The Quest for Simplicity in Human Learning: Identifying the Constraints on Attention

Matthew Galdo, Emily R. Weichart, Vladimir M. Sloutsky, and
Brandon M. Turner*

*Department of Psychology, The Ohio State University,
Columbus, OH, USA*

**Corresponding Author. Email: turner.826@gmail.com*

Abstract

For better or worse, humans live a resource-constrained existence; only a fraction of physical sensations ever reach conscious awareness, and we store a shockingly small subset of these experiences in memory for later use. Here, we examined the effects of attention constraints on learning. Among models that frame selective attention as an optimization problem, attention orients toward information that will reduce errors. Using this framing as a basis, we developed a suite of models with a range of constraints on the attention available during each learning event. We fit these models to both choice and eye-fixation data from four benchmark category-learning data sets, and choice data from another dynamic categorization data set. We found consistent evidence for computations we refer to as “simplicity”, where attention is deployed to as few dimensions of information as possible during learning, and “competition”, where dimensions compete for selective attention via lateral inhibition.

keywords: Dynamic Attention | Category Learning | Eye-tracking | Cognitive Bias

1 Introduction

In our everyday lives, we are constantly bombarded with information that we must navigate to make decisions that are consistent with our goals. Typically, information can be organized into dimensions, which can be prioritized according to their relevance in the decision process. For humans, the challenge of making choices based on many dimensions is aided by our evolved ability to selectively attend. Ideally, selective attention allows us to focus on dimensions of information most relevant to our current goals, while ignoring the dimensions that are irrelevant. For example, when judging the freshness of food, one might focus on color, while ignoring shape, but when identifying tools, we may focus on shape, while ignoring color.

Unfortunately, efficient deployment of selective attention is hindered by the fact that we are often initially ignorant about which dimensions are relevant. In these cases, we learn dimension relevance through experience, and begin to judiciously apply attention toward those dimensions as we encounter new items. The variability in attention to dimensions over time creates a path of attention orientation, and this path is directly affected by the salience of our experiences and our goals.

As with all paths, there may be shortcuts worth considering as they too could offer considerable discounts in the time or effort spent in accomplishing our goals. In many cases, shortcuts arise when our knowledge about the landscape of our goals becomes well-informed. While cresting a hill we may gain sight of where our path is ultimately going and decide to simplify our route by cutting through a meadow of tall grasses rather than remaining on the winding trail. In other cases, constraints on our ability may dictate the paths that we elect to travel. When fatigued, for example, the appeal of a path that leads to a resting place may compete against our desire to see the scenic views along the supplementary mountain pass.

In this article, we investigate the nature of the strategic and capacity-related constraints that are imposed on attention during human learning. Thinking in terms of

43 Marr’s levels of analysis [1], there are many types of computational constraints that
44 will dictate the specific algorithms that are used to extract information from the envi-
45 ronment. Here, we use data consisting of choice and eye fixations to infer what types
46 of constraints are placed on attention, and how these constraints ultimately affect the
47 quality of information acquired during learning.

48 **1.1 Modeling Human Learning**

49 Several models of categorization have been developed to explain how characteristics
50 of the task environment, stimulus set, and the goals and experiences of the observer
51 influence behavior. The strategic manipulation of attention across dimensions often
52 serves as a theoretical centerpiece for explaining how humans rapidly learn to cate-
53 gorize novel stimuli when informed by corrective feedback. Nosofsky’s Generalized
54 Context Model [GCM; 2], for example, describes attention as a means of prioritizing
55 the dimensions that are relevant to the task, such that features can be reliably mapped to
56 the correct category label. Although many models similarly conceptualize attention as
57 a mechanism for optimizing how dimensions are processed, variations in the algorithmic
58 definitions of attention often imbue competing theories with different constraints
59 and assumptions about how the computational goals of the observer are realized.

60 In the current article, we focus on two theoretical assumptions about attention that
61 are implicit in some contemporary models of category learning to varying degrees:
62 *simplicity* and *competition*. When identifying the category membership of a new item,
63 the observer must decide which and how many dimensions of information to consider,
64 prior to making a choice. Some models assume this is achieved in part through compe-
65 tition, whereby increasing attention to the most task-relevant dimension diminishes at-
66 tention to the others (e.g. GCM [2], EXIT [3]). Other models additionally assume that
67 observers seek simplicity by minimizing the expenditure of attention over the course
68 of learning, such as considering the cost of obtaining information relative to its utility

69 [e.g., SEA; 4] or minimizing the number of dimensions they attend to and relying upon
70 a simplified subset of relevant information [e.g., EXIT 3]. Although competition and
71 simplicity are assumed by various models of categorization, the respective contribu-
72 tions of these mechanisms have not been formally evaluated. The goals of the current
73 article, therefore, are to investigate the independent contributions of competition and
74 simplicity during learning, and to introduce a new framework for assessing how these
75 mechanisms dynamically influence attention across paradigms and individuals.

76 Our framework extends a dynamic, instance-based version of GCM [5–7] to al-
77 low for the flexible adaptation of attention based on an observer’s unique experiences.
78 Importantly, the framework includes dissociable mechanisms for competition and sim-
79 plicity that can be modulated in order to investigate their independent effects on atten-
80 tion in various task contexts. Given that computational goals (e.g., maximize accuracy,
81 minimize effort) can be expressed in different algorithmic ways, we will quantitatively
82 evaluate mechanisms for simplicity and competition independently, and compare them
83 against consolidated implementations described by existing models. Importantly, our
84 analysis is the first to use both choice and eye tracking data to quantitatively adjudicate
85 different theories of attention during learning. Although other researchers have used
86 eye tracking data to qualitatively provide evidence for selective attention mechanisms
87 in other contexts [8, 9], we determined that qualitative comparisons were not strong
88 enough to infer computational goals from eye tracking data in the current investiga-
89 tion.

90 In summary, our approach is best understood in terms of Marr’s levels of analysis
91 [1]. Our general purpose is to characterize simplicity and competition at the *compu-*
92 *tational* level by observing their component influences on the path of attention. To do
93 this, we instantiated the mechanisms of interest into sets of canonical and independent
94 *algorithmic* definitions within a modeling framework. Assuming a single *implemen-*
95 *tational* structure whereby latent cognitive processes manifest in behavior and visual

attention via eye-tracking, we use model-fitting and systematic comparison to evaluate the candidate mechanisms. As such, we are not only able to gain insight into *which* computations determine the path of attention, but *how* they are algorithmically instantiated during learning.

1.2 Dynamic Attention

Here, we will describe the general implementation of attention that is relevant to our investigation, starting from GCM. At its core, GCM uses dimension-wise similarity between a stimulus probe and a set of previously stored episodic memories (called “exemplars”) to make categorization decisions. A key insight of GCM was the use of a D -dimensional attention vector $\boldsymbol{\alpha}$ that could modulate the psychological similarity between a stimulus probe and exemplars, where larger amounts of attention to a dimension creates better discrimination of information in that dimension. However, GCM assumes that attention is effectively static; there are no trial-to-trial adjustments of attentional parameters. ALCOVE [10] eliminated this assumption, specifying that attention should adapt after each experience toward values that improve accuracy. This process is formalized by defining error as a function of $\boldsymbol{\alpha}$, denoted “loss($\boldsymbol{\alpha}$),” such that

$$\alpha_{t+1,j} = \alpha_{t,j} - \gamma_0 \frac{\partial}{\partial \alpha_{t,j}} \text{loss}(\boldsymbol{\alpha}), \quad (1)$$

where $\gamma_0 > 0$ is a learning rate parameter. Equation 1 describes a first-order optimization (i.e., gradient descent) process that moves an attention vector from some initial value $\boldsymbol{\alpha}_0$ to a location in “attentional space” that minimizes the loss function after some number of iterations (e.g., trials). To specify the loss function, ALCOVE (and later SUSTAIN [11]) used the humble teaching rule, which is a modified version of a sum of squared error function. Here, we use the cross-entropy loss function, which is more widely used and connected to statistical and machine learning norms of classification [12]. Because we are interested in how humans depart from a normative

120 mathematical procedure, cross-entropy minimization serves as a great baseline or ref-
 121 erence solution to attentional learning. When using a Luce choice rule (e.g., a variant of
 122 a softmax rule), the cross-entropy loss function for a single trial is simply the negative
 123 log likelihood of making the correct categorization decision on that trial [12]. Hence,
 124 we can specify $\text{loss}(\boldsymbol{\alpha}) = -\log(P(\text{correct}))$ and rewrite Equation 1 as

$$\alpha_{t+1,j} = \alpha_{t,j} + \gamma_0 \frac{\partial}{\partial \alpha_{t,j}} \log(P(\text{correct})). \quad (2)$$

125 Then, to specify how attention should adjust with each new experience, we compute the
 126 partial derivative of the cross-entropy loss function in the context of the extended GCM
 127 in Equation 2; we provide this analytic derivation in the Supplementary Materials.

