Compositionality under time pressure

Valerio Rubino[1,2](valerio.rubino@unitn.it), Mani Hamidi[1], Peter Dayan[3], & Charley M. Wu[1,3]
[1] Human and Machine Cognition Lab, University of Tübingen, Tübingen, Germany
[2] Center for Mind/Brain Sciences, University of Trento, Rovereto, Italy
[3] Max Planck Institute for Biological Cybernetics, Tübingen, Germany

Abstract

Compositionality is a central component of the human faculty for generalization and flexibility. However, the computations involved are poorly understood, especially in terms of their cognitive costs. On one hand, compositionality requires searching combinatorially large hypothesis spaces, raising issues of tractability. On the other hand, compositional representations afford efficient and compact compression. To shed light on the cognitive resource required for compositionality, we used a within-subject time pressure manipulation to study how participants navigated a series of mazes, generated using recursive operations over spatial primitives. We find evidence that behavior is guided by the use of primitives and abstract operations over them, where the degree of compositional structure increases performance and speeds up decisions. And while time pressure led to more random errors, it did not impair the capacity for compositionality. Rather, participants increased their reliance on reusing and recombining previous computations, suggesting a remarkable robustness of human compositional reasoning.

Keywords: compositionality; spatial reasoning; time pressure; resource-rationality;

Introduction

The hallmark of compositional reasoning is the ability to understand complex arrangements of objects as a combination of simpler components. This ability facilitates efficient learning via strong generalizations and flexible behavior (Kurth-Nelson et al., 2023). Indeed, it has been proposed that the capacity to reason about compositional representations is a uniquely human feature, distinguishing us from other animals (Dehaene, Al Roumi, Lakretz, Planton, & Sablé-Meyer, 2022) and state-of-the-art deep neural networks (Garnelo & Shanahan, 2019). Interest in this topic has been renewed by recent work demonstrating the use of compositional representations in language (Hahn, Futrell, Levy, & Gibson, 2022), tool use (Thibault et al., 2021), concept learning (Goodman, Tenenbaum, Feldman, & Griffiths, 2008), visuospatial processing (Amalric et al., 2017; Schwartenbeck et al., 2021) and spatial navigation (Sharma, Curtis, Kryven, Tenenbaum, & Fiete, 2021).

However, it is still unclear how humans learn and use compositional representations. Although adjudicating between compositional hypotheses is thought to involve Bayesian inference (Lake, Ullman, Tenenbaum, & Gershman, 2017; Piantadosi, Tenenbaum, & Goodman, 2016), we lack a more complete understanding of the algorithms that make this process tractable and the computational resources required.

In the present study, we study the effect of time pressure on the use of compositionality. In our task, participants navigated through mazes with hidden structure generated by recursively combining spatial primitives that yield templates with different levels of complexity (Fig. 1). We manipulated time pressure (within-subject) in order to shed light on the impact of constrained cognitive resources on the use of compositional reasoning.

Evidence for compositional representations

Two main approaches provide behavioral evidence for the role of compositional representations in human learning.

One approach studies how the *learnability* of different types of problems is sensitive to compositional structure (Dehaene et al., 2022). These learnability studies span multiple domains of concept learning (Piantadosi et al., 2016), spatial geometry (Kumar, Dasgupta, Cohen, Daw, & Griffiths, 2020; Amalric et al., 2017; Sablé-Meyer et al., 2021), and auditory sequences (Planton et al., 2021). For instance, Kumar and colleagues (2022) presented participants with boards made up of tiles, given the goal of uncovering all target tiles. Crucially, the distribution of the hidden tiles were drawn by either compositional grammars (”abstraction boards”) or by an unstructured process which matched the statistical properties of the abstraction boards (”metamer boards”). Participants achieve higher accuracy on abstract boards, corroborating the hypothesis that humans can exploit the compositional structure of a task. A slightly different approach consists of relating the learnability of a task to the length of the shortest Language of Thought (LoT) expression. Amalric and colleagues (2017) asked participants to predict and repeat sequences displayed on a clock-like display. The complexity of a sequence (defined as the length of the shortest LoT expression), correlated with difficulty in predicting and repeating the sequences. This suggests that participants apply a similar LoT to parse sequences.

The other line of research has demonstrated that compositionality provides better *descriptive* models of how people learn concepts (Goodman et al., 2008; Zhou & Lake, 2021), functions (Schulz, Tenenbaum, Duvenaud, Speekenbrink, & Gershman, 2017), and perform spatial navigation (Sharma et al., 2021). For instance, Zhou and Lake (2021) asked participants to categorize programmatically generated images composed of geometric shapes. Only a model endowed with the ability to evaluate the compatibility of images with compositional rules was able to capture participant responses, and did so much better than alternative models that grouped images by their low-level visual similarities. Furthermore, Sharma et al. (2021) had participants explore the floors of a building consisting of repeated units to collect rewards. To perform the task efficiently and anticipate unobserved regions of the floor, participants had to infer the building blocks and their arrangement in order to plan their paths accordingly. Again, human behavior was best accounted for by models that used and exploited the compositional structure of the environment.

