

On the Complexity of Bayesian Generalization

Yu-Zhe Shi^{*1} Manjie Xu^{*1,2} John E. Hopcroft³ Kun He⁴ Joshua B. Tenenbaum⁵
Song-Chun Zhu^{1,2} Ying Nian Wu⁶ Wenjuan Han⁷ Yixin Zhu¹

Abstract

We examine concept generalization at a large scale in the natural visual spectrum. Established computational modes (*i.e.*, rule-based or similarity-based) are primarily studied isolated, focusing on confined and abstract problem spaces. In this work, we study these two modes when the *problem space* scales up and when the *complexity* of concepts becomes diverse. At the **representational level**, we investigate how the complexity varies when a visual concept is mapped to the representation space. Prior literature has shown that two types of complexities (Griffiths & Tenenbaum, 2003) build an inverted-U relation (Donderi, 2006; Sun & Firestone, 2021). Leveraging Representativeness of Attribute (RoA), we computationally confirm: Models use attributes with high RoA to describe visual concepts, and the description length falls in an inverted-U relation with the increment in visual complexity. At the **computational level**, we examine how the complexity of representation affects the shift between the rule- and similarity-based generalization. We hypothesize that category-conditioned visual modeling estimates the co-occurrence frequency between visual and categorical attributes, thus potentially serving as the prior for the natural visual world. Experimental results show that representations with relatively high subjective complexity outperform those with low subjective complexity in rule-based generalization, while the trend is the opposite in similarity-based generalization.

1. Introduction

What is a cucumber? One may respond by a *deep green colored slim-long cylinder with trichomes on the surface is a cucumber*, or directly pick a cucumber—*see, something*

^{*}Equal contribution ¹Peking University ²National Key Laboratory of General Artificial Intelligence, BIGAI ³Cornell University ⁴Huazhong University of Science and Technology ⁵MIT ⁶UCLA ⁷Beijing Jiaotong University. Correspondence to: Yu-Zhe Shi <syz@perslab.co>, Yixin Zhu <yixin.zhu@pku.edu.cn>.

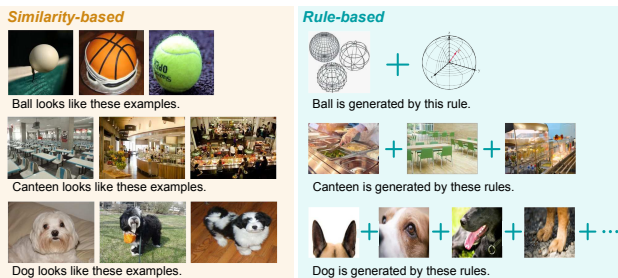


Figure 1. Concepts can be described either directly by similar examples or indirectly by a set of related rules. Here, we demonstrate this intuition through the concepts of *ball*, *canteen*, and *dog*.

looks like this is a cucumber. Given either answer as prior knowledge, you can easily identify cucumbers; you may check whether a candidate meets the rules described or judge whether it is visually similar to the known example. Such capability is concept generalization, and the approaches used to identify cucumber are rule- and similarity-based generalization (Sloman & Rips, 1998; Shepard, 1987), respectively.

Can both approaches always be applied to all concepts we see? Hardly. Let us consider how people learn to identify *dog*, *canteen*, and *ball*. We can easily capture the main feature of a dog given very few examples, yet we may get confused with the complex rules to identify a dog. Conversely, people may easily tell the limited rules that shape the concept of canteens, such as serving windows, tables, and chairs; and the concept of the ball is such a simple concept that can be easily captured by examples or a single rule. We refer the readers to Fig. 1 for an illustration. This observation naturally leads to a hypothesis that *whether people generalize through rules or similarity* has something to do with *how complex the concept instances look like*. “Look like,” complexity, and generalization—these three elements shape the hypothesis. In this work, we systematically look into these dimensions of concept generalization.

We contextualize the problem in the literature of concept generalization—the framework of Bayesian inference unifies rule- and similarity-based generalization (Tenenbaum, 1999; 1998; Tenenbaum & Griffiths, 2001a; Xu & Tenenbaum, 2007). Based on perception (Kersten et al., 2004), this paradigm reconstructs human’s hypothesis space consisting of abstract features and incubates modern concept learning algorithms (Tenenbaum et al., 2011; Lake et al., 2015; Ellis et al., 2020). However, as most concept learners have