128 Equation 1 generally, and Equation 2 specifically, detail the concepts of “land-
 129 scape” and “path” of attention discussed in the Introduction. Here, landscape is de-
 130 fined by the cross-entropy loss function, whereas the path of attention is defined by
 131 incremental adjustments of attention (i.e., $\boldsymbol{\alpha}_t$). Assuming a learner is well intentioned,
 132 attention will orient to the location on the landscape that will maximize their learning
 133 objectives.¹ In the next section, we discuss a set of mechanisms that distort the cross-
 134 entropy gradient toward simpler representations and constraints of attentional capacity.
 135 These mechanisms can either affect the landscape a learner traverses or the path on
 136 which they orient attention.

137 1.3 Mechanisms of Selective Attention and Constraint

138 Equation 2 defines an “unconstrained” learner because it defines a path for attention that
 139 will maximize accuracy with no direct constraints on $\boldsymbol{\alpha}_t$, which is an assumption that
 140 is used in models such as ALCOVE and SUSTAIN. Although unconstrained attention
 141 works well in optimizing learning problems, it may fail to describe important aspects

¹We note that the current version of the model only maximizes learning objectives for the current trial, and does not make forward computations as in the recently developed SEA model [4].

142 of human behavior, such as our limited attentional capacity. In the next sections, we
 143 discuss a few algorithms that specify constraints on α_t . Throughout these sections, we
 144 will use the unconstrained learner in Equation 2 as a baseline model, and will modify
 145 Equation 2 to impose various constraints on attention.

146 1.3.1 Sum-To-Constant Constraint

147 Across contemporary models of category learning, it has become commonplace to im-
 148 pose constraints on attention. In light of work demonstrating that attention is a limited-
 149 capacity resource [see 13, for review], many models assume a sum-to-constant con-
 150 straint on the total quantity of attention that is distributed across dimensions. Based
 151 on seminal animal discrimination learning work by Sutherland and Mackintosh [14],
 152 GCM was, to our knowledge, the first model of human category learning to assume
 153 $\sum_j \alpha_j = k$, where k is assigned to a value of one for ease of interpretation. When $k = 1$,
 154 the sensitivity parameter δ in GCM interacts with the attention vector to determine the
 155 total attention $\sum_j \delta \alpha_j$. Because $\delta \alpha_j$ is the measure of sensitivity for Dimension j , the
 156 total attention ($\sum_j \delta \alpha_j$) informs us about the total sensitivity to differences across all
 157 dimensions, and so we will make use of this quantity when comparing across different
 158 attentional constraints.

159 Despite its ubiquity among existing models and its roots in irrefutable evidence
 160 of limited-capacity attention, the sum-to-constant assumption has critical implications
 161 that deserve reconsideration. First, it assumes that attention is used up to its limit on
 162 every single trial, regardless of task demands, condition-level difficulty, or changes in
 163 response efficiency over the course of learning. In other words, participants are ex-
 164 pected to expend the maximum possible quantity of attention at all times, regardless
 165 of how easy or difficult a particular categorization decision is perceived to be. Sec-
 166 ond, the sum-to-constant constraint naturally imposes competitive inhibition between
 167 dimensions, such that increasing attention to one dimension necessitates a reciprocal

168 decrease in attention to the others. Although evidence of attentional inhibition has been
 169 observed in a variety of contexts, it may be overly constraining to assume that inhibi-
 170 tion necessitates a net-zero change in total attention. Despite being an early pioneer of
 171 the sum-to-constant constraint in models of learning, subsequent work by Mackintosh
 172 [15], for example, showed that increasing attention to one dimension did not neces-
 173 sarily inhibit learning of the others as would be expected from the sum-to-constant
 174 assumption.

175 1.3.2 Norm-to-Constant Constraint

176 In subsequent variations of ALCOVE, such as the EXIT model [3, 16], a more general
 177 form of attentional constraint was considered that we refer to as the norm-to-constant
 178 constraint. The norm-to-constant relies on the extension of the geometric definition of
 179 distance, such that the p -norm of a vector is defined as

$$\|\boldsymbol{\alpha}\|_p = \left(\sum_j \alpha_j^p \right)^{1/p}, \quad (3)$$

180 where different values of p produce different types of distances (e.g., $p = 2$ creates
 181 the Euclidean distance whereas $p = 1$ creates the Manhattan distance). The norm-
 182 to-constant constraint implies that the p -norm must always equal a constant value,
 183 such that $\|\boldsymbol{\alpha}_p\| = k$. In EXIT, k was set equal to one in a manner consistent with the
 184 traditional sum-to-constant constraint, which is just a special case of this more general
 185 constraint (i.e., $p = 1$ in the sum-to-constant constraint).

186 Defining distance with the p -norm provides some additional flexibility over the
 187 set of possible attention values and can be treated as a free parameter [3, 16]. Kr-
 188 uschke [3] originally referred to this constraint as competition, because as the size of
 189 p as decreases, the competition between dimensions increases. In the extreme case
 190 where $p = 0$ (and $k = 1$), the learner can only attend to a single dimension. As such,
 191 when $0 < p < 1$, the model is a dimension reducer, consistent with a computation

that imposes simplicity, but this simplicity manifests through extreme competition. To elaborate, when $k = 1$, the maximal total attention is δ , which can only be achieved when attention is the constant k for a single dimension, but is set to 0 for all other dimensions. This setting biases learners to attend to fewer dimensions when improving accuracy, even in the case where multiple dimensions are predictive of the category. When $p = 1$, simplicity is not algorithmically instantiated because total attention is always maximized (i.e., $\sum_j \alpha_j = 1$), but competition still exists through normalization. Finally, when $p > 1$, less competition is imposed such that the maximal total attention is achieved when attending equally to all possible dimensions (e.g., when $p = \infty$). Because these ranges of parameters ($0 < p < 1$, $p = 1$, and $1 < p < \infty$) instantiate different computational goals as we have defined them, we will treat these three parameter ranges as three separate possible constraints on attention.

Figure 1 provides an illustration of how sum-to-constant and norm-to-constant constrain the set of available paths on the landscape of learning problems. Figure 1a illustrates the landscape of an example learning problem, where two dimensions of information must be used to solve a categorization problem. Both dimensions are important, but Dimension 2 is slightly more diagnostic in its relationship to the category label than is Dimension 1. If the learner has a goal of maximizing accuracy, the landscape would reflect this goal by making regions consistent with the goal more valuable (yellow areas). If there are no additional goals considered or constraints imposed, a learner can easily adjust their attention to the region that maximizes the value within their landscape (i.e., the red circle), which would simultaneously maximize categorization accuracy (i.e., the red triangle).

However, the norm-to-constant algorithm imposes constraints on the path of values that attention can take during learning. The imposition imbues the algorithm with a way to instantiate competition, and when $p < 1$, simplicity, but in doing so, the algorithm reduces the space of possible attentional states within the landscape. Starting with

Figure 1: **Mechanisms of Attention.** (a) In a two-dimensional learning problem, when there are no constraints on attention, a learner can reach an attentional state that maximizes their goal (i.e., the optimum of the learner’s objective function, red circle) of achieving highest accuracy (i.e., the statistically global optimum of the learning problem, red triangle). (b) Five sets of solutions are shown with $p = 1$ as the diagonal line, $p < 1$ in the lower triangle, and $p > 1$ in the upper triangle. In each set, although the maximum of the learner’s goal can be achieved within the set, the maximum may deviate from the statistical optimum depending on the value of p . (c) When a learner is trying to satisfy two learning goals simultaneously, such as (1) maximize categorization accuracy, and (2) minimize the total number of attended dimensions, the attentional state that maximizes the learner’s goals (red circle; optimum) departs from the attentional state that would maximize accuracy (red triangle). This type of attentional constraint is called LASSO regularization, where the purple color gradient in the bottom panel shows the optimum attentional state of several learning goals that vary along the importance of the second goal’s contribution to the learning problem. (d) Ridge regularization can cause the best attentional state for the learner’s goals to deviate from the statistical optimum, but this deviation does not necessarily reduce the total number of dimensions attended, just the amount of attention applied to each dimension. (e) Effects of competition in a “rule-based” task where only Dimension 2 is relevant to the categorization problem. (f) “Information integration” task where both dimensions must be used to achieve high categorization accuracy. In both examples, the most relevant dimension attracts attention, causing a deviation of the path of attention for the competitive model (dashed line) relative to the baseline model (solid white line). See the Supplementary Video Files for videos corresponding to each of these illustrations.

