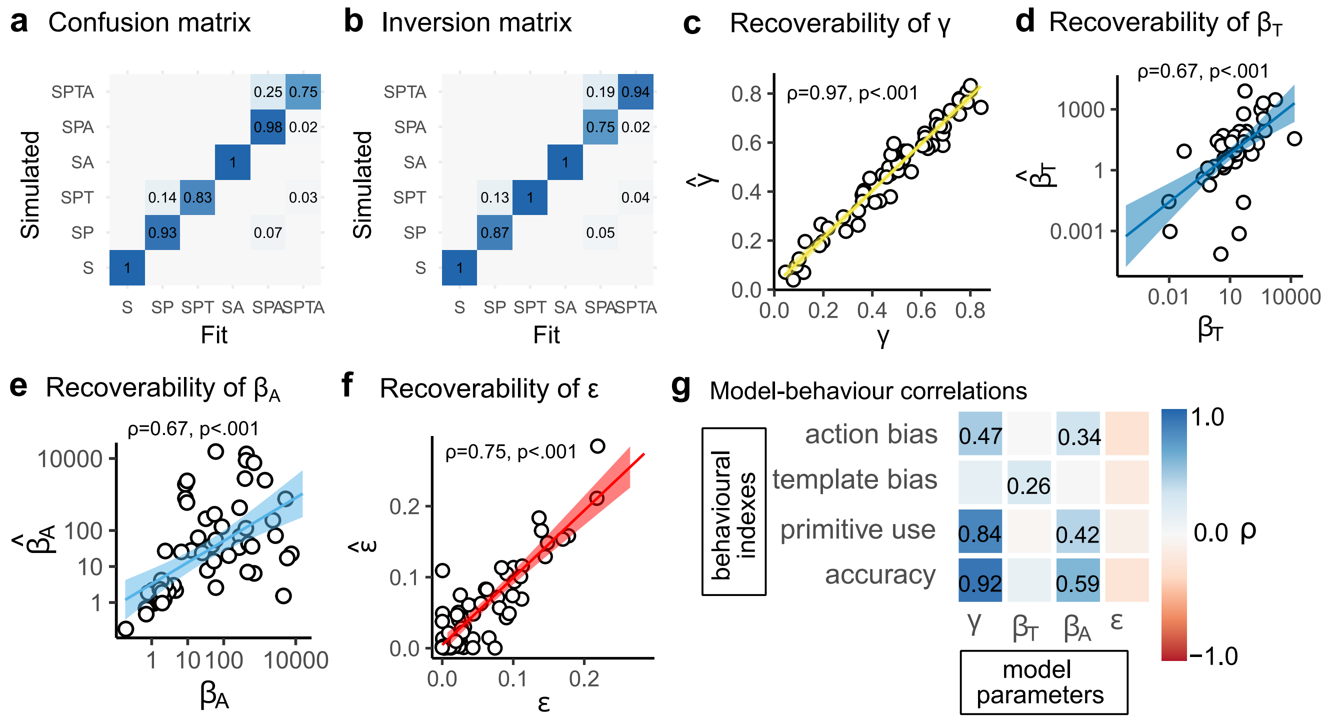

Figure S6. **Exp. 2 supplementary modelling results**. (**a**) Confusion matrices showing model recovery, (fitted model|simulated model). (**b**) Inversion matrices showing model inference, p(simulated model|fitted model) Model recovery and inversion were performed using the same simulation, refitting, and BIC-based model selection procedure as in Exp. 1. (**c-f**) Parameter recovery for the SPTA model. The panels show recovered versus generating parameters for γ (**c**), β_T and β_A (**d** and **e**, respectively, both shown on a logarithmic scale) and ε (**f**). (**g**) Spearman's ρ correlation coefficients between behavioural indices and model parameters.

Exp. 3

Materials

For Exp. 3, we generated actions sequences using the conditional and marginal probabilities of actions in the paths for Exp. 1 and Exp. 2 (Fig. S1c,d). We constrained sequences so that they satisfied the same requirements as those of Exp. 1 and Exp. 2 and so that marginal (Fig. S1a) and conditional probabilities (Fig. S1b) of actions matched the previous ones.