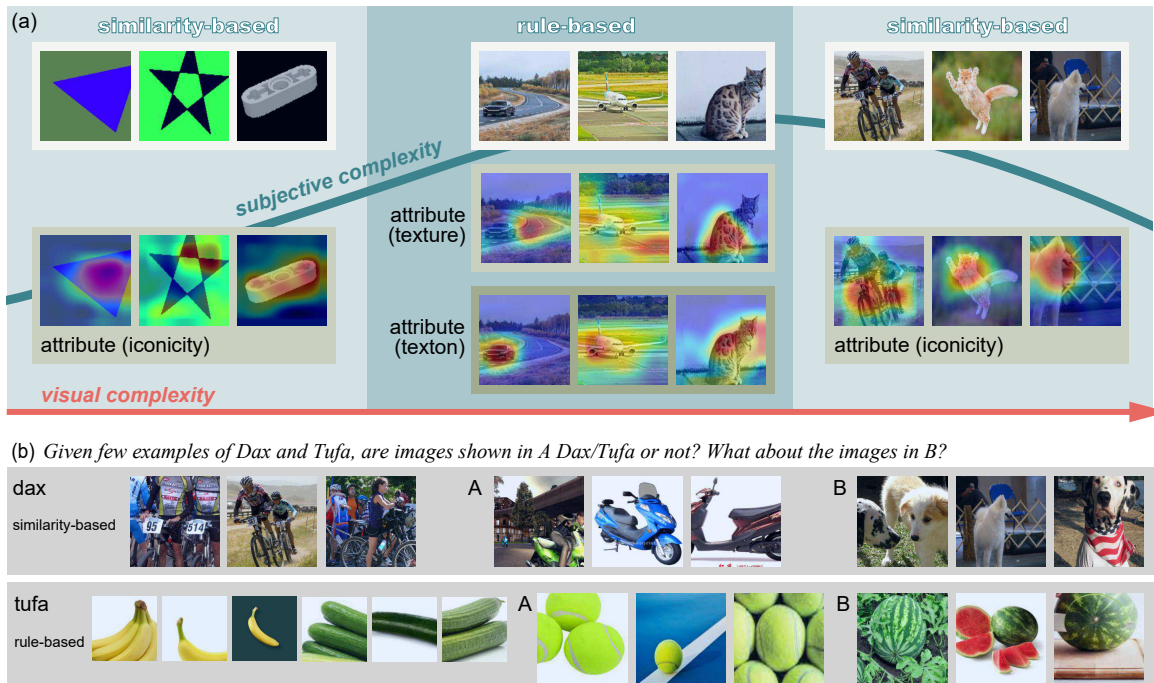


Figure 2. **The landscape of the computation-mode-shift vs. the concept complexity.** (a) *Representation level*: original visual concepts of diverse complexity and visualization of their representative attributes (around the peaks of heatmaps). (b) *Computation level*: an illustration of similarity- and rule-based generalization. The former is similar to word learning (Xu & Tenenbaum, 2007; Jiang et al., 2023): Given very few examples of known concept *dax*, tell which is most likely to be *dax* in unseen examples. The latter is akin to concept learning (Salakhutdinov et al., 2012; Zhang et al., 2019a): Given a rule *tufa* over two known concepts, tell how *tufa* generates the examples of unknown concepts. As **concept visual complexity** increases, **concept subjective complexity** first increases, then decreases—and the computation mode shifts from similarity to rules as **subjective complexity** increases.

only demonstrated in confined and abstract problem space, a challenging problem remains: When the problem space *scales up* (e.g., using data collected from the natural world), is there a unified concept representation that combines the two established modes (i.e., rule- and similarity-based)? If it does, how does the generalization shift between the two modes w.r.t. the *complexity* of concepts?

One hypothesis (Donderi, 2006; Wolfram et al., 2002) is that we tend to describe visual concepts (i.e., **visual complexity**, the complexity coded by pixels) by simple visual patterns with explicit semantics (i.e., **subjective complexity**, the coding length for describing certain concepts). For a simple concept, we may only need one attribute; e.g., the shape *circle* for *ball*. As the concepts become more complex, we adopt more attributes, such as *canteens are rooms with serving windows, tables, and chairs*. When the concept becomes even more complicated, we would choose *not* to describe it—the description would be much too long to be appropriate for communication if we describe it with the attributes generated by the complex rules. Hence, we capture the main feature and view it as an “icon” for the concept, such as *dog looks like dog*.

Together, we observe a shift in the continuous space spanned by the rule- and similarity-based approaches w.r.t. the in-

crease of concept complexity. Intuitively, both very simple and complex concepts have a lower description length, generalized by similarity. Conversely, concepts neither too simple nor too complex have a higher description length, generalized by rules. This observation echos modern literature in information theory and psychology, demonstrating that subjective and visual complexity (Griffiths & Tenenbaum, 2003) come in an inverted-U relation (Donderi, 2006; Sun & Firestone, 2021).

In essence, we seek to quantify the relation between the prior-studied but mostly isolated modes (i.e., rule- and similarity-based): What are the *relations* between the computation-mode-shift and the concept complexity, illustrated also in Fig. 2? Specifically, we disassemble the above question into two on the basis of Marr’s (Marr, 1982) *representational level* and *computational level*, respectively: (i) How does the complexity change when a visual concept is mapped to the representation space? (ii) How does the complexity of representation affect the shift between rule- and similarity-based generalization? By answering these two questions, we hope to provide a new perspective and the very first pieces of evidence on unifying the two computational modes by mapping out the landscape of the concept complexity vs. the computation mode.

Representation vs. complexity Representing the natural visual world merely with human prior is insufficient (Griffiths et al., 2016) and oftentimes brittle to generalize. Despite that hierarchy empowers large-scale Bayesian word learning (Miller, 1998; Abbott et al., 2012), extending it to visual domains is still challenging (Jiang et al., 2023). In comparison, modern discriminative models trained for visual categorization by leveraging large-scale datasets can capture the rich concept of attributes (Xie et al., 2020). These observations and progresses naturally lead to the problem of concept representation complexity: If we distinguish visual concepts using attributes, at least how many attributes should we use (Chaitin, 1977; Li et al., 2008)?

To tackle this problem, we bridge the subjective complexity with the visual complexity via Representativeness of Attribute (RoA), consisting of (i) the probability of recalling an attribute z when referring to a concept c , and (ii) the probability of recalling other concepts \hat{c} when referring to an attribute z . This design echoes the principles in rational analysis (Tenenbaum & Griffiths, 2001b) yet can be obtained by frequentist statistics for large problem spaces (e.g., natural visual world (Abbott et al., 2011)).