Manipulating cognitive resources with time pressure. Whereas there is ample evidence that humans learn compositional representations and use them to solve novel problems, the computations involved and their cost relative to non-compositional strategies are still poorly understood. This motivates our use of a time pressure manipulation to explore whether compositional reasoning is easier or harder to deploy when limiting the availability of cognitive resources.

Goals and Scope

In our study, participants were tasked with navigating through a series of mazes, where a hidden path was generated using a series of spatial primitives combined into an abstract, recursive template (Fig. 1). Whereas previous work composed sequences over single actions (Planton et al., 2021), our task facilitates composition of actions over two levels of abstraction: First, spatial primitives consisting of four individual actions (e.g., *left-left-up-up*); and second, the composition of two different primitives over the length of the maze. This design allows us to assess the extent to which the compositional structure of higher-level templates influences the ability of participants to navigate each maze efficiently. We present both behavioral and model-based analyses to show that participants solve the task by deploying a rich form of compositional reasoning, which allows them to exploit previously observed subsymmetries and repeats. Although time pressure negatively affects performance by introducing more random errors, it does not diminish the capacity for compositionality.

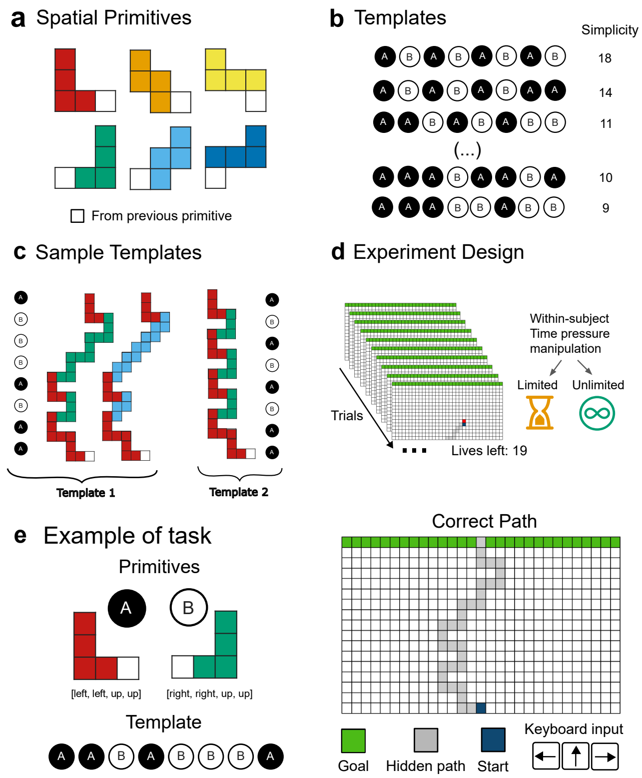

Figure 1: Task overview. **a**) The 6 spatial primitives used in the task, where the white block indicates the position of the last tile of the previous primitive. **b**) Examples of sequence templates ordered in decreasing simplicity from top to bottom. **c**) Three examples of paths, with the first and the second sharing the same underlying template and the third one having the same primitives pair as the first but a different template. **d**) Experiment design. Participants navigated a series of 40 mazes with hidden structure, where we manipulated the presence or absence of time pressure within-subject in randomized, interleaved order. **e**) An illustration of the task, showing the generative primitives and template.

Methods

Participants navigated a series of 40 mazes, each with a single admissible hidden path. Stepping off of the path resulted in losing a "life", with participants incentivized to solve each maze with as many lives remaining as possible. Each hidden path was generated using a pair of spatial primitives (Fig. 1a) that were combined using a binary algebraic patterns (Fig 1b; similar to Planton et al., 2021). A within-subject time pressure manipulation (randomly interleaved) gave participants either *unlimited* time or a *limited* 20 second budget to navigate each maze (Fig. 1d).

Participants and Design. We recruited 71 participants on Prolific. After excluding 5 with missing data or performing below chance, we had a total sample of $n = 66$ (33 female; M_{age}=38.00; *SD*=10.95). Participants were paid £3.75 for taking part in the experiment and a performance contingent bonus of up to £3.75. Participants spent 29.0 $\pm$ 10.68 minutes on the task and earned £4.71 $\pm$ 0.59 in total. The study was approved by the Ethics in Psychological Research Commission of the University ot Tübingen (Wu_2021/0124/213) and informed consent was obtained from all subjects.