Logistic mixed-effects regressions of accuracy on complexity

We modeled accuracy in Experiment 3 using a series of mixed-effects logistic regressions. As round number did not have a significant effect on accuracy (OR: 0.98 [0.97, 1.00], $p = .051$), we used a baseline model (BIC = 83990) that only included the positive effect of step number (OR: 1.06 [1.05, 1.08], $p < .001$). We compared the same complexity measures used in Exp. 1: actions entropy (BIC = 83988; $\chi^2(1) = 12.71$, $p < .001$), actions LZW (BIC = 83999; $\chi^2(1) = 1.21$, $p = .272$), actions change complexity (BIC = 83999; $\chi^2(1) = 1.93$, $p = .164$), actions LoT Complexity (BIC = 83978; $\chi^2(1) = 22.51$, $p < .001$), actions simplicity (BIC = 83947; $\chi^2(1) = 53.99$, $p < .001$). We then added the interaction term for complexity measures significantly improving the fit: actions entropy (BIC = 83990; $\chi^2(1) = 9.34$, $p = .002$), actions LoT Complexity (BIC = 83979; $\chi^2(1) = 10.68$, $p = .001$), actions simplicity (BIC = 83952; $\chi^2(1) = 5.37$, $p = .020$). In the winning model (BIC = 83952), accuracy increases with step number (OR: 1.06 [1.05, 1.08], $p < .001$) and with actions simplicity (OR: 1.06 [1.04, 1.08], $p < .001$); the two regressors interact negatively (OR: 0.98 [0.97, 1.00], $p = .020$).

Step-by-step completion analysis

In the completion analysis, we applied a different approach from that used in Experiments 1 and 2. For each step, we generated the three possible completions by appending each of the three possible actions to the observed subpath. For each completion, we calculated its relative simplicity by averaging the pairwise differences in action-level simplicity with respect to the other two alternatives. We then used a logistic regression to predict participants' choices based on this simplicity measure and additional covariates. We found that adding interaction terms for these regressors with step and round number improved the fit significantly (BIC: 78901, $\chi^2(8) = 703.99$, $p < .001$). We found that the probability of choosing an action increases with its marginal (OR: 4.95 [4.71, 5.20], $p < .001$) or conditional (OR: 1.10 [1.04, 1.16], $p < .001$) probability and if it is the same as the previous one (OR: 1.17, [1.09, 1.25], $p < .001$). Furthermore, it scales with the action simplicity of its completion (OR: 2.25 [2.16, 2.35], $p < .001$). Marginal probability effect decreases with round number (OR: 0.98 [0.98, 0.99], $p < .001$), perhaps because participants come to rely more on conditional probability, although we did not find any significant interaction for it (OR: 1.00 [1.00, 1.01], $p = .264$). Step number interacts positively with marginal probability (OR: 1.17 [1.14, 1.20], $p < .001$) and negatively with conditional probability (OR: 0.93 [0.91, 0.96], $p < .001$), perhaps because participants used the action *up* (which is the highest in marginal probability) more when close to the upper goal row. The repetition strategy (repeating the last action) did not interact either with step (OR: 0.98 [0.95, 1.01], $p = .235$) or round (OR: 1.00 [0.99, 1.00], $p = .233$) number. Lastly, the effect of actions simplicity decreased with step number (OR: 0.84 [0.82, 0.86], $p < .001$), but increased with round number (OR: 1.01 [1.01, 1.02], $p < .001$). These interactions suggest that participants corrected their bias towards simplicity within each maze, noticing the lack of regularities, but did not override it in next rounds.

Model parameters and behavioural correlates

We examined Spearman correlations between behavioural measures and the model parameters β and ε. Accuracy was not significantly correlated with either β ($\rho = -0.17$, $p = .209$) or ε ($\rho = -0.16$, $p = .227$). In contrast, action simplicity bias was positively associated with both β ($\rho = 0.47$, $p < .001$) and ε ($\rho = 0.39$, $p = .003$). We also found that ε and β_A were strongly correlated ($\rho = 0.90$, $p < .001$). Therefore, we also used partial correlations (Fig. S9f) and found action simplicity bias remained positively associated with β_A after controlling for ε ($\rho = 0.31$, $p = .017$). Instead, the association between action simplicity bias and ε, controlling for β_A, was not ($\rho = -0.10$, $p = .440$).

We examined associations between behavioural measures and model parameters across all rounds (Figure S9g). Accuracy was not significantly related to ε (OR = 0.98, [0.93, 1.03], $p = .417$) or to β_A (OR = 1.00, [0.94, 1.05], $p = .936$).