Computation vs. complexity Modern statistical learning methods have demonstrated strong expressiveness in concept representation by implicitly calculating the co-occurrence frequency between visual attributes and categories (Wu et al., 2019; Xie et al., 2016), even when scaling up to the complex and large-scale visual domain—the learned representation fits the prior distribution of visual concepts conditioned on categorical description (Xie et al., 2020). It can also bridge sensory-derived and language-derived knowledge (Bi, 2021). Hence, this learning paradigm should somehow have inherent semantic properties in addition to visual properties, such as iconicity (Fay et al., 2010; 2013; Qiu et al., 2022) and disentanglement (Allen & Hospedales, 2019; Gittens et al., 2017; Mikolov et al., 2013; Edmonds et al., 2019).

To properly evaluate the computation, we extend the problem domain from generalization over single concepts to that across multiple concepts. This is because, in the natural visual world, we cannot precisely answer how a concept is generated by rules or which examples are sufficient to represent a concept. Hence, instead of considering the absolute measurements for single concepts, we consider the relative measures between concepts; for example, *cucumber to banana is watermelon to what*, or *dog is more similar to cat or to bike*—only the significant differences are considered. We argue that *rule-based* and *similarity-based* generalization reflects the *analogy* and *similarity* properties in psycholinguistics (Gentner & Markman, 1997; Zhang et al., 2019b; 2021b; 2022; Holyoak et al., 2022), where the former pairs are two ends of a continuum of concept

representation, and the latter pairs are two ends of a continuum of literal meaning. Visual categorization brings these two pairs together because linguistic analogy and similarity come from generalizing the corresponding appearance instead of pure literal meaning—concepts with more easy-to-disentangle attributes (e.g., shape and color) are more likely to generalize by rules, while concepts represented with more iconicity (Fay et al., 2014) (i.e., those more likely to be viewed holistically) tend to generalize by similarity.

Computationally, the above hypothesis is consistent with the findings by Wu et al. (Wu et al., 2008). Specifically in the natural visual scenes, textons (low-entropy) (Zhu et al., 2005) can be composed by very simple concepts (Wu et al., 2010), akin to rule-based generalization. In comparison, textures (high-entropy) (Julesz, 1962) cannot be represented by rules (Zhu et al., 1997); instead, they are evaluated and generalized in terms of similarity by “pursuit” (Zhu et al., 1998). As such, we hypothesize that generalization shifts from similarity to rules as subjective complexity increases.

In the remainder of the paper, we first present the new metrics, Representativeness of Attribute (RoA), to measure the subjective complexity and analyze the computation-mode-shift in Sec. 2. Next, through a series of experiments, we provide strong evidence to support our hypotheses in Sec. 3; we have two primary **findings** in response to the two problems raised at the beginning: (i) **Representation:** the subjective complexity significantly falls in an inverted-U relation with the increment of visual complexity. (ii) **Computation:** rule-based generalization is significantly positively correlated with the subjective complexity of the representation, while the trend is the opposite for similarity-based generalization.

2. Bayesian generalization and complexities

In this section, we formulate Bayesian generalization for visual concept learning (Sec. 2.1), followed by the definitions of subjective complexity and visual complexity (Sec. 2.2).

2.1. Bayesian generalization for large-scale visual concept learning

Concept-conditional modeling Let us consider $f : \mathbb{R}^D \mapsto \mathbb{R}^d$, which maps the input $\mathbf{x} \in \mathbb{R}^D$ to a representation vector $\mathbf{z} \in \mathbb{R}^d$. Here, f might be part of a discriminative model trained for visual categorization tasks, such as a prefix for a convolutional neural network without the last fully-connected layer for mapping \mathbf{z} to the category vector $\mathbf{c} \in \mathbb{R}^c$; \mathbf{z} is a collection of independent attributes $\{z_1, z_2, \dots, z_d\}$ where d is the dimensionality of the attribute space. These attributes are not one-hot encoders but are relaxed to normalized weights. Without loss of generality, we assume all samples share the same attribute space \mathcal{Z} . Training a discriminator for image categorization is to esti-

mate the likelihood of concept c given a set of samples X : $P(c|X) = \prod_{x \in X} P(c|x)$. Here, we assume that f provides a good estimation of $P(z|X; \theta)$, where θ is the parameter of f ; Tishby *et al.* (Tishby & Zaslavsky, 2015) provides empirical evidence that a discriminative model may first learn how to extract proper attributes to model images X conditioned on c , then learn to discriminate their categories based on the attribute distribution. Some dimensions of \mathbf{z} (usually 5% ~ 10% of the total dimensions) capture concrete semantic attributes of visual concepts when the activation score $f_z(X)$ is relatively high (Bau *et al.*, 2020). Combining this concept-conditional measurement with attribute modeling, we rewrite the category prediction considering the attribute as a latent variable z and marginalize the observable joint distribution (X, c) over z :

$$P(c|X) = \sum_{z \in \mathcal{Z}} P(c, z|X; \theta) = \sum_{z \in \mathcal{Z}} P(c|z)P(z|X; \theta), \quad (1)$$

where c is assumed to be conditionally independent of X given z , which is naturally consistent with the definition of attributes for visual concepts—concepts are all identified by attributes whether holistically or partially. This expression is essentially a Bayesian prediction view of visual categorization, which can be derived to Bayesian generalization in the natural visual world.