219 the standard sum-to-constant rule (i.e., $p = 1$), Figure 1b illustrates the set of possible
220 attentional states as the diagonal line overlaid on the original landscape of the learning
221 problem. Although the learner still orients their attention to the region that maximizes
222 their goal (i.e., they are “well intentioned”), this region (i.e., the red circle) may deviate
223 from the single goal of maximizing accuracy (i.e., the statistical optimum shown as the
224 red triangle). As examples, Figure 1b shows the subsets of attention that can be used
225 under different values of p . The lower triangle shows two subsets where $p < 1$, whereas
226 the upper triangle shows two subsets where $p > 1$. Changing p alters the set of possible
227 solutions by warping the diagonal line in the sum-to-constant rule. This warping of
228 the attentional states allows for attentional dynamics that instantiate computations that
229 are competitive or dimension-reducing (when $p < 1$). However, both dynamics are
230 controlled by a single parameter p .

231 In the norm-to-constant specification when p is reasonably small, shifting attention
232 is necessarily a zero-sum game: increasing attention to one dimension implies a re-
233 duction in at least one other dimension. However, this may not be the most accurate
234 description of human behavior. Indeed, an extensive literature has shown that atten-
235 tion clearly fluctuates as a consequence of trial difficulty [17–19], learning-related ef-
236 ficiency [20–22], and fatigue [23–25]. The results from these studies hint at a possible
237 theoretical shortcoming: if the total amount of attention (i.e., the sum) has been shown
238 to fluctuate according to experimental or biological properties, then perhaps attention
239 should not be constrained to a constant value.

240 To examine the nature of attention allocation in detail, we offer an alternative
241 framework to the norm-to-constant algorithm. Specifically, we propose to decouple
242 the computations of competition and simplicity by replacing the strict constraint used
243 within the norm-to-constant algorithm with an alternative constraint we refer to as *reg-*
244 *ularization*. In addition, to instantiate our second computation of interest, we specify
245 competition through lateral inhibition of stimulus dimensions to allow for flexible nav-

246 igation of the attentional landscape. We next describe both of these alternative algo-
 247 rithms in turn.

248 1.3.3 Theoretical Alternative: Regularization

249 In statistics, regularization is the addition of a bias term to an objective function (e.g.,
 250 cross-entropy loss) that reduces complexity. In standard regression models, when a co-
 251 efficient approaches zero, the corresponding predictor variable will have less of an ef-
 252 fect on the model’s predictions. Regularization methods such as LASSO [26] or Ridge
 253 [27] effectively bias regression coefficients for linear models toward zero. Therefore,
 254 regularization works as a feature selection method, remedies issues of multicollinear-
 255 ity in dimensions, and instantiates the computation of simplicity because only the most
 256 predictive variables will be able to overcome the bias toward zero imposed by regular-
 257 ization.

258 Here we consider two forms of regularization due to their connections to exist-
 259 ing theories in the cognitive modeling literature: LASSO and Ridge. The analogy of
 260 these algorithms to human learning is that they impose different attentional constraints
 261 for achieving simplicity. Whereas LASSO regularization instantiates simplicity by re-
 262 ducing the number of attended stimulus dimensions, Ridge regularization instantiates
 263 simplicity by reducing the total *amount* of attention. As derived in the Supplementary
 264 Materials, to apply LASSO regularization to α_t , Equation 2 is adjusted to

$$\alpha_{t+1,j} = \alpha_{t,j} + \gamma_0 \frac{\partial}{\partial \alpha_{t,j}} \log(P(\text{correct})) - \gamma_0 \lambda, \quad (4)$$

265 where λ determines the LASSO regularization bias on α_t . Here, α_t has a constant bias
 266 toward zero, and this bias affects the attention vector in proportion to each dimension’s
 267 diagnosticity: when a dimension is not particularly predictive of a category label, reg-
 268 ularization will bias the corresponding element of α_t toward zero, and consequently,
 269 that dimension will have less influence on subsequent categorization decisions. As

such, LASSO regularization has a clear cognitive interpretation: it is a bias for low-dimensional (simplified) representations of the category structure.

Ridge regularization [27, 28] applies a penalty that is proportional to the value of the coefficient itself. Interestingly, this type of regularization is already common in dynamic theories of decision making, such as in Decision Field Theory [29] and the Leaky Competing Accumulator model [30]. As we show in the Supplementary Materials, to incorporate Ridge regularization, Equation 4 need only be adjusted to the following:

$$\alpha_{t+1,j} = \alpha_{t,j} + \gamma_0 \frac{\partial}{\partial \alpha_{t,j}} \log(P(\text{correct}) - \kappa \alpha_{t,j}), \quad (5)$$

where $\kappa \in [0, 1]$. In theories of decision making, the parameter κ is often referred to as “leakage” because it allows information that has been accumulated to passively decay away if new evidence is not acquired. In contrast to LASSO, Ridge regularization does not necessarily converge to a simplified low-dimensional representation of category structure, but will still significantly bias the total amount of attention toward zero, producing less total attention compared to the unconstrained learner.

With a regularization bias in place, a dimension must help accomplish the learner’s goals in order for a dimension to continue attracting attention (i.e., the gradient must exceed the regularization penalty λ or $\kappa \alpha_{t,j}$). As a consequence, regularization could potentially facilitate learning in the context of high-dimensional stimuli by mitigating the “curse of dimensionality.” Furthermore, LASSO and Ridge regularization impose important constraints on the attention vector related to finite capacity views of attention. Specifically, the LASSO penalty sets an upper bound on the total attention vector such that $\sum_j \alpha_{t,j} < k$, whereas the Ridge penalty sets an upper bound on the squared sum of the attention vector such that $\sum_j \alpha_{t,j}^2 < k$. The upper bound value k is a complex function of the penalty term (i.e., either λ or κ) and the data, but conceptually it suffices to say that both regularization methods impose a capacity bias on total attention, rather than a strict constraint as in the norm-to-constant assumption discussed above.

Figure 1c and Figure 1d illustrate the effects of LASSO and Ridge regularization, respectively, on the illustrative categorization problem introduced above. In both panels, regularization imposes simplicity by trying to minimize the total attention used during learning. When combined with the initial goal of maximizing accuracy, the goal of minimizing effort distorts the landscape of the learning problem by increasing the values associated with regions in which the total attention vectors are shorter in length (i.e., nearer toward the origin). In both panels, when regularization is applied, the attentional state that maximizes the learner’s goals (i.e., the red circle) departs from the attentional state that would maximize the single goal of maximizing categorization accuracy, and the magnitude of this departure is proportional to the regularization parameter.

To visualize this relationship more generally, the purple color gradient in each panel shows the location of the attentional state that maximizes the learner’s goals (i.e., maximizing accuracy while keeping attention small). In both panels, by increasing the regularization bias (i.e., by decreasing the regularization parameter), the attentional state that maximizes the learner’s goals is nearer to the origin, reflecting a decrease in the total amount of allocated attention. However, an important difference between LASSO and Ridge regularization is the path of these optimal attentional states: a LASSO regularization prioritizes the reduction of dimensions of information (i.e., a low-dimensional solution), whereas a Ridge regularization prioritizes a reduction in the total length of the attentional vector.

1.3.4 Theoretical Alternative: Competition

When learning which dimensions predict category labels, one efficient strategy would be to allow knowledge about one dimension to influence attention allocation to other dimensions. As we will show, this encourages a more direct path of attention orientation. By itself, Equation 2 only allows attention to orient toward a dimension in an

independent fashion, with no interactions among dimensions. If attention capacities are finite [31, 32], removing attention from the less relevant dimensions and reallocating it to the more relevant dimensions would be a prudent investment in terms of attention deployment.