We then fit a mixed-effects linear regression predicting RTs, including random intercepts for participants (Figure S9h). RTs were not significantly associated with ε ($b = 0.09$, [-0.02, 0.20], $p = .122$) or with β_A ($b = 0.00$, [-0.11, 0.12], $p = .973$).

RTs in distinguishing steps

We investigated RTs of choices in distinguishing steps (Figure S9i). RTs decreased with action simplicity bias ($b = -0.01$, [-0.01, -0.00], $p = .001$), whereas β_A showed no significant main effect ($b = 0.07$, [-0.02, 0.16], $p = .114$). We also found a negative interaction between action simplicity bias and β_A ($b = -0.02$, [-0.02, -0.01], $p < .001$).

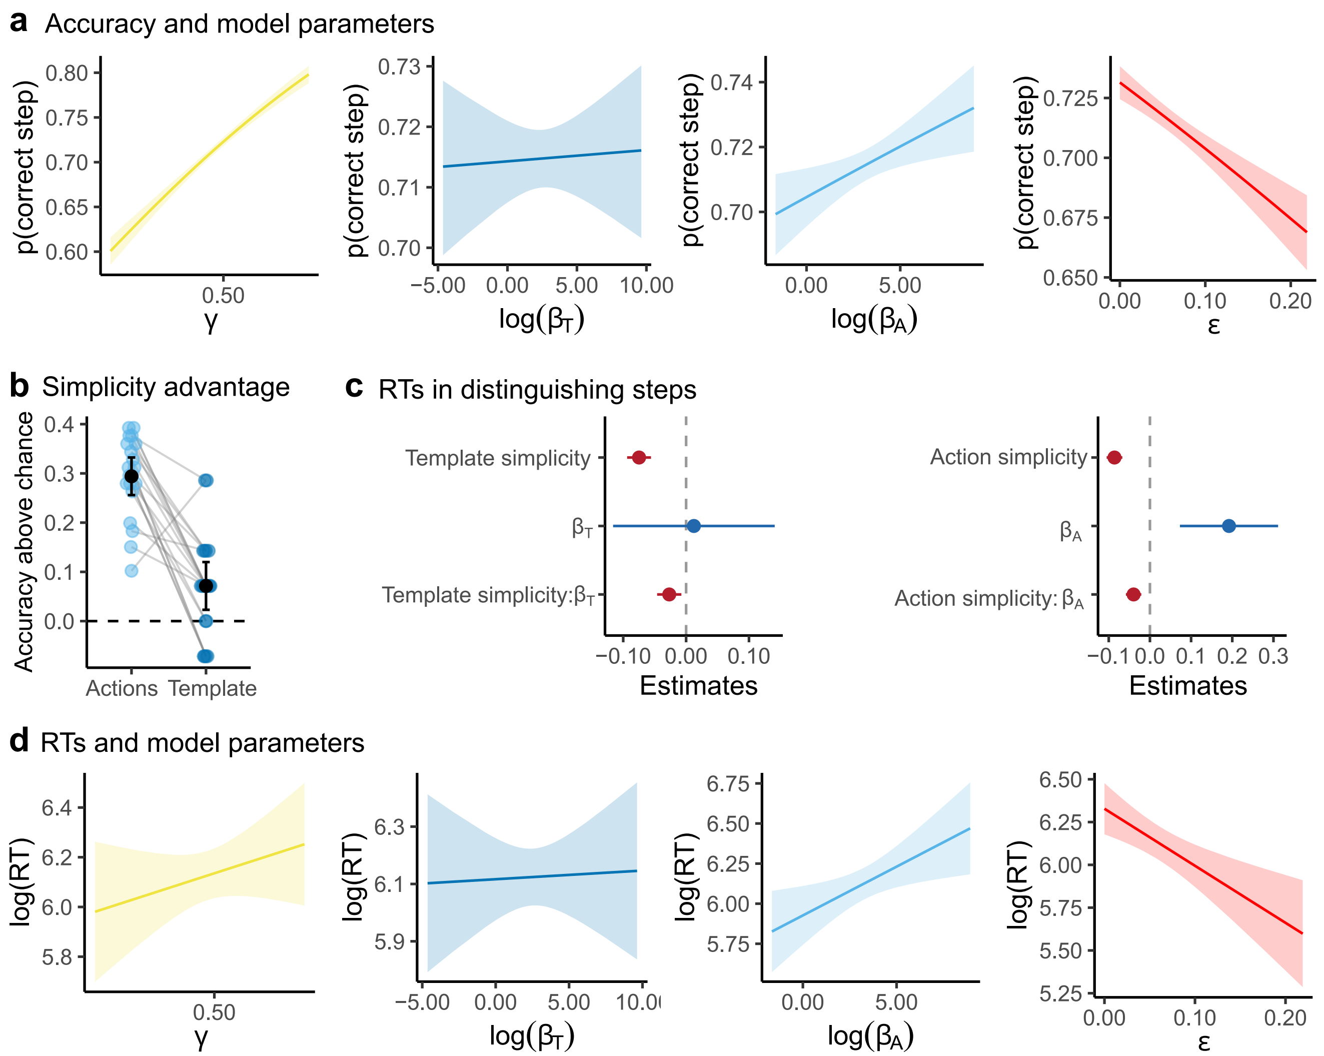

Figure S7. **Exp. 2 Accuracy, RTs, and model parameters** (**a**) Mixed-effects logistic regression predicting accuracy on the basis of model