The experiment used a within-subject design manipulating time pressure (Fig. 1d). 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.

Materials and Procedure. Mazes were constructed on a 17 row by 33 column grid world, each with single hidden solution that was generated with a compositional structure (see below). Each round began from the center tile on the bottom row, and the goal was to reach the top row colored in green (Fig. 1e). On each trial, participants chose an action using the left, up, and right arrow keys to take a step. The experiment ignored *invalid* choices corresponding to moving outside of the boundaries of the grid or "backtracking" (i.e., selecting left after the previous correct step was right, and vice versa). All models also exclude invalid choices.

After each valid choice, participants received feedback, with successful steps marked in blue and incorrect steps displayed in red. After each incorrect step, participants would lose a life (from a total of 20 on each round) and remain on the last correct tile. Participants were incentivised to complete each maze with as few lost lives (i.e., errors) as possible, with the number of remaining lives after each completed maze determining their reward bonus.

Each maze had a different correct path, which was compositionally generated as a binary *template* of length eight. Each path used two out of the six spatial primitives (Fig. 1a), which were communicated to participants during the tutorial. For instance, the template *ABABABAB* specifies an alternation between primitive A and B, which then repeats four times. In turn, each primitive consisted of four actions (e.g., *up-left-left-up*), which were chosen to be comparable in complexity and able to generate a diverse set of sequences (Fig. 1c).

Initially, there are $2^7 = 128$ possible templates. However, not all templates produce valid solutions (e.g., reaching the goal or staying within the boundaries), thus necessitating additional filters to balance several key factors. We first generated all possible sequences (3840 in total; 30 possible primitives pairings $\times$ 128 templates) and then retained only valid solutions that reached the goal row without exceeding the grid boundaries. Additionally, we filtered sequences to ensure that templates covered a range of complexity levels. Furthermore, we subsampled templates to generate sequences that were balanced in terms of the left-right symmetry of the marginal $p(\text{action})$ and conditional probabilities $p(\text{action}|\text{prevAction})$ of each action. This resulted in a set of 20 unique hidden mazes, which were each presented to participants in both limited and unlimited time conditions.

Results

We first examine how time pressure influences learnability based on the accuracy of choices. We show that participants select actions on the basis of primitives and compare different complexity measures in capturing the learnability of each maze. We are then able to use the winning complexity measure to capture a simplicity bias in how participants complete the maze. Then, our analysis of reaction times (RTs) shows how time pressure influences the speed of choices as as function of template complexity and the type of decision. Lastly, we compare computational models in predicting choices.

Behavioral results

Figure 2a shows that participants achieved lower accuracy (proportion of correct actions) under time pressure ($t(65) = 8.4$, $p < .001$, $d = 0.9$), with an average of $P(\text{Correct}|\text{Unlimited}) = .74$ and $P(\text{Correct}|\text{Limited}) = .69$. Performance in both conditions was also superior to several baseline models, using the true marginal probabilities of actions $P(\text{action}) = .39$ and the true conditional probabilities of actions $P(\text{action}|\text{prevAction}) = .45$.

To get a crude understanding of whether participants used compositional structure, we measured the number of *primitive-inconsistent actions*, which were defined as actions inconsistent with either of the two primitives within the round. Participants made inconsistent actions 14% of the time, which was significantly less than chance ($t(65) = -37.8$, $p < .001$, $d = 4.7$) and higher in limited time rounds ($t(65) = 8.5$, $p < .001$, $d = 0.8$). These results suggest that participants were indeed using primitives, which was impaired to some degree by time pressure. Next, we test various measures of complexity as evidence for a richer form of compositionality.

Which measure of complexity predicts learnability? We consider several measures of complexity, defined either at the level of individual actions or at the template level. If participants are learning to parse sequences by composing primitives, their performance should be sensitive to complexity measures defined at the template level, rather than only at the action level. Later, we examine the complexity of completion patterns, defined only on observed sequences prior to each decision.

Shannon (1948) **Entropy** provides a non-compositional baseline, by measuring the level of uncertainty based on the distribution of actions or primitives in a sequence: $H(X) = -\sum_{x \in X} p(x) \log p(x)$. Entropy is agnostic to the order of elements in a sequence, capturing only frequency with which each element occurs. The Lempel-Ziv-Welch (**LZW**) compression algorithm provides an estimate of Kolmogorov complexity, describing the shortest program that can reproduce a string (Welch, 1984). LZW does so by progressively building a dictionary of symbols that exploit the repetition of substrings in a greedy fashion. **Simplicity** is a measure inspired by Alexander and Carey (1968), which counts how often substrings (of any length) are repeated or mirrored in a sequence. Thus, more repetitions or mirroring of substrings correspond to higher simplicity. **Change complexity** (Aksentijevic, Mihailovic, & T. Mihailovic, 2020) defines the number of changes (i.e., the number of times two consecutive substrings are different), at all possible scales (i.e., substring lengths). For instance, the string ABABABAB can be analyzed at four scales, based on substrings of length 1,2,3 or 4. The number of changes computed at higher scales are downweighted to discount the influence of higher-level structure. Lastly, **LoT complexity** (Planton et al., 2021) defines