Representativeness of Attribute (RoA) as an informative prior Statistically, we treat the concept-conditional attribute activation score as an estimation of the probability $P(z|c)$ that recalls an attribute z when referring to a concept c , similar to answering “Describe how a dog looks like.” In the context of the natural visual world, we also have all activation scores generated by an attribute as an estimation of the probability $P(\hat{c}|z)$ that recalls all concepts $\hat{c} \neq c$ when referring to the attribute z , akin to answering “What do you recall when seeing a blue thing in ball shape?”

Formally, we define the RoA of a specific attribute z_i for concept c (Tenenbaum & Griffiths, 2001b):

$$\text{RoA}(z_i, c) = \log \frac{P(z_i|c)}{\sum_{\hat{c} \neq c} P(\hat{c})P(z_i|\hat{c})}, \quad (2)$$

where $P(\hat{c})$ is the prior of concepts in the context. We hypothesize that humans estimate $P(\hat{c})$ through both language derivation and visual experience, essentially calculating the co-occurrence between visual attributes and categorical attributes over the joint distribution $P(z_i, \hat{c})$. Hence, modeling RoA with large-scale image datasets and language corpora should yield human-level prior modeling. On this basis, we use f to statistically estimate $P(z_i, \hat{c})$ (Xie *et al.*, 2020):

$$\text{RoA}(z_i, c) = \log \frac{P(z_i|c)}{\sum_{\hat{c} \neq c} P(\hat{c})P(z_i|\hat{c})} \propto \log \frac{P(z_i|c; \theta)}{\sum_{\hat{c} \neq c} P(\hat{c}|z_i; \theta)}, \quad (3)$$

where $P(z|c; \theta)$ and $P(\hat{c}|z; \theta)$ are estimations of $P(z_i|c)$ and $P(z_i|\hat{c})$, respectively. Please refer to Appendix B.2 for additional details about implementing RoA.

Generalize to the unseen Given an appropriate modeling of $P(z|X; \theta)$, the goal is to generalize an unknown concept c' to a small set of unseen examples $\hat{X} = \{x_1, \dots, x_n\}$, where n tends to be a small integer. The generalization function $P(c'|\hat{X})$ is given by:

$$\begin{aligned} P(c'|\hat{X}) &= \sum_{z \in \mathcal{Z}} P(c'|z)P(z|\hat{X}; \theta) \\ &= \sum_{z \in \mathcal{Z}} \frac{P(c')P(z|c')}{\sum_{c \in \mathcal{C}} P(c)P(z|c)} P(z|\hat{X}; \theta) \\ &\propto \underbrace{P(c')}_{\text{uninformative prior}} \sum_{z \in \mathcal{Z}} \underbrace{\exp(\text{RoA}(z, c'))}_{\text{informative prior}} P(z|\hat{X}; \theta), \end{aligned} \quad (4)$$

where the uninformative prior $P(c')$ encodes the computation-mode-shift. Specifically, the similarity-based generalization $\langle c :: c' \rangle$ between a concept pair is defined as $\exists c \in \mathcal{C}, \sigma_0(c, c') < \delta$, where δ is a relative small neighbour. Similarly, the rule-based generalization $\langle c_1 : c_2 :: c_3 : c' \rangle$ over a quadruple of concepts is defined as $\exists c_1, c_2, c_3 \in \mathcal{C}, \sigma_1(c_1 - c_2, c_3 - c') < \delta$, where $\sigma_N(\cdot, \cdot)$ is an arbitrary metric measurement with an N-order input. Further, we define $P(c') \propto \sigma_0 \sigma_1 / (\sigma_0 + \sigma_1)$ (Tenenbaum, 1999), resulting in the simplest hypotheses of concepts: The harmonic property keeps guide to similarity-based generalization if σ_0 is dominating, and vice versa.

2.2. Complexities

Visual complexity

Visual concepts come with diverse complexity, from very simple geometry concepts such as *squares* and *triangles* to very complex natural concepts such as *dogs* and *cats*. Inspired by Wu *et al.* (2008), we indicate concept-wise visual complexity by Shannon’s information

entropy (Shannon, 1948). Formally, for a set of images $X = \{x_1, x_2, \dots\}$ belonging to a concept c , the concept-wise entropy is $H(X|c) = \mathbb{E}_{X \sim P(\cdot|c)}[\log P(X|c)]$. Fig. 3 computes the visual complexity of some image datasets: 2D geometries (El Korchi & Ghanou, 2020), single objects (Tatman, 2017), compositional-attribute objects (Zhang *et al.*, 2021a; Johnson *et al.*, 2017), human-made objects (Deng *et al.*, 2009), scenes (Zhou *et al.*, 2017), and animals (Xian *et al.*, 2018; Deng *et al.*, 2009).