parameters. (**b**) Accuracy advantage of choosing actions based on action and primitives on template simplicity with respect to chance (1/3 for actions, 1/2 for primitives). (**c**) Effects of β_T (left) and β_A (right) on RTs in distinguishing trials (where participants chose between primitives) and their interactions with completion simplicity. (**d**) Mixed-effects linear regressions predicting RTs on all trials on the basis of model parameters

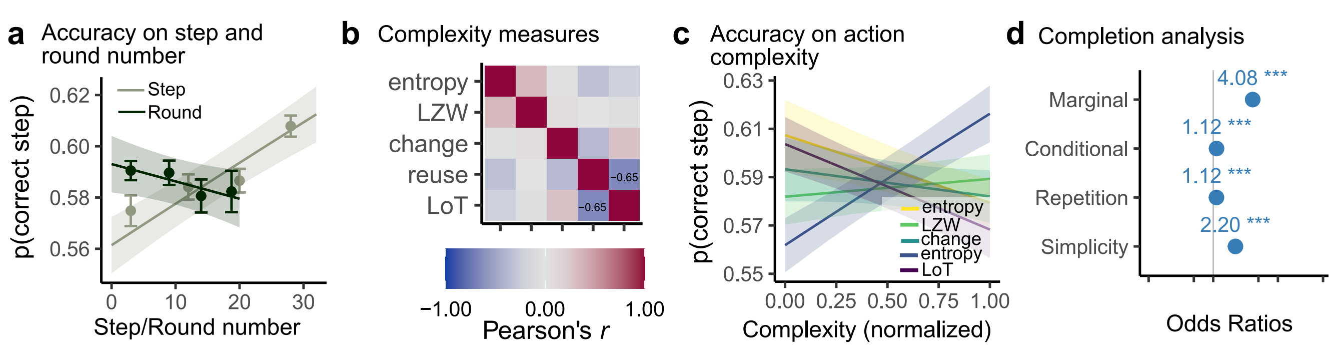

Figure S8. **Exp. 3 supplementary results**. (**a**) Baseline logistic regression predicting participants' accuracy. (**b**) Pairwise correlations among complexity measures. (**c**) Effects of action-level complexity measures on accuracy. (**d**) Completion analysis: main effects from a logistic regression predicting participants' choices.