When attention allocation is mathematically specified in this interactive way, the dynamic is often mechanistically referred to as competition in the attention literature [33, 34], and is implemented in various ways in theories of cognition [7, 11, 29, 30, 35]. To instantiate competition within our current framework, we modify Equation 2 to

$$\alpha_{t+1,j} = \alpha_{t,j} + \gamma_0 \frac{\partial}{\partial \alpha_{t,j}} \log(P(\text{correct})) - \beta \sum_{k \neq j} \frac{\partial}{\partial \alpha_{t,k}} \log(P(\text{correct})).$$

where the inhibitory parameter β is positive and determines the strength of inhibition and excitation (see Supplementary Material for more details). The effect of this mechanism is to accelerate the degree to which more predictive dimensions attract attention and less predictive dimensions repel attention. Increasing β allows for more competitive interactions across dimensions, which can lead to an acceleration in prioritizing dimensions [7, 8]. In contrast to the norm-to-constant rule discussed above, this form of competition operates on the stimulus dimensions themselves rather than on the sum of attention over dimensions.

Figures 1e and 1f illustrate how the competition mechanism enables shortcuts when allocating attention in two pedagogical examples. Figure 1e illustrates the landscape of a classic “rule-based” task where only Dimension 2 is relevant for a categorization decision, whereas Figure 1f shows the landscape of an “information integration” task where both dimensions are relevant. In both cases, the unconstrained baseline model correctly orients attention to the global optimum, maximizing the learner’s goal of

340 high accuracy. By contrast, the competitive mechanism accelerates the orientation of
341 the baseline model, creating a more direct path to the optimum that will allow accuracy
342 to increase more quickly during learning.

343 **2 Results**

344 We present the results in three sections. First, to illustrate the effects of different at-
345 tentional constraints on learning, we simulated accuracy and attention data from the
346 models using the classic Shepard et al. [36] experimental designs. Second, we fit the
347 models to four data sets using rule-based and information integration designs with ei-
348 ther two or four response alternatives. These first four data sets were chosen as “bench-
349 mark” empirical data as they involve standard category learning problems, but also have
350 the compelling advantage of using eye tracking as a direct measure of overt attention.
351 Third, we fit the models to a more challenging experimental design from Mack et al.
352 [37], which involves within-participant changes in the type (i.e., difficulty) of the cat-
353 egorization rule used. In both empirical applications, we fit all the models discussed
354 above to provide evidence not only for the computational principles of simplicity and
355 competition, but also evidence for the specific algorithms used to instantiate those prin-
356 ciples.

357 **2.1 Shepard et al. (1961) Simulation**

358 To illustrate the consequences of different attentional constraints on learning, we sim-
359 ulated models under different configurations of a classic theoretical benchmark: the
360 stimuli from Shepard et al. [36]. Figure 2a shows the set of eight stimuli, which vary
361 along three binary dimensions. Shepard and colleagues [36] used these stimuli to create
362 a set of six experimental designs by varying the complexity of the feature-to-category
363 map. The design numbers reflect category complexity, where increasing numbers sug-
364 gest increasing complexity. Figure 2a illustrates the designs by color coding nodes of

365 the cube according to which category the stimuli in the left panel are assigned.

366 To illustrate the effects of attentional constraints, we simulated three models in each
367 of the six stimulus types: a model unconstrained by attention (see Equation 2), a model
368 constrained by LASSO regularization (see Equation 4) and competition, and a model
369 constrained by the norm-to-constant assumption (see Equation 3).

370 The left panels of Figures 2b-d show the average accuracy values for the uncon-
371 strained (b), regularized (c), and sum-to-constant (d) models for each stimulus set.
372 Although the unconstrained model reaches an asymptotic level of accuracy faster than
373 the regularized model, they eventually asymptote at similar values (i.e., near perfect
374 performance). The sum-to-constant model's accuracy increases quickly and asymp-
375 totes. Because the model had the same sensitivity parameter (i.e., δ) across stimulus
376 conditions, increases in category complexity produce decreases in categorization accu-
377 racy.

378 The right panels of Figures 2b-d show the sum of attention across all three dimen-
379 sions for each condition. As expected, the unconstrained attention model continues
380 to increase its attention over time, even when accuracy is maximized (left panel). By
381 contrast, the model with regularized attention increases its attention in each condition
382 until arriving at an asymptotic level, a level that increases with increasing stimulus
383 complexity. Finally, the sum-to-constant model produces equal *total* attention in each
384 task by construction. More generally, other norm-to-constant models with $p \neq 1$ will
385 also produce attentional vectors that will sum to a constant c if the distance of the at-
386 tentional vector is calculated with a Minkowski p -metric. The differences in the total
387 attention profiles among the unconstrained, regularized, and norm-to-constant models
388 will be of critical importance when discriminating among the models in the final em-
389 pirical study which uses a subset of these Shepard designs to alter the complexity of
390 the feature-to-category map over blocks of the experiment within a participant.

Figure 2: Results of Theoretical Benchmark Simulation. (a) The stimuli used in the simulation consisted of three stimulus dimensions (left) and two category responses. Six different experimental designs were constructed through unique feature-to-category maps of increasing difficulty (right). (b) The choice accuracy (left) and total amount of attention applied (right) for each of the six simulated experimental designs when no attentional constraints are applied (see Equation 2). (c) The choice accuracy (left) and total amount of attention applied (right) for each of the six simulated experimental designs with a model assuming both LASSO regularization and competition. (d) The choice accuracy (left) and total amount of attention applied (right) for each of the six simulated experimental designs for the sum-to-constant model (i.e., the norm-to-constant model with $p = 1$) over time. Total attention is defined as $\sum_j \delta \alpha_{t,j}$, where $\delta = 1$ for regularized and unconstrained models. In each panel b-d, curves are color coded according to the category structure key in Panel (a). In the right panel (d), the total attention was equivalent across conditions, however we separated the lines for visual clarity.

391 2.2 Empirical Data: Benchmark Data Sets

392 We now look to data to provide evidence for each alternative constraint. We begin with
393 a set of four empirical “benchmarks” containing choice and eye tracking measures over
394 time in four standard experimental paradigms. In all experiments, stimuli consisted of
395 three, continuously valued features within dimensions that were spatially separated,
396 such that fixations to dimensions is a direct measure of overt attention. Because both
397 eye tracking and choice data provide measures of attention allocation over time, the
398 purpose of these first analyses is to establish whether or not learning data can indeed
399 be characterized as an attention optimization problem.

400 For generalizability, we selected four data sets that vary on three important char-
401 acteristics. The first characteristic is the relevance of dimensions; the data sets we
402 selected either have one or two relevant dimensions. The second characteristic is the
403 complexity of the category structure. For this characteristic, the data sets we selected
404 either have a “rule-based” structure (Figure 1c), or an “information integration” struc-
405 ture (Figure 1f). Rule-based category structures can be described by rules that are easily
406 verbalized, whereas information integration category structures are optimally learned
407 by integrating feature values across dimensions [38] in a manner difficult to verbalize
408 without equations (e.g., a weighted linear combination). The third characteristic is the
409 number of categories; data sets we selected either have two or four response options.
410 Each of the four previously-published data sets contained choice and eye-tracking data,
411 and were made freely available by McColeman et al. [39]. In the Methods and Mate-
412 rials section, we describe the main details of each data set, but the reader is directed to
413 McColeman et al. [39] for additional details.