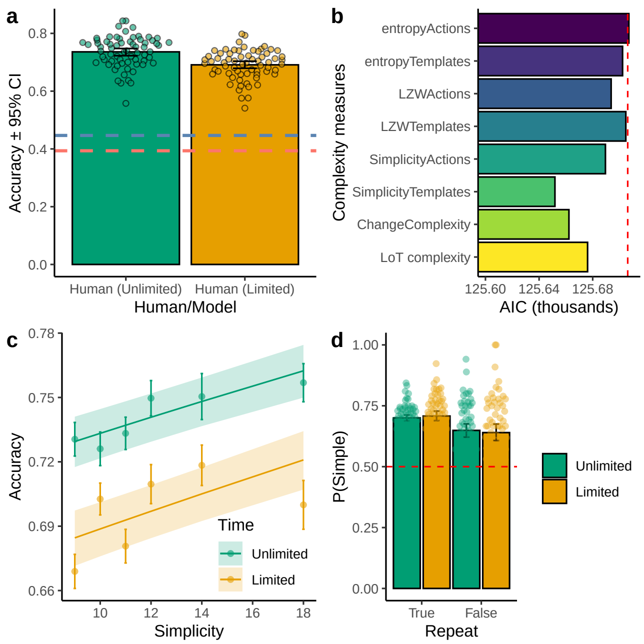

Figure 2: Behavioral results. **a**) Choice accuracy. Each dot is a single participant, with the bars showing the aggregate mean ± 95% CI. Dashed lines illustrate baseline performance, using the true marginal (red) or conditional probabilities (blue). **b**) The AIC of mixed-effect logistic regressions using different complexity measures (along with time pressure) to predict choice accuracy, where template simplicity was the best predictor. The red dashed line is the AIC using time pressure alone. **c**) Three examples of the hidden paths of mazes, the first two using the same template and different pairs of primitives and the latter using another template and the same primitives pair as the first one. **d**) Completion analysis. Probability of actions consistent with selecting a simpler completion (template simplicity). The data is separated based on whether the simpler completion also entails repeating the last completed primitive. Each dot is one participant, with the bars and whiskers showing the aggregate mean (±95% CI).

complexity as the length of a minimal LoT expression, which recursively applies operations such as *repeat* and *vary*. For change complexity and LoT complexity, we use implementations that are only defined over binary variables; thus, we only apply it at the level of templates, whereas all other measures are applied separately for individual actions and at the template level.

We then fit a series of mixed-effects logistic regressions to determine which complexity measure (along with time pressure) best predicted the accuracy of each choice (Fig. 2b). Template simplicity was the best overall predictor (AIC = 125652; Likelihood ratio test against a model with only time pressure: $\chi^2(1) = 56.30$, $p < .001$), followed by change complexity (AIC = 125662; $\chi^2(1) = 45.89$; $p < .001$) and LoT complexity (AIC = 125676; $\chi^2(1) = 31.94$; $p < .001$). Indeed, these three measures all correlate significantly with each other (at least, Pearson's $r = -.75$, $p < .001$, between simplicity and LoT complexity).$r = -0.92$ However, neither model was improved by adding an interaction with time pressure (all $p > .05$). Figure 2c shows the predictions of the winning model plotted against raw data, where we find a linear effect of template simplicity increasing accuracy (Odds Ratio (OR): 1.17 [1.12,1.22], $p < .001$), while time pressure reduced accuracy across the board (OR: 0.80 [0.78,0.83], $p < .001$).

Time pressure and simplicity bias. Next, we explored potential biases in how participants tested which primitive comes next in the sequence. For instance, in rounds pairing the primitives A = "*left-left-up-up*" and B = "*right-right-up-up*" (Fig. 1e), it is possible to distinguish the hypothesized primitive on trials where participants were completing the first action of the primitive (e.g., at the first, fifth, eighth steps of the maze, and so forth).

After filtering out trials where choices were inconsistent with either primitive, we created a set of distinguishing trials, where choices correspond either to primitive A or primitive B. For each of these trials, we appended the chosen primitive (A or B) to the template observed up to that point and computed the simplicity of the resulting hypothesized subtemplate. If ABA has already been observed, selecting an action consistent with primitive A would correspond to a hypothesized subtemplate ABAA.