Subjective complexity We quantify the subjective complexity over the prior model by Kolmogorov Complexity (Li *et al.*, 2008). We leverage the minimum description length, *i.e.*, the minimum number of attributes to

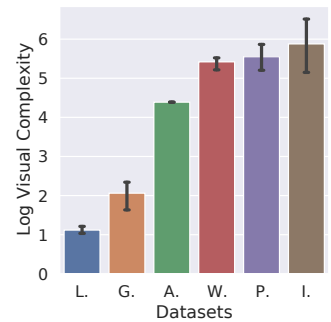


Figure 3. **Visual complexity of datasets, in increasing order.** L: LEGO, G: 2D-Geo, A: ACRE, W: AWA, P: Places, I: ImageNet.

discriminate a concept: For each concept c , we rank all attributes $z \in \mathcal{Z}$ by $\text{RoA}(z, c)$ decreasingly, such that $\forall i, j \in [1, d], i < j, \text{RoA}(z_i, c) \geq \text{RoA}(z_j, c)$. Starting from $K = 1$, we select the top- K attributes for each iteration and check whether these attributes can distinguish the concept c from the others. This process continues if the current iteration cannot distinguish it from the others. Formally, we define subjective complexity of visual concept $L(\hat{c})$ as:

$$L(\hat{c}) = \min_K \mathbb{1}(P(\hat{c} \neq c) < \epsilon | c = \arg \max_c P(c | z_1, \dots, z_K; \phi)), \quad (5)$$

where ϵ is the error rate threshold, and ϕ is the parameter of f 's suffix in the same discriminative model for visual categorization (e.g., the fully-connected layer). We calculate $P(c | z_1, \dots, z_K; \phi)$ by removing the neurons' effects corresponding to z_{K+1}, \dots, z_d (Bau et al., 2020). Instead of maintaining all error rate thresholds, we leverage the *accuracy gain* between every two iterations to search for the minimum K . Of note, although accuracy may indeed be affected by the size of the training set or training epochs, the metric of subjective complexity is intended to capture the essence of the attribute's representation in a concept, making it a stable measure of subjective complexity. Thus, our measurement is fair because the subjective complexity of the visual concept is influenced by ϵ , positively correlated to model accuracy. This process yields the concept-wise subjective complexity in RoA; see also Appendix B.3.

3. Results and analysis

This section provides results and analyses that validate the above hypotheses.¹ We conduct empirical analyses at both the *representation* (Sec. 3.1) and *computational* (Sec. 3.2) level, quantitatively analyze the computation-mode-shift w.r.t. the concept complexity in Sec. 3.2), and qualitatively interpret the results via natural image statistics in Sec. 3.3.

3.1. Representation vs. complexity

This experiment investigates the visual concepts' subjective complexity by visual categorization. Our predictions were that models use attributes with high RoA to describe visual concepts, and the description length falls in an inverted-U relation with the increment of visual complexity.

Method Six groups of discriminative models were trained from scratch on six datasets (see Fig. A2) with the supervision of concept labels: LEGO (Tatman, 2017), 2D-Geo (El Korchi & Ghanou, 2020), ACRE (Zhang et al., 2021a), AWA (Xian et al., 2018), Places (Zhou et al., 2017), and ImageNet (Deng et al., 2009), ordered as the increment of concept-wise visual complexity. All models were opti-

¹The experiment source code and the dataset information are available at <https://github.com/YuzheSHI/bayesian-generalization-complexity>.

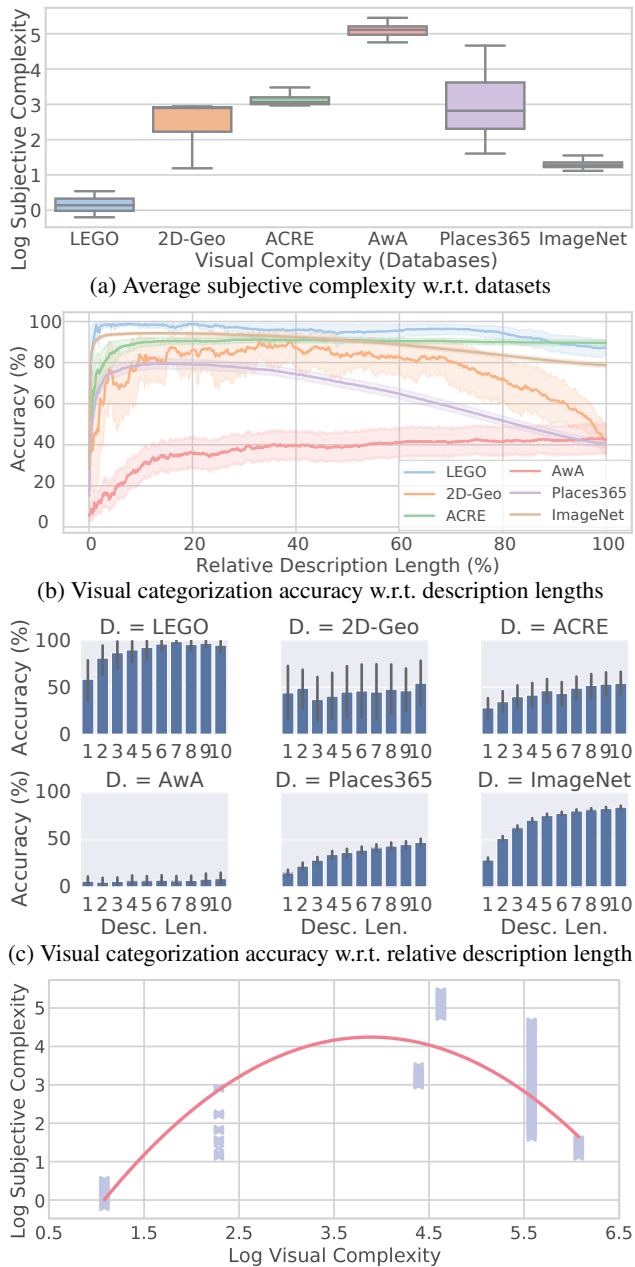


Figure 4. **Quantitative results of Representation vs. Complexity.**

mized to converge on the training set and tuned to the best hyper-parameters on the validation set. Please refer to Appendices B.1 and D for details about datasets and training.