414 2.2.1 Switchboard Analysis: Assessing Relative Model Performance

415 To examine which computations and specific algorithms were likely to explain human
416 learning, we performed a “switchboard analysis” by factorially testing every combina-

Figure 3: **Switchboard Analysis.** (a) Model diagram representing the organizational structure of each class of models. On the left, the models are organized into either existing (top) or alternative candidate mechanisms (bottom). On the right, the models are organized via a Venn diagram, illustrating how models can either have computational goals of simplicity, competition, both, or neither. C=Competition, R=Ridge, L=LASSO, U=Unconstrained, and P = p -norm to constant (b) Each panel shows the relative model performance using the Akaike Information Criterion (AIC) for different data sets, where the number of alternatives is organized by rows and the type of decision rule is organized by columns. The AIC is rescaled within an experiment such that values reflect units of AIC relative to the best performing model.

tion of mechanism within each model structure [7, 40–43]. Using the unconstrained model as a baseline, we systematically added every combination of attentional constraint and competition to create a total of nine model variants. In the norm-to-constant model variants, because competition was inherent to the model structure, we only tested variants with $p < 1$, $p = 1$, and $p > 1$, creating three model variations, as these models have different implications for the amount of dimension reduction and competition they impose. Regarding the new algorithms, we crossed LASSO, Ridge, and no regularization with competition via lateral inhibition to create a total of six variations, and we refer to this set of nine models as the Adaptive Attention Representation Model (AARM). Figure 3a represents each model as a node organized either as a lattice (left) or as a Venn Diagram (right). We fit all nine models separately to data from 117 participants in total across the four benchmark experiments (i.e., 1,053 model fits). In the Supplementary Materials we also crossed each of these nine model variants with different assumptions about the temporal profile of memory, but as these results did not produce interesting differences in the relative performances of the models, we only report a subset of the results here.

Figure 3b shows the relative model performance for each of the nine models in terms of the Akaike information criterion (AIC) with a small sample size correction [44], color coded according to the legend in Figure 3a. In an absolute sense, the $p < 1$ model performed best in the two-alternative tasks, whereas the competitive LASSO model performed best in the four-alternative tasks. Both of these models performed well overall, indicating that each can capture the basic learning profiles in these benchmark studies. By contrast, the $p > 1$ model and the $p = 1$ performed the worst overall, and models with only Ridge or Lasso performed worse than their competitive counterparts. Finally, the unconstrained model variants also performed poorly, performing worse than models that added either competition or any type of regularization except in a single case (i.e., in the two alternative information integration task, the unconstrained

444 model performed worse than the Ridge regularization model).

445 **2.2.2 Absolute Fits to Data**

446 Having determined which models performed best in a relative sense, we also wished to
447 verify that the the competitive LASSO model also performed well in an absolute sense.
448 Figure 4 shows different aspects of the model fits to each of the data sets (rows) for
449 the competitive LASSO model. Although both the $p < 1$ and the competitive LASSO
450 models performed well across all data sets relative to the other model variants, because
451 the competitive LASSO is the central development of our article and it ultimately per-
452 forms best across all studies, we show only the competitive LASSO fits to data here.
453 The Supplementary Materials provides a congruent plot for the $p < 1$ model.

454 Figures 4a, 4e, 4i, and 4m show the models fits to the accuracy data for each data
455 set. Although there is a slight under-performance of the model in the rule-based two-
456 alternative data, the model generally captures the increase and asymptote of the accu-
457 racy data well. Figures 4b, 4f, 4j, and 4n show the model’s fits to the fixation propor-
458 tion data for each task for the relevant (red) and irrelevant (blue) dimensions over time.
459 Here, the competitive LASSO model precisely captures the pattern of attention orienta-
460 tion over time. Figures 4c, 4g, 4k and 4o show the latent attention curves predicted by
461 the model for relevant (red) and irrelevant (blue) dimensions over time, with the total
462 attention (i.e., the sum of attention over all dimensions) shown as the shaded yellow
463 region. Here, the model shows that a steady increase in attention to the relevant dimen-
464 sion over time and a rapid decrease in attention to the irrelevant dimension. Figures 4d,
465 4h, 4i, and 4p show the distribution of stimuli used in the task, color coded according
466 to the category label for each stimulus.

Figure 4: **Fits to the Empirical Benchmark Data Sets.** (a, e, i, m) Aggregated (across participants) choice accuracy in the data (black lines; gray area designates the 95% confidence interval of the mean) and model fits (green lines) from the competitive LASSO model variant. (b, f, j, n) Fixation proportions for the relevant (red) and irrelevant (blue) dimensions, where the dark lines correspond to the observed mean fixation proportion while the lighter color lines represent the mean predictions from the winning model. In each panel, the shaded bands represent the 99% confidence intervals of the observed fixation proportions' means. (c, g, k, o) The median predictions of the latent attention vector for the winning model. The red line corresponds to the relevant dimensions, whereas the blue line corresponds to the irrelevant dimensions. The total amount of attention (i.e., the sum of the attention vector across all dimensions) is shown as the yellow line. (d, h, l, p) Three dimensional scatter-plot of the stimuli used in each task, color coded by category membership.

467 **2.2.3 Summary**

468 After fitting the models to data we found that the norm-to-constant model with $p <$
469 1 and the regularization model with competition performed best across the four data
470 sets. Although these models contain different algorithms for allocating and orienting
471 attention, they each are derived on the same basic computational goals. Namely, they
472 both attempt to create parsimony by reducing the total number of attended dimensions,
473 and they both impose a form of competition among the dimensions of information.
474 Because both model variants perform well across the four data sets, we take this as
475 converging evidence for the presence of both simplicity and competition. However,
476 because the experimental designs were fairly standard, we unfortunately were not able
477 to discriminate among these two specific algorithms. In the next study, we attempt such
478 a discrimination to further refine our results.

479 **2.3 Algorithmic-level Distinctions: Data from Mack et al. (2016)**

480 The simulation study above (see Figure 2) revealed an interesting prediction about
481 the attentional profiles across the Shepard et al. [36] designs between the regulariza-
482 tion and sum-to-constant models: as the task difficulty increased, the regularization
483 models increased the total amount of attention whereas sum-to-constant models were
484 constrained to apply the same amount of attention throughout. However, this result
485 occurs with a fixed degree of complexity in the categorization rule over time. What
486 were to happen if the complexity of the task increased, such that more dimensions be-
487 came necessary to complete the categorization task accurately? We might envision two
488 competing hypotheses. First, one could imagine that humans use the same amount of
489 attention on each trial, regardless of the complexity of each stimulus. Although the
490 attention vector could move and orient (e.g., along the line in Figure 1b) to different
491 weighting of dimensions, the total amount of attention would be constant. Second,
492 one could imagine that when the task is simple, not much attention would be needed

493 to perform the task well, but if the task were to suddenly increase in complexity, one
494 would need to increase attention to the task in order to maintain some desired level of
495 accuracy.

496 In our framework, the first hypothesis corresponds to the sum-to-constant algo-
497 rithm, whereas the second hypothesis corresponds to the regularization algorithm. The
498 relationship becomes complicated for the norm-to-constant model because the total
499 amount of attention can increase or decrease as attention orients², but the basic intuition
500 that the constraint on attention creates a restricted path on the attentional landscape still
501 holds. We hypothesized that if data in which the complexity of the categorization rule
502 changed from time to time, those data could be used to provide further discrimina-
503 tion among the models at the algorithmic level. Luckily, such data were collected by
504 Mack et al. [37], and so we now turn to these data in an attempt to provide further
505 evidence about the specific nature of the interaction between simplicity and dimension
506 competition in human learning.

507 2.3.1 Data Overview

508 Although we describe the high-level details of the experiment here, we direct the reader
509 to the full description in presented in [37] for additional details. The dataset for this
510 study is freely available via the Open Science Foundation <https://osf.io/5byhb/>.
511 The dataset contains MRI and behavioral data from 23 right-handed participants (12
512 males, age 18-31 years) with normal or corrected-to-normal vision who completed a
513 category learning task using classic stimuli from Shepard et al. [36]. Participants com-
514 pleted 12 runs of the learning task in an fMRI scanner, which included four runs each
515 using stimuli from categorization designs of Shepard Types 1, 2, and 6. During the
516 learning task, participants viewed insect stimuli with three binary features (legs, anten-
517 nae, and mouths) and were asked to indicate which of two groups each insect belonged

²The complication is that total attention is defined as the length of the attention vector using a Minkowski distance of $p = 1$, whereas the attentional constraint can be any value of p in the norm-to-constant model.

to by pressing a button. Importantly, the stimulus remained while corrective feedback was provided after each response. The learning task contained three phases of four runs each, where each phase employed a different classification rule. For all participants, Phase 1 used the Type 6 classification rule. The classification rule order (Type 1 or Type 2) for Phases 2 and 3 was counterbalanced across participants. Participants were informed that a rule change occurred at the beginning of Phases 2 and 3. Each of the eight unique stimuli were presented four times per run for a total of 32 trials per run.