We then refined the set of distinguishing trials by removing all choices where both completions yielded the same simplicity. This allows us to define *simplicity bias* as the probability of choosing the primitive yielding a simpler subtemplate, instead of choosing the primitive yielding a more complex one. This allows us to quantify the extent participants apply abstract rules (such as *repeat* and *mirror*) on primitives (and their arbitrary compositions) in informing which action they select, where higher simplicity corresponds to a stronger bias towards compositionality. A potential confound is whether the simpler completion corresponds to repeating the last observed primitive, which we also control for.

Figure. 2d shows simplicity bias separated by time pressure and whether the simpler completion corresponded to a repeat (of the last primitive). Overall, participants chose simpler completions at higher than chance level ($t(65) = 24.7$, $p < .001$, $d = 3.0$). We then conducted a two-way repeated measures ANOVA, where we only found a significant effect of repeat, ($F(1,65) = 33.6$, $p < .001$) but not time pressure ($F(1,65) = 0.002$, $p = .961$). This suggests that thesimplicity bias is stronger (but not fully accountable for) when a simpler completion corresponds to repeating the last correct primitive, but time pressure does not diminish this bias.

Reaction times. We first analyzed RTs to observe the influence of time pressure and template simplicity on the speed of choosing each action (Fig. 3a). We entered these two regressors into a mixed regression predicting log-transformed RTs, finding both time pressure (-0.14 [-0.18, -0.11], $p < .001$)

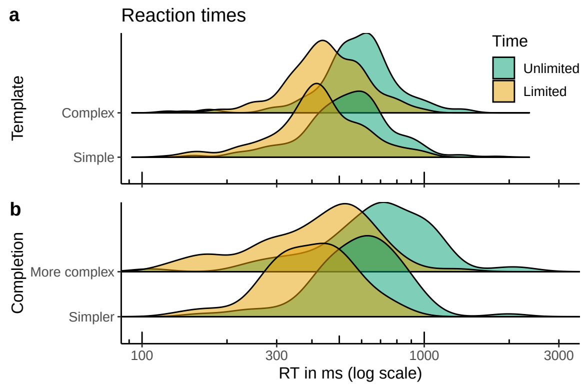

Figure 3: Reaction times. **a**) Distribution of RTs (log scale) as a function of time pressure and binarized template simplicity. **b**) RTs for simpler vs. more complex completions.

and simplicity (-0.10 [-0.12, -0.08], $p < .001$) produced faster RTs. Additionally, we find a significant interaction (-0.07 [-0.09, -0.04], $p < .001$), with even faster RTs for limited time rounds featuring simple templates.

Next, we use the same set of distinguishing trial from the completion analysis to test whether RTs were influenced by simpler or more complex completions (Fig. 3b). We find that both time pressure (-0.45 [-0.50, -0.40], $p < .001$) and simpler completions produced faster RTs (-0.11 [-0.15, -0.07], $p < .001$), along with a significant positive interaction (0.12 [0.05, 0.18], $p < .001$). Whereas simpler completions produce faster RTs with unlimited time, this effect disappears under time pressure, potentially due to a floor effect.

Model results

Lastly, we fit a series of computational models predicting participant choices. We compare several baselines using the marginal and conditional probabilities of actions against a model that exploits the compositional structure of the task.

The *marginal* probability model includes a free parameter ε to combine the true marginal distribution of actions with a uniform policy:

$$\pi_m = (1-\varepsilon)p(\text{action}) + \varepsilon\frac{1}{3} \tag{1}$$

Similarly, the *conditional* probability model uses the true conditional probability of actions, where ε again controls the level of random errors:

$$\pi_c = (1-\varepsilon)p(\text{action}|\text{prevAction}) + \varepsilon\frac{1}{3} \tag{2}$$

We also present a *compositional* model that predicts actions based on sampling hypotheses about the structure of the environment. This takes the form of a particle filter, where a finite number of sampled particles ($n = 100$) represent hypotheses about the structure of the task. Data from task feedback is used to refine these hypotheses, based on removing invalidated particles and resampling from the remaining valid hypotheses.