During the evaluation, RoA was calculated for each attribute in the context of all concepts for each dataset. Following the protocol described in Sec. 2.2, the models were tasked with visual categorization, leveraging from only one attribute with the highest RoA to the entire attribute space.

Results The main quantitative results are illustrated in Fig. 4. Subjective complexity shows significant diversity between the datasets. The logarithm values are as follows;

see also Fig. 4a. LEGO: .10 ($CI = [-.10, .52]$, $p < .05$), 2D-Geo: 2.91 ($CI = [1.21, 2.95]$, $p < .05$), ACRE: 3.08 ($CI = [2.99, 3.46]$, $p < .05$), AWA: 5.08 ($CI = [4.82, 5.36]$, $p < .05$), Places: 2.74 ($CI = [1.63, 4.72]$, $p < .05$), and ImageNet: 1.28 ($CI = [1.16, 1.51]$, $p < .05$). All models’ performances plateau after leveraging $< 20\%$ of all attributes; see Fig. 4b. Most models (5 out of 6) exploit very few (less than 5% of all) attributes to reach a higher accuracy than that of all attributes; see Fig. 4b. The models for the simplest dataset (*i.e.*, LEGO) and the most complex dataset (*i.e.*, ImageNet) obtain a large accuracy gain (over 10%) with the description length from 1 to 4 and obtain smaller accuracy subsequently. In comparison, the models for ACRE and Places obtain relatively small accuracy gain (about 5%) with description length from 1 to 10; see Fig. 4c. Fig. 4d shows the estimated inverted-U relation between subjective complexity and visual complexity. Following the “two-lines” test (Simonsohn, 2018), the relation is relatively robust across the datasets, decomposing the non-monotonic relation via a “breakpoint”; the positive linear relation ($b = 1.10$, $z = 253.76$, $p < 1e - 4$) and the negative linear relation ($b = -2.57$, $z = -659.26$, $p < 1e - 4$) are both significant. Please refer to Appendix C.1 for assumptions in applying the two-lines test approach.

Discussion The above results reveal that (i) the deep features help the models to describe concepts with very few attributes, (ii) representation trained from very simple or very complex datasets usually have a shorter concept description length than those trained on other datasets, and (iii) the subjective complexity significantly comes in an inverted-U relation with the visual complexity.

3.2. Computation vs. complexity

This experiment evaluates the rule- and similarity-based generalization using the representations in Sec. 3.1. We hypothesized that representations with relatively high subjective complexity outperform those with low subjective complexity in rule-based generalization, while the trend is the opposite in similarity-based generalization.

Method The evaluation of generalization is designed with two phases: in-domain and out-of-domain generalizations. The former consists of unseen samples from the test set of ACRE and ImageNet, whereas the latter contains unseen samples of unknown concepts collected from the internet. Each phase has a dataset with pairs for similarity-based generalization evaluation and a dataset with quadruples for rule-based generalization evaluation.

The evaluation protocol for similarity-based generalization extends its definition in Sec. 2.1. Formally, given unknown concept c' and known concepts $c \in \mathcal{C}$, the ranking of the pairwise metric measurement is $S = \{\sigma_0(c_i, c') \geq \sigma_0(c_j, c') | c_i, c_j \in \mathcal{C}\}$. The representation ranking S_r is ob-

tained by the cosine similarity between two representation vectors $\cos(z_i, z')$. The ground-truth ranking S_h is obtained by human judgment. Hence, the generalization capability of the representation can be quantified by the rank correlation coefficient (Spearman, 1961) as an accuracy measurement.

Similarly, the evaluation protocol for rule-based generalization is defined as follows. Given incomplete rule $r'(c_3, c')$ and known rules $r_i(c_1, c_2) \in \mathcal{R}$, the ranking score R_r of representation is reduced to a cosine similarity calculation $\cos(z_2 - z_1 + z_3, c')$, and the ground-truth ranking R_h is obtained by human judgment (Mikolov et al., 2013). We obtain the ground-truth concepts by literal meanings through the language representation model GloVe (Pennington et al., 2014). The image examples are retrieved from datasets (in-domain) or the internet (out-of-domain) with label embedding matching (Vendrov et al., 2015). The details of human study, approved by the Institutional Review Board (IRB) of Peking University, can be found in Appendix D.3.

Results The quantitative results for in-domain generalization evaluation are illustrated in Fig. 5. In similarity-based generalization, the representation trained from ImageNet outperforms others (over 15%), and LEGO outperforms its more complex counterparts 2D-Geo and ACRE (over 10%). In rule-based generalization, the representation trained from ACRE outperforms its more complex counterpart, ImageNet