2.3.2 Switchboard Analysis: Relative Fits

After fitting each of the nine models to the data, we compared the model fits using the AIC metric with the small sample size correction [44]. Figure 5c shows the relative model performance color coded according to the legend on the right-hand side. Here, the figure shows that the competitive Ridge model performs best, with the competitive LASSO model as a close second. Models that were unconstrained, or had competition or LASSO performed next best, followed by Ridge alone. The norm-to-constant models performed the worst in the set, with the $p < 1$ model performing worst of all. If we sum up the AIC across all five experiments, the competitive LASSO model wins by a considerable margin (AIC=231322.92), with the next best model having an AIC 1178 greater ($p < 1$; AIC=232501.416). When we compute the AIC-based conditional probabilities [45] for all nine models, the competitive LASSO model's probability is numerically 1. Therefore, taking a holistic view of all five data sets, the competitive LASSO model is the best representation of the data of the models we considered in the information-theoretic sense.

2.3.3 Absolute Fits: Comparing Norm-to-constant and Regularization

Having checked the relative model fits, we also examined whether the worst performing $p < 1$ model had clear, qualitative differences from the best (across all studies)

543 performing competitive LASSO model. Both models qualitatively performed simi-
 544 larly for the condition where the categorization Type I rule was taught before the Type
 545 II. However, the models make drastically different predictions when Type II is taught
 546 before Type I. Figure 5a shows the latent attention curves for the $p < 1$ model, and Fig-
 547 ure 5b shows the latent attention curves for the competitive LASSO model over time.
 548 Both models predict the proper ordering of attention allocation: for Type 6 stimuli, all
 549 dimensions are prioritized; for Type 2 stimuli, only the two relevant dimensions are
 550 prioritized (blue and purple lines); for Type 1 stimuli, only the one relevant dimension
 551 is prioritized (green lines). However, interestingly, the total amount of attention differs
 552 considerably across the models, with the $p < 1$ model predicting that there was more
 553 total attention applied to Type 1 stimuli, and the competitive LASSO model predict-
 554 ing a clear ranking of total attention that is in line with the complexity of the stimulus
 555 types: Type 6 is greatest, Type 2 is second, and Type 1 is the least. These differ-
 556 ences materialize in the accuracy curves over time, shown in Figure 5d for the $p < 1$
 557 model and Figure 5e for the competitive LASSO model. Here, whereas the competitive
 558 LASSO model closely tracks the accuracy profiles across time, the $p < 1$ model shows
 559 a consistent under prediction of the accuracy curve across all trials. This misprediction
 560 comes directly from the type of attentional constraint: because the $p < 1$ model can
 561 only increase attention through orientation (i.e., by reducing at least one dimension), it
 562 is unable to increase attention freely in the most complex Type 6 task. In effect, when
 563 fitting the model the parameter estimate for p converges to a value that maximizes the
 564 fit over all the data, but this fit is compromised in some conditions (e.g., Type 6 in
 565 this case). By contrast, because the competitive LASSO model is able to flexibly in-
 566 crease the total amount of attention when it is needed (e.g., when the task increases in
 567 difficulty), it can easily capture the accuracy profiles of human subjects in this task.

Figure 5: Model Evaluation for the Mack et al. Data (a, b) Latent attention curves for each stimulus dimension (green, blue, and purple) for the $p < 1$ (a) and the competitive LASSO (b) models across time. In the background of each panel, the total sum of attention is shown as the shaded yellow region. In each panel, the shaded region designates the type of Shepard stimuli used during the task (i.e., either a Type 6, 2 or 1). (d, e) Aggregated (across participants) choice accuracy in the data (black lines; blue shaded area designates the 95% confidence interval of the mean) and model fits (green lines) for the $p < 1$ (d) and competitive LASSO (e) model variants. (c) The relative model performance as measured by the Akaike Information Criterion is shown for each model variant, color coded according to the key on the right hand side: C=Competition, R=Ridge, L=LASSO, U=Unconstrained, and P = p -norm to constant. Each model's AIC is rescaled to reflect units of AIC relative to the best-performing model.

568 3 Discussion and Conclusions

569 In this article, we defined learning as an attention-optimization problem: to learn the
570 best action to take (e.g., choice) given dimensions of information, we must deploy at-
571 tention in proportion to how well we believe those dimensions will help us in achieving
572 our goals. Here, our approach was to assume that each learner was “well intentioned”,
573 meaning that they would always orient attention in the direction that maximized their
574 learning goals. Then, we used the classic Generalized Context Model [GCM; 2] as the
575 basis for defining how category structure could be learned, which created a baseline
576 landscape relating attention to categorization accuracy. We then specified that addi-
577 tional computational goals such as reducing dimensions might exist, and we defined
578 these goals with respect to attention itself. We specified a few different algorithms that
579 would behave in ways that were consistent with these overarching goals, fit models
580 equipped with those algorithms to data from five different studies, and provided evi-
581 dence for the relative fidelity of each computation as well as the specific algorithms for
582 accomplishing those goals.

583 The first four studies were considered “benchmark” studies because they were
584 based on fairly simple and classic categorization problems, which crossed rule-based
585 and information integration rules against two- and four-alternative response environ-
586 ments. A compelling advantage of these previously published data sets was that they
587 also provided explicit measures of attention through eye tracking. By fitting both
588 choice and eye tracking data simultaneously, we were able to provide strong evidence
589 that attention orientation involves both competition and some type of simplicity, where
590 fewer dimensions were preferred. This conclusion was drawn because models that
591 used competition and reduced the number of stimulus dimensions provided superior
592 fits to all four data sets relative to every other combination of model mechanisms.
593 The norm-to-constant model provided good fits to the data when $p < 1$, supporting
594 both dimension-reducing and competitive dynamics. The competitive LASSO regular-

595 ization model also provided good fits to the data, providing converging evidence for
 596 dimension-reducing and competitive dynamics within a different algorithm. As we ex-
 597 plained (also see Figure 1), the norm-to-constant model can be thought of as defining
 598 a specific subset of attentional states that are nested within the possible values defined
 599 by the general form of the alternative framework presented here. When regularization
 600 is applied, a bias toward zero is instantiated, resulting in behavior that is similar to the
 601 $p < 1$ constraint within the norm-to-constant model. For the first four data sets, the
 602 profile of attention was consistent with both model representations producing similar
 603 absolute fits to the data, but the norm-to-constant model could achieve good fits with
 604 fewer parameters, and won in some cases after penalties for model complexity were
 605 applied.

606 To discriminate between the norm-to-constant model with $p < 1$ and the LASSO
 607 regularization model with competition, we fit the models to data from Mack et al. [37].
 608 In this task, participants had to allocate attention to either three dimensions (i.e., Type
 609 6 stimuli), two dimensions (i.e., Type 2 stimuli), or one dimension (i.e., Type 1 stimuli)
 610 in different blocks of the experiment; importantly, this manipulation was performed
 611 *within* participants. We argued that this type of difficulty manipulation [see 36, and
 612 our Simulation Study] was critical because it provided indirect evidence for the path
 613 of the attention vector across blocks through the accuracy measures. There are at least
 614 two hypotheses for how attention could be distributed across blocks of the experiment.
 615 First, consistent with the norm-to-constant algorithm, participants could orient their
 616 attention in a manner consistent with the stimulus type, but could not increase attention
 617 without reducing sensitivity to at least one dimension. Second, consistent with the
 618 mechanism of regularization, participants could reflexively allocate more attention in
 619 conditions where the task required attention to more dimensions, and simply relax
 620 their attention in easier conditions. After fitting the models to data, we found both
 621 qualitative and quantitative evidence that the norm-to-constant model with $p < 1$ was

622 unable to capture the profile of attention (gleaned through accuracy data) over blocks
623 of the experiment. By contrast, the LASSO regularization model with competition was
624 able to amplify attention when the task became more difficult, and withdraw attention
625 when the task was easier in a manner that was consistent with human participants.