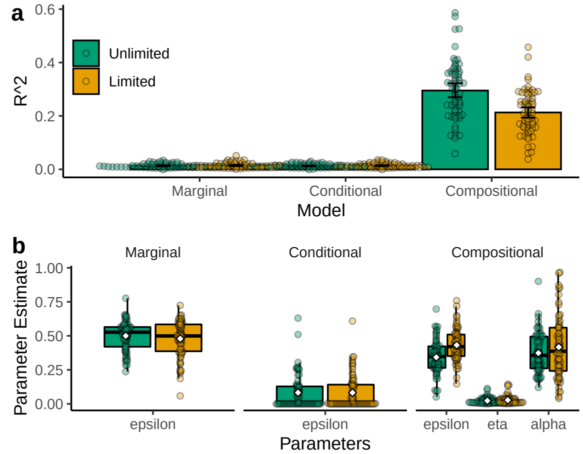

Figure 4: Model results. **a**) Model comparison showing R^2, where $R^2 = 0$ is equivalent to a random baseline and $R^2 = 1$ is a theoretically perfect model. Each dot is one participant, with bars showing the aggregate mean (± 95% CI). **b**) Parameter estimates for each model. Each dot is one participant. The boxplot shows the interquartile range and the white diamond is the aggregate mean. ε captures random errors, η is the probability of selecting invalid primitives, and α is the simplicity bias.

On every between-primitive decision, we sample from the set of primitive pairs that have not been invalidated by the data, with a bias to re-use previously observed substrings:

$$p(\text{primitive}) \approx (1-\eta)p(\text{pair}|\text{data},\alpha) + \eta p(\text{pair}) \tag{3}$$

We include η as an error rate for sampling from all possible primitive pairs (including invalidated pairs) to account for errors in learning which primitives are still valid given the observed feedback. The bias towards re-using substrings is controlled by the parameter α, where substrings are generated using the simplicity measure described previously. With probability $p(\alpha)$, hypothesized primitives are generated by sampling from these substrings, and with $p(1-\alpha)$ we sample from the set of valid primitive pairs. Higher α estimates thus correspond to a stronger bias towards simplicity.

We then sample actions from the distribution of primitives corresponding to the current step, where we again have ε as a source of random error:

$$p(\text{action}) \approx (1-\varepsilon)p(\text{primitive}|\text{data},\text{step}) + \varepsilon\frac{1}{3} \tag{4}$$

Model comparison and parameter estimates. We computed a maximum likelihood estimate for each model and compare them based on BIC. McFadden's Pseudo-R^2 provides an interpretable measure of model fit for any model k relative to a random baseline: $R^2 = 1 - \frac{\text{BIC}_k}{\text{BIC}_{\text{random}}}$. Intuitively, $R^2 = 0$ is equivalent to a random model, while $R^2 = 1$ is a theoretically perfect model. Figure 4a provides a visualization of our model comparison, where although both baselines are better than chance (all $p < .001$), the compositional model wins by an uncontested margin (Unlimited time: $R^2 = .29$; Limited time: $R^2 = .21$), and with better fits with unlimited time ($t(65) = 9.7$, $p < .001$, $d = 0.9$).

Examining the parameters of the winning model, we find that time pressure produced higher rates of random errors ε ($t(65) = 7.7$, $p < .001$, $d = 0.7$), but with no differences in considering invalid primitives η ($t(65) = -1.6$, $p = .114$, $d = 0.3$). However, we find that time pressure induced a higher simplicity bias α ($t(65) = 2.9$, $p = .006$, $d = 0.4$).

Discussion

Compositionality can afford more compressed representations that facilitate better generalization and improved performance (Kurth-Nelson et al., 2023). However, these benefits may come with higher ostensible costs of having to search a combinatorially vast hypothesis space (Fränken, Theodoropoulos, & Bramley, 2022). Here, we investigated how participants navigated compositionally structured mazes, where a within-subject time pressure manipulation allowed us to probe how limiting cognitive resources influences the capacity for compositionality. Our results consistently reveal that rather than being impaired by time pressure, the ability to leverage compositionality remained robust.

Using both behavioral and model-based analyses, we find that although time pressure reduced accuracy (Fig. 2a), this can be largely attributed to an increased rate of random errors (Fig. 4b). Accuracy was best predicted using a compositional measure of *simplicity* (Alexander & Carey, 1968), which is operationalized as the number of subsymmetries and repeats at the template level (Fig. 2b). Surprisingly, this correspondence between simplicity and improved accuracy was not impaired under time pressure (Fig. 2c). This same simplicity measure allowed us to uncover a bias towards simpler completions (i.e., actions that corresponded to hypothesized templates with higher simplicity on distinguishing trials), which persisted even when controlling for choices where the simpler completion entailed repeating the previous primitive (Fig. 2d). Again, time pressure did not alter this bias.

Simpler mazes also generated faster RTs (which was even more pronounced under time pressure; Fig. 2a), with participants also responding faster when making choices that produced simpler completions (Fig. 2b). These results suggest that leveraging compositionality by re-using previously encountered substrings affords quicker responses. Indeed, the best model for predicting choices (Fig. 4a) leveraged the compositional structure of the task and incorporated a simplicity bias capturing the propensity for re-use (Fig. 4b). And while our compositional model uncovered higher rates of random errors under time pressure (ε), we also find an increased reliance on simplicity (α).