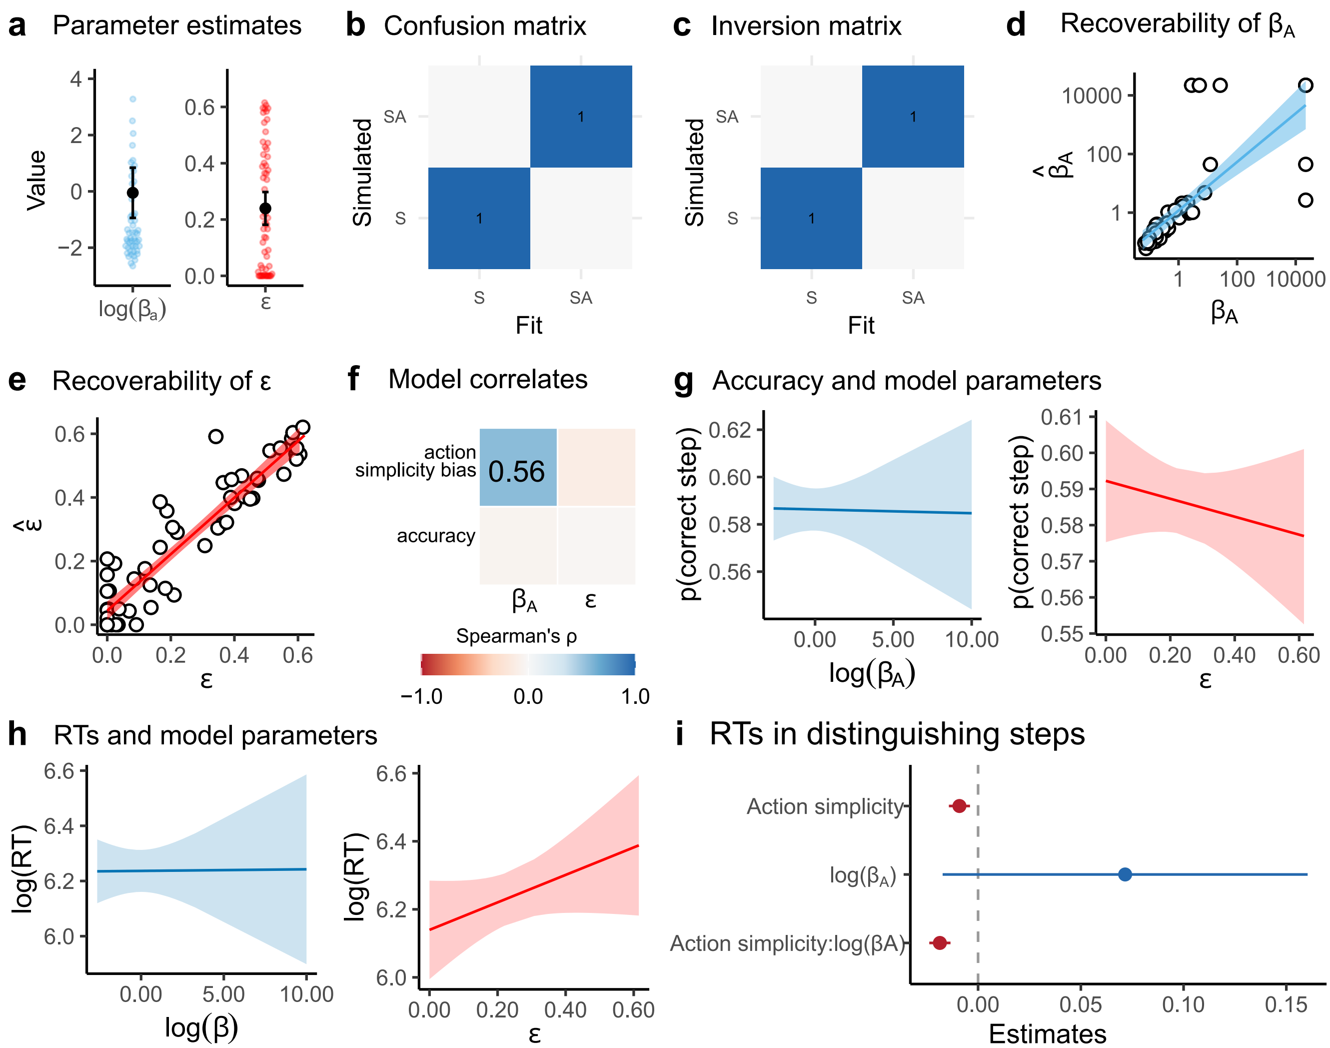

Figure S9. Exp. 3 supplementary modelling results. (**a**) Parameters inspection of the SA model, which best fit participants' choices in Exp. 3. (**b**) Confusion matrices showing model recovery, $(fitted model|simulated model)$. (**c**) Inversion matrices showing model inference, $p(simulated model|fitted model)$ We estimated model confusion and inversion with the same approaches used for the previous experiments. (**d-e**) Parameter recovery for the SA model. The panels show recovered versus generating parameters for β_A (**d**) (shown on a logarithmic scale), and ε (**e**). (**f**) Spearman's ρ correlation coefficients between behavioural indices and model parameters. Given that β_A and ε were correlated across participants, we report partial correlations here. (**g**) Mixed-effects logistic regression predicting accuracy on the basis of model parameters. (**h**) Mixed-effects linear regressions predicting RTs on all trials on the basis of model parameters. (**i**) Effects of β_A and its interaction with completion action simplicity.