Supplementary Materials

S1. Bayesian Models

We modeled behavior in the task as Bayesian belief updating with respect to a hypothesis space $\mathcal{H}$. In the PERM model, $\mathcal{H}$ was comprised of 24 possible orderings of the four monsters shown in Figure 1. In the CA strategy, $\mathcal{H}$ was comprised of 8 possible hierarchical rules involving the two cues (color and shape).

The four monsters were $BS = (blue, square)$, $BC = (blue, circle)$, $GS = (green, square)$, and $GC = (green, circle)$. On each trial of the learning phase a “race" X was observed between two monsters, with a total of six possible unique races. Each hypothesis $h \in \mathcal{H}$ specifies the deterministic likelihood of observing a outcome of a race between two monsters, $p(winner(x) = m|h)$ for $m \in X$, based on their relative rank (PERM model) or the order and direction of the two cues (CA model). On trial t, the learner has observed a set of races $\mathcal{D}$ and their joint likelihood under a hypothesis is $p(\mathcal{D}|h)$. The posterior distribution is given by Bayes rule,

$$p(h|\mathcal{D}) = \frac{p(\mathcal{D}|h)p(h)}{\sum_{h' \in \mathcal{H}} p(\mathcal{D}|h')p(h')}, \tag{1}$$

where the prior $p(h)$ is assumed to be uniform over the hypothesis space. The probability of a new race resulting in the outcome $winner(x) = m$ is then given by the predictive distribution

$$p(winner(x) = m|\mathcal{D}) = \sum_{h \in H} p(winner(x) = m|h)p(h|\mathcal{D}). \tag{2}$$

Selections during the active learning phase. Selections were modeled using stepwise expected information gain (EIG), which is the expected reduction in Shannon entropy measured over the posterior distribution as a result of observing the outcome of a race. Shannon entropy is given by

$$H(\mathcal{D}) = -\sum_{h \in \mathcal{H}} p(h|\mathcal{D})\ log\ p(h|\mathcal{D}), \tag{3}$$

where $\mathcal{D}$ is the set of observations thus far. Entropy is maximized when all hypotheses have equal probability and is equal to zero when one hypothesis has $p(h|\mathcal{D}) = 1$. The value of selecting test X is the expected decrease in entropy resulting from each outcome of the test, weighted by its probability of occurring,

$$EIG(X) = \sum_{m \in X} p(\text{winner}(X) = m|\mathcal{D})\,[H(\mathcal{D}) - H((\text{winner}(X) = m), \mathcal{D})]. \tag{4}$$

Generalization test. There were three trials in the generalization test in which the participant predicted the winner of a race between monsters that were not seen during the learning phase. For the cue-abstraction hypothesis space, the likelihood of each outcome was simply based on the cue order on the known dimension. For example, in the Shape trial participants predicted the outcome of a race between a red square (RS) and red circle (RC). The probability $p(RS > RC)$ was determined by the cue direction for shape under each hypothesis.

For the permutation-based hypothesis space, predictions were based on the relative ranking of monsters that were matched on the feature dimension with an unfamiliar value. For example, in the Shape trial, the probability $p(RS > RC)$ was the proportion of races between monsters of the same color in which the square was higher ranked than the circle. If a given hypothesis specified that a square was faster than a circle for both values of the color dimension (i.e., $BS > BC$ and $GS > GC$), the probability of a red square being faster than a red circle was 1. In contrast to the CA strategy, under the PERM strategy it was possible to have reversed orderings (e.g., the square being faster when the shapes were blue, but slower when the shapes were green, as in the ordering $BS > BC > GC > GS$) which would lead to more uncertainty about the outcome for the novel generalization pair.

Podium test. In the podium task the participant is asked to identify the correct ordering of the four monsters seen during the learning phase. Under the Bayesian model, the probability of each response is simply the posterior probability of the corresponding hypothesis.

S2. Latent Mixture Modeling in Study 2

We used hierarchical Bayesian latent mixture modeling to estimate the relative probabilities of cue-abstraction and permutation-based representations during the learning and test phases. Performance in each phase was modeled as a mixture of three strategies: a random strategy (RAND) that corresponded to random search during learning and guessing at test; the permutation-based model (PERM); and the cue-abstraction model (CA).

The model specifications are shown in Figure S1 for each phase. There were six groups in Study 2 (3 age groups $\times$ 2 conditions). The probability of each strategy in group k was denoted by the mixture probability θ_k, with each strategy assigned an equal prior probability. The mixture probabilities determine the probability of an individual adopting a particular strategy, where the chosen strategy for participant i is denoted by z_i. A value of $z_i = 0$ corresponds to the RAND strategy, under which all search and test choices have equal probability. If an individual adopts the PERM ($z_i = 1$) or CA ($z_i = 2$) representations, the corresponding hypothesis space was used to evaluate the likelihood of their selections (learning phase) or test responses (test phase) based on the Bayesian updating model described in Section S1. Choices in each phase were modeled using softmax functions with group-specific inverse-temperature parameters λ_k (see below for details of each phase). Inverse-temperature parameters were assigned a prior distribution of $Gamma(\alpha, \beta)$, where α and β were hyperparameters that were common to all groups.