Previous work has shown that humans are sensitive to the manipulation of cognitive costs, and adopt simpler learning strategies when put under working memory load (Cogliati Dezza, Cleeremans, & Alexander, 2019) or time pressure (Wu, Schulz, Pleskac, & Speekenbrink, 2022), often falling back to habitual behaviors (Kool, Cushman, & Gershman, 2018). However, compositional representations can afford more efficient compression (Planton et al., 2021) and facilitate learning because primitives and their compositions can be cached and reused (Wingate, Diuk, O'Donnell, Tenenbaum, & Gershman, 2013; Cheyette & Piantadosi, 2017).

Indeed, being able to re-use past computations (i.e., amortization) is thought to be a core feature of how humans navigate the cost-benefit trade-off of performing Bayesian inference (Gershman & Goodman, 2014; Dasgupta, Schulz, Goodman, & Gershman, 2018). Our results suggest that rather than being a costly process requiring increased deliberation time, compositionality may be a core feature of how humans cope with limited cognitive resources.

Limitations and future directions

Our results support previous findings showing that compositional structure predicts the learnability of binary sequences (Planton et al., 2021), while extending the scope to more complex primitives comprised of multiple actions. However, a limitation of our experimental design is that while participants were introduced to the primitives during the tutorial, they were not trained on them. Although it is known that people can acquire new primitives automatically and without explicit training (Tano, Romano, Sigman, Salles, & Figueira, 2020), the capacity for compositionality has also been linked to hippocampal replay (Schwartenbeck et al., 2021; Kurth-Nelson et al., 2023), which may require longer timescales of training. Future studies should investigate how prior knowledge about primitives effects the capability to deploy them in a compositional fashion.

Another promising future direction is to extend our models by i) incorporating meta-learning and ii) refining how we apply simplicity (e.g., over limited subsequences). Whereas all models incorporated elements that relied on the true underlying distribution of actions, primitives, and/or primitive pairs, future extensions could model learning using a Dirichlet distribution parameterized by the number of past observations. Additionally, while we find strong evidence for a simplicity bias, future work could add limits to the number of past observations used to define simplicity (i.e., a limited subsequence) or the number of abstract operations that can be applied.

Lastly, while we focused on time pressure, there are multiple forms of computational costs, which are not mutually exchangeable (Zenon, Solopchuk, & Pezzulo, 2019). Future work could examine alternative manipulations, such as memory load (Dasgupta et al., 2018) or increased task complexity (Haridi, Wu, Dasgupta, & Schulz, 2022).

Conclusions

Rather than impairing the capacity for compositionality, we find that time pressure increases our reliance on re-using previously observed subsymmetries and repetitions, enriching our understanding of the computational mechanisms underlying this singular human ability.

Acknowledgments

CMW is supported by the German Federal Ministry of Education and Research (BMBF): Tübingen AI Center, FKZ: 01IS18039A and funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy–EXC2064/1–390727645. PD is a member of the same cluster of excellence, and is funded by the Max Planck Society and the Humboldt Foundation. VR was funded by the Humboldt Foundation. We thank Philipp Schwartenbeck for insightful discussions and feedback on an early version of the experiment and Santiago Figueira for his help with the computation of LoT complexity.

References

Aksentijevic, A., Mihailovic, A., & T. Mihailovic, D. (2020). Time for change: implementation of aksentijevic-gibson complexity in psychology. *Symmetry*, *12*(6), 948.

Alexander, C., & Carey, S. (1968). Subsymmetries. *Perception & Psychophysics*, *4*(2), 73–77.

Amalric, M., Wang, L., Pica, P., Figueira, S., Sigman, M., & Dehaene, S. (2017). The language of geometry: Fast comprehension of geometrical primitives and rules in human adults and preschoolers. *PLoS computational biology*, *13*(1), e1005273.

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

Cogliati Dezza, I., Cleeremans, A., & Alexander, W. (2019). Should we control? the interplay between cognitive control and information integration in the resolution of the exploration-exploitation dilemma. *Journal of Experimental Psychology: General*, *148*(6), 977.

Dasgupta, I., Schulz, E., Goodman, N. D., & Gershman, S. J. (2018). Remembrance of inferences past: Amortization in human hypothesis generation. *Cognition*, *178*, 67–81.

Dehaene, S., Al Roumi, F., Lakretz, Y., Planton, S., & Sablé-Meyer, M. (2022). Symbols and mental programs: a hypothesis about human singularity. *Trends in Cognitive Sciences*.

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

Garnelo, M., & Shanahan, M. (2019). Reconciling deep learning with symbolic artificial intelligence: representing objects and relations. *Current Opinion in Behavioral Sciences*, *29*, 17–23.