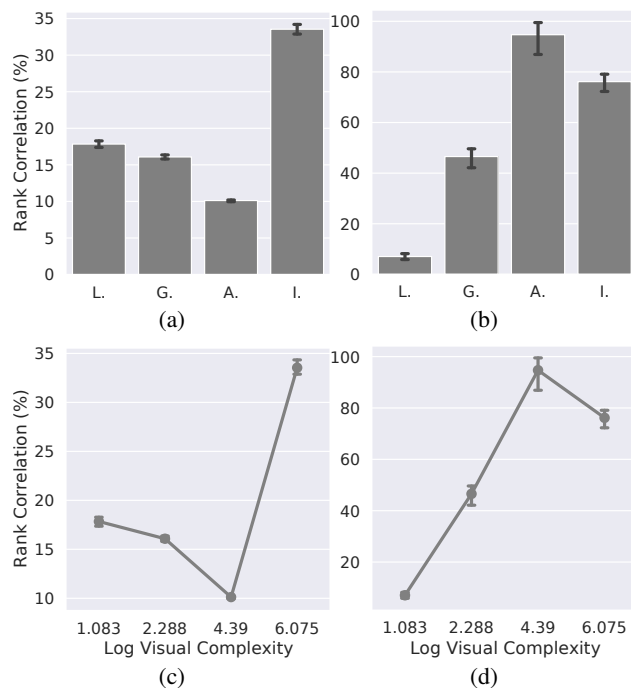


Figure 5. Quantitative results of Computation vs. Complexity. (a)(b) The rank correlation of similarity- and rule-based generalization with the four representations trained from four datasets. (c)(d) The rank correlation of similarity- and rule-based generalization according to the visual complexity. These plots reflect the landscape in Fig. 2. (L: LEGO, G:2D-Geo, A: ACRE, I: ImageNet)