626 In summary, our results provide strong evidence that humans exhibit simplicity
627 through a tendency toward low-dimensional explanations when learning about the world,
628 and importantly, these dimensions compete for attention. Fixation and choice data pro-
629 vided signatures of these learning computations (reducing attention, and dimension
630 competition through inhibition) across five tasks. We instantiated these computations
631 by using the previously developed theoretical algorithm of norm-to-constant, as well as
632 developing novel mechanisms of regularization and competition. We then used a more
633 complex design from Mack et al. [37] to directly test the specific algorithms identified
634 within the first four studies and found evidence against the hypothesis that attention
635 normalizes to a constant value. Evidently, a persistent bias toward a parsimonious so-
636 lution is always present in human learning, but when the task demands more attention,
637 we are able to flexibly allocate attention so as to increase our performance. Although
638 both regularization and competition can at times provide inferior learning solutions
639 when considering accuracy compared to a learner with no constraints, regularization
640 and competition more accurately describe human profiles of attention, highlighting the
641 multifaceted goals used during learning that arise from the bounded nature of our cog-
642 nitive capabilities. However, we argue that these bounds are largely beneficial as they
643 provide humans with a way to parsimoniously represent and adapt to a vastly complex
644 and ever-changing world.

4 Methods and Materials

4.1 The Adaptive Attention Representation Model

When learning about categories, the observer’s task is to assign a D -dimensional stimulus vector $\mathbf{e}_t$ on the t^{th} trial to one of C categories. To model the learning process, we take an adaptive approach [5–7, 46] where new episodic traces are added to the representation after each experience, a procedure consistent with the instance theory of automatization [47–49]. To do this, we assume that after each experience with a stimulus $\mathbf{e}_t$, an episodic trace $\mathbf{x}_i = [x_{i,1} \ x_{i,1} \ \dots \ x_{i,D}]$ (i.e., an exemplar) is stored. Associated with this exemplar is a memory salience weight $m_{t,i}$, and feedback about the true category membership $f_i \in \{1, 2, \dots, C\}$. These entities are all contained within matrices that evolve over time. On Trial t , the stored exemplars are contained within $\mathbf{X}_t = [\mathbf{x}_1 \ \dots \ \mathbf{x}_N]^\top$, the memory saliences are contained within $\mathbf{M}_t = [m_{t,1} \ m_{t,2} \ \dots \ m_{t,N}]^\top$, and the feedback associated with each experience is contained within $\mathbf{F}_t = [f_1 \ \dots \ f_N]$. These matrices expand with each new experience (e.g. trial), and together they form the representation by which a feature-to-category map is constructed. When a new stimulus is presented, the stimulus activates the set of stored exemplars based on their similarity to the stimulus. The GCM assumes that similarity is computed by way of a factorizable exponential similarity kernel [2, 50], such that activation $a_{t,i}$ of the i^{th} exemplar in response to the stimulus $\mathbf{e}_t$ on Trial t is

$$a_{t,i} = \exp \left(-\delta \sum_{j=1}^D \alpha_{t,j} |e_{t,j} - x_{i,j}| \right) m_{t,i} \quad (6)$$

where δ is the specificity of the similarity kernel function, and $\alpha_{t,j}$ is the attention applied to the j th dimension of information on Trial t . The adaptive attention components of the model are discussed at a high-level in the introduction, and the technical details are provided in Section 2 of the Supplementary Materials, as different assumptions about these processes can instantiate different computational goals (e.g., maximizing

669 accuracy with resource constraints).

670 Once the exemplars are activated, a Luce choice rule is used to compute category
 671 activation in the context of all stored exemplars. Specifically, the probability of making
 672 a Category c response is

$$P("c"|\boldsymbol{\alpha}_t, \mathbf{e}_t, \mathbf{F}_t, \mathbf{X}_t, \mathbf{M}_t) = \frac{\sum_{i=1}^N a_{t,i} \mathbb{I}_{[f_i=c]}}{\sum_{i=1}^N a_{t,i}}, \quad (7)$$

673 where $\boldsymbol{\alpha}_t = [\alpha_{t,1} \ \alpha_{t,2} \ \dots \ \alpha_{t,D}]$ and $\mathbb{I}_{[f_i=c]}$ is an indicator function that returns a one if
 674 the i th exemplar $\mathbf{x}_i$ has an associated category label (e.g., from feedback) of Category
 675 c :

$$\mathbb{I}_{[f_i=c]} = \begin{cases} 1 & f_i = c \\ 0 & \text{otherwise} \end{cases}$$

676 Thus the probability of choosing c is the summed similarity of the exemplars associated
 677 with the c -th category, normalized by the total activation of all exemplars.

678 The adaptive GCM formalization above allows for learning to occur in two ways.
 679 First, learning can occur by simply adding summaries (i.e., episodic traces) of new
 680 experiences into the memory matrix $\mathbf{X}_t$ because when making choices, the activations
 681 of past memories are aggregated according to a summed similarity principle (Equation
 682 7). Second, learning can occur even when no new memories are formed by reallocating
 683 attention to dimensions that are most relevant or diagnostic to the categorization task
 684 via Equation 2. With larger values of $\alpha_{t,j}$, the psychological distance stretches in the j th
 685 dimension, creating better discriminability in that dimension. By contrast, decreases
 686 in attention to a dimension create contraction and hence less discrimination in that
 687 dimension. For example, if $\alpha_{t,j} = 0$ for the j -th dimension, that dimension will have
 688 no influence when computing distance and hence will not activate exemplars based on
 689 the information in the j th dimension.

690 4.2 Functional Form of Temporal Memory Bias

691 Category learning models typically often assume a memory component [7, 51]. To
692 build memory into the model, we used a simple weighting function that can incorporate
693 primacy and recency biases in the memory salience [52]:

$$m_{t,i} = [1 - (1 - \epsilon_p^i)(1 - \epsilon_r^{N_t - i + 1})](1 - \eta) + \eta, \quad (8)$$

694 where ϵ_p and $\epsilon_r \in [0, 1]$ are primacy and recency weighting parameters, respectively,
695 and $\eta \in [0, 1]$ is a lower bound parameter for memory weights, and N_t is the number
696 of exemplars stored on Trial t . For simplicity, we present only models that assume a
697 single memory bias term subject to the constraint that $\epsilon_p = \epsilon_r$ when fitting models to the
698 benchmark data sets. In the Supplementary Materials, we report model fits that cross
699 the set of models discussed here with different configurations of recency and primacy.
700 For the data from [37], we freely estimated both ϵ_p and ϵ_r for all models.

701 Data Availability

702 The choice and eye-tracking data that support the analyses presented in *Empirical*
703 *Data: Benchmark Data Sets* are available in the original publication [39] in PLoS
704 One with the identifier DOI: 10.1371/journal.pone.0083302.s007. The choice data that
705 support the analyses presented in *Algorithmic-level Distinctions: Data from Mack et*
706 *al. (2016)* are hosted by the Open Science Foundation and is discoverable with the
707 identifier: osf.io/5byhb.

708 Code Availability

709 The code used to complete the analysis presented here will be available upon publica-
710 tion at <https://github.com/MbCN-lab>.

References

- [1] Marr, D. *A computational investigation into the human representation and processing of visual information* (WH San Francisco: Freeman and Company, 1982).
- [2] Nosofsky, R. M. Attention, similarity, and the identification–categorization relationship. *Journal of experimental psychology: General* **115**, 39 (1986).
- [3] Kruschke, J. K. Toward a unified model of attention in associative learning. *Journal of mathematical psychology* **45**, 812–863 (2001).
- [4] Braunlich, K. & Love, B. C. Bidirectional influences of information sampling and concept learning. *Psychological Review* (2021).
- [5] Nosofsky, R. M. & Alfonso-Reese, L. A. Effects of similarity and practice on speeded classification response times and accuracies: Further tests of an exemplar-retrieval model. *Memory and Cognition* **27**, 78–93 (1999).
- [6] Turner, B. M., Van Zandt, T. & Brown, S. D. A dynamic, stimulus-driven model of signal detection. *Psychological Review* **118**, 583–613 (2011).
- [7] Turner, B. M. Toward a common representational framework for adaptation. *Psychological review* **126**, 660 (2019).
- [8] Rehder, B. & Hoffman, A. B. Eyetracking and selective attention in category learning. *Cognitive psychology* **51**, 1–41 (2005).
- [9] Rehder, B. & Hoffman, A. B. Thirty-something categorization results explained: selective attention, eyetracking, and models of category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition* **31**, 811 (2005).
- [10] Kruschke, J. K. ALCOVE: an exemplar-based connectionist model of category learning. *Psychological review* **99**, 22 (1992).
- [11] Love, B. C., Medin, D. & Gureckis, T. M. SUSTAIN: a network model of category learning. *Psychological review* **111**, 309 (2004).