Gershman, S., & Goodman, N. (2014). Amortized inference in probabilistic reasoning. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 36).

Goodman, N. D., Tenenbaum, J. B., Feldman, J., & Griffiths, T. L. (2008). A rational analysis of rule-based concept learning. *Cognitive science*, *32*(1), 108–154.

Hahn, M., Futrell, R., Levy, R., & Gibson, E. (2022). A resource-rational model of human processing of recursive linguistic structure. *Proceedings of the National Academy of Sciences*, *119*(43), e2122602119.

Haridi, S., Wu, C. M., Dasgupta, I., & Schulz, E. (2022). The scaling of mental computation in a sorting task. *PsyarXiv*. doi: 10.31234/osf.io/8hqtv

Kool, W., Cushman, F. A., & Gershman, S. J. (2018). Competition and cooperation between multiple reinforcement learning systems. *Goal-directed decision making*, 153–178.

Kumar, S., Dasgupta, I., Cohen, J. D., Daw, N. D., & Griffiths, T. L. (2020). Meta-learning of structured task distributions in humans and machines. *arXiv preprint arXiv:2010.02317*.

Kumar, S., Dasgupta, I., Marjieh, R., Daw, N. D., Cohen, J. D., & Griffiths, T. L. (2022). Disentangling abstraction from statistical pattern matching in human and machine learning. *arXiv preprint arXiv:2204.01437*.

Kurth-Nelson, Z., Behrens, T., Wayne, G., Miller, K., Luettgau, L., Dolan, R., ... Schwartenbeck, P. (2023). Replay and compositional computation. *Neuron*.

Lake, B. M., Ullman, T. D., Tenenbaum, J. B., & Gershman, S. J. (2017). Building machines that learn and think like people. *Behavioral and brain sciences*, *40*.

Piantadosi, S. T., Tenenbaum, J. B., & Goodman, N. D. (2016). The logical primitives of thought: Empirical foundations for compositional cognitive models. *Psychological review*, *123*(4), 392.

Planton, S., van Kerkoerle, T., Abbih, L., Maheu, M., Meyniel, F., Sigman, M., ... Dehaene, S. (2021). A theory of memory for binary sequences: Evidence for a mental compression algorithm in humans. *PLoS computational biology*, *17*(1), e1008598.

Sablé-Meyer, M., Fagot, J., Caparos, S., van Kerkoerle, T., Amalric, M., & Dehaene, S. (2021). Sensitivity to geometric shape regularity in humans and baboons: A putative signature of human singularity. *Proceedings of the National Academy of Sciences*, *118*(16), e2023123118.

Schulz, E., Tenenbaum, J. B., Duvenaud, D., Speekenbrink, M., & Gershman, S. J. (2017). Compositional inductive biases in function learning. *Cognitive psychology*, *99*, 44–79.

Schwartenbeck, P., Baram, A., Liu, Y., Mark, S., Muller, T., Dolan, R., ... Behrens, T. (2021). Generative replay for compositional visual understanding in the prefrontal-hippocampal circuit. *bioRxiv*.

Shannon, C. E. (1948). A mathematical theory of communication. *The Bell system technical journal*, *27*(3), 379–423.

Sharma, S., Curtis, A., Kryven, M., Tenenbaum, J., & Fiete, I. (2021). Map induction: Compositional spatial submap learning for efficient exploration in novel environments. *arXiv preprint arXiv:2110.12301*.

Tano, P., Romano, S., Sigman, M., Salles, A., & Figueira, S. (2020). Towards a more flexible language of thought: Bayesian grammar updates after each concept exposure. *Physical Review E*, *101*(4), 042128.

Thibault, S., Py, R., Gervasi, A. M., Salemme, R., Koun, E., Lövden, M., ... Brozzoli, C. (2021). Tool use and language

share syntactic processes and neural patterns in the basal ganglia. *Science*, *374*(6569), eabe0874.
Welch, T. A. (1984). A technique for high-performance data compression. *Computer*, *17*(06), 8–19.
Wingate, D., Diuk, C., O'Donnell, T., Tenenbaum, J., & Gershman, S. (2013). Compositional policy priors.
Wu, C. M., Schulz, E., Pleskac, T. J., & Speekenbrink, M. (2022). Time pressure changes how people explore and respond to uncertainty. *Scientific reports*, *12*(1), 1–14.
Zenon, A., Solopchuk, O., & Pezzulo, G. (2019). An information-theoretic perspective on the costs of cognition. *Neuropsychologia*, *123*, 5–18.
Zhou, Y., & Lake, B. M. (2021). Flexible compositional learning of structured visual concepts. *arXiv preprint arXiv:2105.09848*.