Learning phase. On each trial in the learning phase participants could select one of six possible observations (pairwise match-ups of monsters). Under the random strategy, all six observations were equally likely on every trial. Under the PERM and CA strategies, the probability $\rho_{i,t}$ of choosing observation i on trial t was modeled with a softmax choice function over the EIG of all possible observations on that trial. The strategy adopted by an

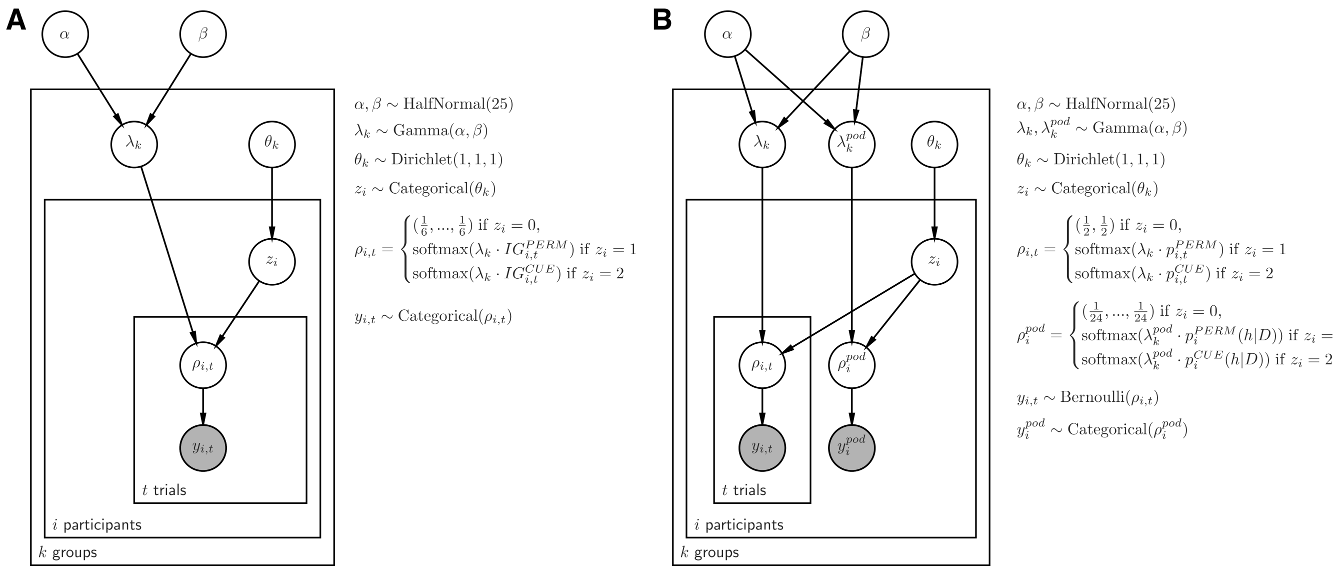

Figure S1. Plate diagrams for hierarchical mixture models in the learning phase (A) and test phase (B) for Study 2.

individual determined whether EIG was calculated according to the PERM or CA strategy.

Test phase. During the test phase, each participant made four responses: three generalization responses (shape, color, and order) and one podium response. For the random strategy, all choices were equally likely. For the PERM and CA strategies, the probability of possible responses were based on the respective hypothesis space and the evidence an individual observed during the learning phase (see Section S1). Choices are again modeled using a softmax function with an inverse-temperature parameter λ_k^{test}.

For responses in the podium task, the probability of selecting a podium was similarly defined via a softmax choice function over the posterior distribution for the corresponding model, with a separate inverse-temperature parameter λ_k^{pod}.

Parameter estimation. Parameters were estimated using Markov Chain Monte Carlo (MCMC) with the PyMC Python library (Salvatier, Wiecki, & Fonnesbeck, 2016). Parameters were estimated separately for the learning and test phases. For both models we used four MCMC chains with 20,000 samples and 2,000 burn-in iterations. All estimates converged as indicated by Gelmin-Rubin ratios under 1.05.

			Mean difference [95% HDI]	
Age	Memory load	Comparison	Active learning phase	Test phase
5	High	PERM - RAND	-.36 [-.93, .23]	-.48 [-.82, -.13] *
		CA - RAND	-.60 [-.92, -.26] *	-.47 [-.79, -.12] *
		CA - PERM	-.24 [-.60, .06]	.02 [-.34, .37]
	Low	PERM - RAND	.40 [-.16, .96]	-.19 [-.70, .38]
		CA - RAND	-.18 [-.52, .13]	-.38 [-.75, .03]
		CA - PERM	-.57 [-.96, -.18] *	-.19 [-.61, .26]
6	High	PERM - RAND	.67 [.20, .99] *	-.32 [-.70, .09]
		CA - RAND	-.07 [-.35, .18]	-.20 [-.66, .28]
		CA - PERM	-.73 [-.99, -.42] *	.13 [-.32, .58]
	Low	PERM - RAND	.26 [-.32, .81]	.30 [-.06, .65]
		CA - RAND	-.25 [-.58, .11]	-.10 [-.38, .18]
		CA - PERM	-.51 [-.86, -.14] *	-.40 [-.76, .01]
7	High	PERM - RAND	.51 [.04, .95] *	-.14 [-.61, .45]
		CA - RAND	.07 [-.32, .48]	-.18 [-.61, .38]
		CA - PERM	-.43 [-.94, .13]	-.03 [-.48, .36]
	Low	PERM - RAND	.73 [.39, .99] *	.05 [-.21, .34]
		CA - RAND	-.04 [-.28, .17]	.41 [.12, .70] *
		CA - PERM	-.77 [-.99, -.50] *	.36 [-.03, .73]

Table S1

Pairwise differences between mixture probabilities θ for each strategy in the active learning and test phases. Differences where 95% HDIs do not overlap with 0 are marked with an asterisk (*).