- 736 [12] Goodfellow, I., Bengio, Y., Courville, A. & Bengio, Y. *Deep learning* **2** (MIT
737 press Cambridge, 2016).
- 738 [13] Chun, M, Golomb, J & Turk-Browne, N. A taxonomy of external and internal
739 attention. *Annual Review of Psychology* **62**, 73–101 (2011).
- 740 [14] Sutherland, N & Mackintosh, N. *Mechanisms of animal discrimination learning*
741 (Academic Press, 1971).
- 742 [15] Mackintosh, N. A theory of attention: Variations in the associability of stimuli
743 with reinforcement. *Psychological Review* **82**, 276 (1975).
- 744 [16] Paskewitz, S. & Jones, M. Dissecting EXIT. *Journal of Mathematical Psychol-*
745 *ogy* **97**, 102371 (2020).
- 746 [17] Lavie, N & Tsal, Y. Perceptual load as a major determinant of the locus of
747 selection in visual space. *Perception & Psychophysics* **56**, 183–197 (1994).
- 748 [18] Lavie, N. Perceptual load as a necessary condition for selective attention. *Jour-*
749 *nal of Experimental Psychology: Human Perception and Performance* **21**, 451
750 (1995).
- 751 [19] Lavie, N & Cox, S. On the efficiency of visual selective attention: Efficient visual
752 search leads to inefficient distractor rejection. *Psychological Science* **8**, 395–396
753 (1997).
- 754 [20] Warm, J, Parasuraman, R & Matthews, G. Vigilance requires hard mental work
755 and is stressful. *Human Factors* **50**, 433–441 (2008).
- 756 [21] Awh, E, Belopolsky, A & Theeuwes, J. Top-down versus bottom-up attentional
757 control: A failed theoretical dichotomy. *Trends in Cognitive Sciences* **16**, 437–
758 443 (2012).
- 759 [22] Turner, B *et al.* *Cognitive inertia: How loops among attention, representation,*
760 *and decision making distort reality* DOI: 10.31234/osf.io/8zvey. 2021.

- 761 [23] Smallwood, J & Schooler, J. The restless mind. *Psychological Bulletin* **132**, 946
762 (2006).
- 763 [24] Mittner, M, Boekel, W, Tucker, A & Turner, B. When the brain takes a break: A
764 model-based analysis of mind wandering. *Journal of Neuroscience* **34**, 16286–
765 16295 (2014).
- 766 [25] Turner, B, Van Maanen, L & Forstmann, B. Informing cognitive abstractions
767 through neuroimaging: The neural drift diffusion model. *Psychological Review*
768 **122**, 312 (2015).
- 769 [26] Tibshirani, R. Regression shrinkage and selection via the lasso. *Journal of the*
770 *Royal Statistical Society: Series B (Methodological)* **58**, 267–288 (1996).
- 771 [27] Hoerl, A. E. & Kennard, R. W. Ridge regression: Biased estimation for nonorthog-
772 onal problems. *Technometrics* **12**, 55–67 (1970).
- 773 [28] Park, T. & Casella, G. The bayesian lasso. *Journal of the American Statistical*
774 *Association* **103**, 681–686 (2008).
- 775 [29] Busemeyer, J. R. & Townsend, J. T. Decision field theory: a dynamic-cognitive
776 approach to decision making in an uncertain environment. *Psychological review*
777 **100**, 432 (1993).
- 778 [30] Usher, M. & McClelland, J. L. The time course of perceptual choice: the leaky,
779 competing accumulator model. *Psychological review* **108**, 550 (2001).
- 780 [31] Kahneman, D. *Attention and effort* (Citeseer, 1973).
- 781 [32] Lavie, N. & Tsal, Y. Perceptual load as a major determinant of the locus of
782 selection in visual attention. *Perception & psychophysics* **56**, 183–197 (1994).
- 783 [33] Klein R. M .and Taylor, T. L. Categories of cognitive inhibition with reference
784 to attention. (1994).
- 785 [34] Klein, R. M. Inhibition of return. *Trends in cognitive sciences* **4**, 138–147 (2000).

- 786 [35] Sederberg, P. B., Howard, M. W. & Kahana, M. J. A context-based theory of
787 recency and contiguity in free recall. *Psychological review* **115**, 893 (2008).
- 788 [36] Shepard, R. N., Hovland, C. I. & Jenkins, H. M. Learning and memorization of
789 classifications. *Psychological monographs: General and applied* **75**, 1 (1961).
- 790 [37] Mack, M. L., Love, B. C. & Preston, A. R. Dynamic updating of hippocampal
791 object representations reflects new conceptual knowledge. *Proceedings of the*
792 *National Academy of Sciences* **113**, 13203–13208 (2016).
- 793 [38] Ashby, F. G. & Gott, R. E. Decision rules in the perception and categorization of
794 multidimensional stimuli. *Journal of Experimental Psychology: Learning, Mem-*
795 *ory, and Cognition* **14**, 33 (1988).
- 796 [39] McColeman, C. M. *et al.* Learning-induced changes in attentional allocation
797 during categorization: a sizable catalog of attention change as measured by eye
798 movements. *PloS one* **9**, e83302 (2014).
- 799 [40] Van den Berg, R., Awh, E. & Ma, W. Factorial comparison of working memory
800 models. *Psychological Review* **121**, 124–149 (2014).
- 801 [41] Heathcote, A., Loft, S. & Remington, R. W. Slow down and remember to re-
802 member! A delay theory of prospective memory costs. *Psychological Review*
803 **122**, 376–410 (2015).
- 804 [42] Turner, B. M. *et al.* On the neural and mechanistic bases of self-control. *Cerebral*
805 *Cortex*, 1–19 (2018).
- 806 [43] Turner, B. M., Schley, D. R., Muller, C. & Tsetsos, K. Competing theories of
807 multialternative, multiattribute preferential choice. *Psychological review* **125**,
808 329 (2018).
- 809 [44] Burnham, K. P. & Anderson, D. R. A practical information-theoretic approach.
810 *Model selection and multimodel inference* **2** (2002).

- 811 [45] Wagenmakers, E.-J. & Farrell, S. AIC model selection using Akaike weights.
812 *Psychonomic bulletin & review* **11**, 192–196 (2004).
- 813 [46] Turner, B. M. & Van Zandt, T. Hierarchical approximate Bayesian computation.
814 *Psychometrika* **79**, 185–209 (2014).
- 815 [47] Logan, G. D. Toward an instance theory of automatization. *Psychological Re-*
816 *view* **95**, 492–527 (1988).
- 817 [48] Logan, G. D. Shapes of reaction-time distributions and shapes of learning curves:
818 A test of the instance theory of automaticity. *Journal of Experimental Psychol-*
819 *ogy: Learning, Memory, and Cognition* **18**, 883–914 (1992).
- 820 [49] Logan, G. D. An instance theory of attention and memory. *Psychological Review*
821 **109**, 376–400 (2002).
- 822 [50] Shepard, R. N. Toward a universal law of generalization for psychological sci-
823 ence. *Science* **237**, 1317–1323 (1987).
- 824 [51] Nosofsky, R. M. & Alfonso-Reese, L. A. Effects of similarity and practice on
825 speeded classification response times and accuracies: Further tests of an exemplar-
826 retrieval model. *Memory & Cognition* **27**, 78–93 (1999).
- 827 [52] Pooley, J. P., Lee, M. D. & Shankle, W. R. Understanding memory impairment
828 with memory models and hierarchical Bayesian analysis. *Journal of Mathemat-*
829 *ical Psychology* **55**, 47–56 (2011).

830 Acknowledgements

831 This research was supported by an NSF CAREER award (B.M.T.) and NIH grant
832 R01HD078545 (V.M.S.). We would like to thank our colleagues for the construc-
833 tive feedback on this manuscript; in particular, we thank Alexandra Greenberg, Nate
834 Blanco, Daniel Evans, and Giwon Bahg.

835 Author Contributions

836 M.G. and B.M.T. designed the analysis and developed new analytic tools. M.G. ana-
837 lyzed the data. All authors interpreted results and wrote the paper.

838 Competing Interests

839 The authors report no competing interests.

840 Materials and Correspondence

841 Correspondence regarding this article and associated materials should be addressed to
842 B.M.T.