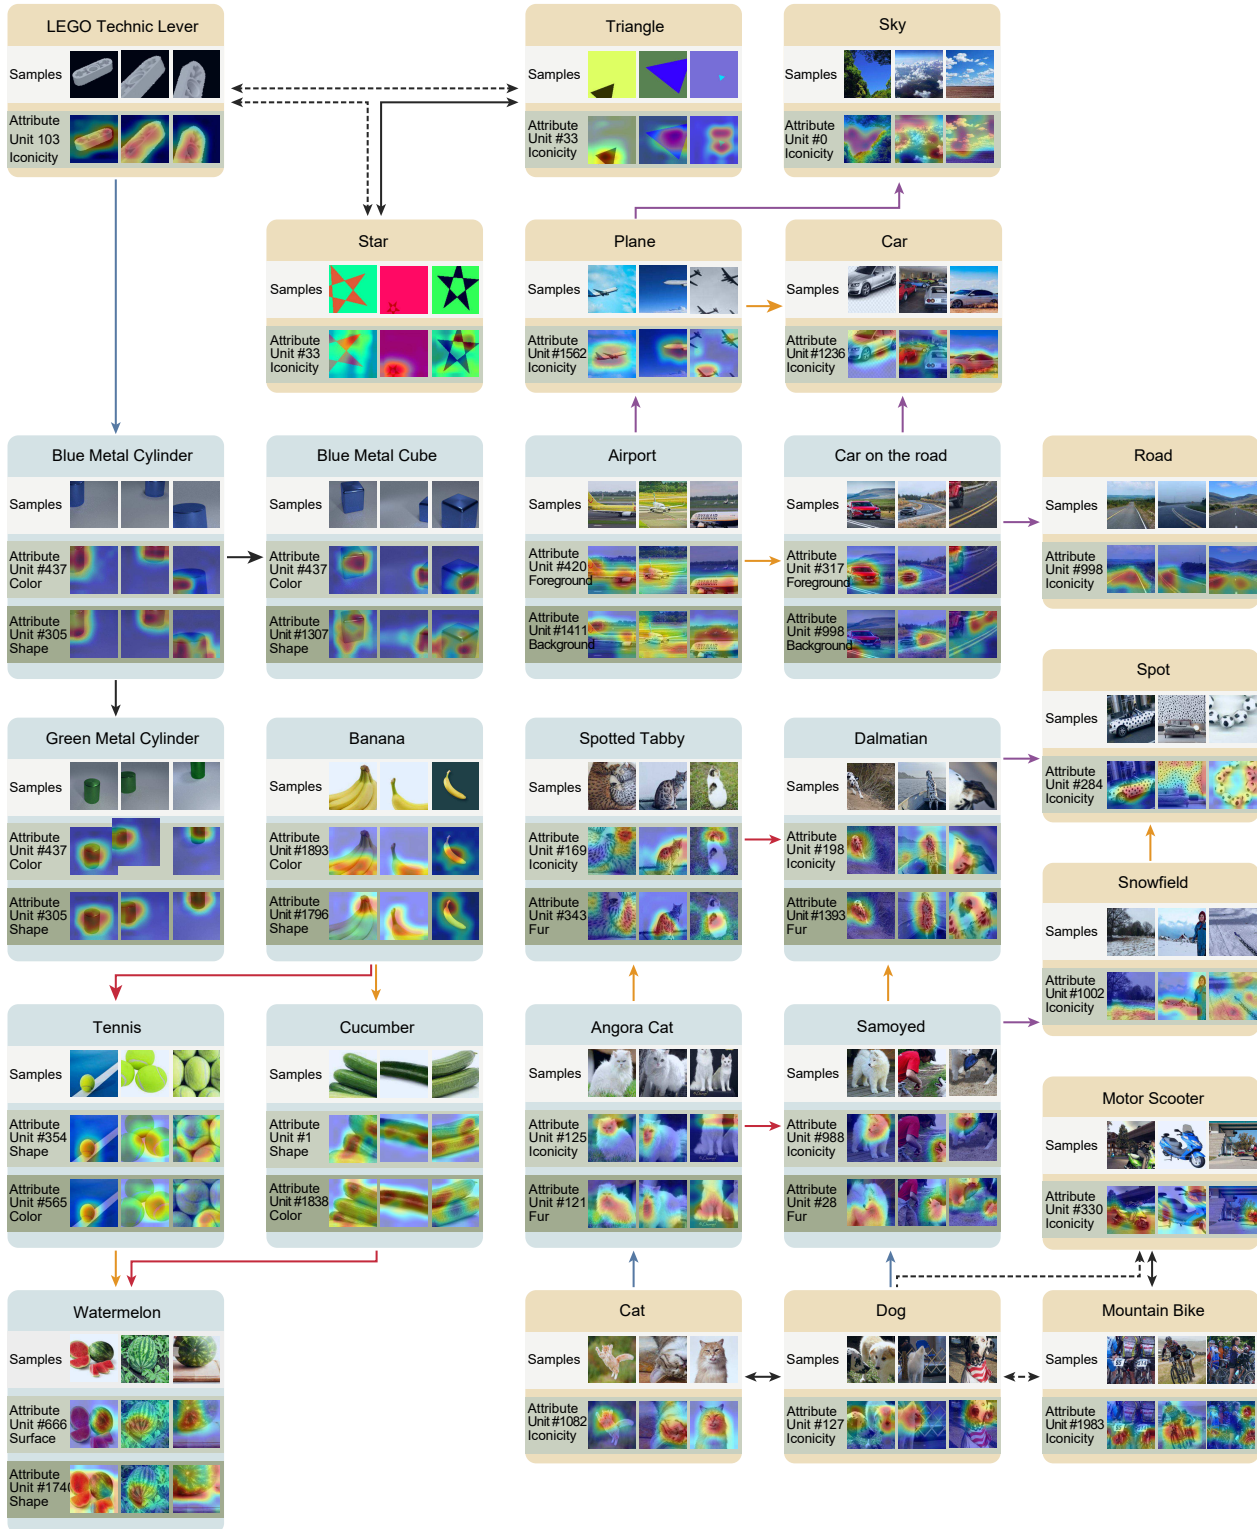


Figure 6. A landscape of similarity- and rule-based generalization over concepts with relatively high and low subjective complexity, considering both concept complexities and concept hierarchy. Zoom in for more details. Bidirectional arrows denote the similarity judgment between concepts, wherein concepts linked by solid lines are more similar than those linked by dashed lines. Arrows denote rules over concepts. Rule-based generalization in basic-level generalizes given rules to unknown rules. Similarity shifts to rules when the sample hierarchy goes from superordinate-level to subordinate-level (e.g., from block to blue cylinder, from cat to angora cat). Rules shift to similarity as the sample hierarchy goes from subordinate-level to superordinate-level (e.g., from car on the road to car, from dalmatian to spot). We also note a confusing similarity judgment between blue cylinder, blue cube, and green cylinder.

(over 20%). Though the models trained on ImageNet and ACRE reach the highest accuracy on similarity- and rule-based generalization, this is not likely due to over-fitting in training: The objective of visual categorization is different from that of generalization, thus the over-fitting on one visual categorization would not result in an over-fitting on other objectives. Intuitively, representations trained on more complex datasets span more complex attribute spaces. However, the result implies that the shift between similarity- and rule-based generalization is non-monotonic as the dataset complexity increases; it is more correlated to the subjective complexity based on Sec. 3.1. Hence, there is a significant negative relationship between similarity-based generalization and subjective complexity ($r = -.48, p < .05$), and a significant positive relationship between rule-based generalization and subjective complexity ($r = .68, p < .01$).

Fig. 6 illustrates the qualitative results for out-of-domain generalization. As shown in Fig. 7, though never tuned on the unseen examples, the representation model also captures representative attributes for unknown concepts, which supports our argument in Sec. 2.1 that RoA has the potential to serve as a prior for Bayesian generalization. Further, we visualize the most representative attributes of each concept by upsampling the activated feature vector to the size of the original image (Bau et al., 2020); the attributes are located around the peaks. Most attributes with high RoA are explainable, such as the shape attribute shared by *blue cylinder* and *green cylinder*, shape and color captured by two distinct attributes in *banana* and *watermelon*, and foreground object (plane, car) and background (road, field) attributes in *airport* and *car on the road*. Those concepts with more than one meaningful attributes are sensitive to rule-based generaliza-




Figure 7. **The RoA matrix.** Most (21 out of 25) concepts are unknown; high saturation indicates high RoA value. The diagonal elements are the most representative attributes of all concepts. Please refer to Appendix E for the full diagram.

tion. By contrast, those concepts with only one meaningful attribute, such as *dog-like face* for *dog*, *car-like shape* for *car*, are sensitive to similarity-based generalization.

Discussion The above experiment reveals that (i) both similarity- and rule-based generalizations are not significantly related to the visual complexity of datasets, (ii) the capability of similarity-based generalization has a significant negative relationship with the subjective complexity of representation, and (iii) the capability of rule-based generalization has a positive relationship with the subjective complexity of representation. We empirically articulate that the computation-mode-shift significantly exists, and similarity shifts to rules as the subjective complexity increases; please refer to Appendix E for more details.

3.3. A statistical interpretation

Subjective complexity in natural image statistics According to algorithmic information theory (Chaitin, 1977), a concept’s subjective complexity is proportional to the probability of perceiving this concept. This is consistent with the subjective complexity of visual concepts defined in our work. An attribute z is representative for concept c when $\text{RoA}(z, c)$ is relatively high; we have a high probability of observing the attribute by the concept (e.g., $P(z|c) = 1$) or only by the concept (i.e., $\sum_{\hat{c} \neq c} P(\hat{c}|z)$ is small). Specifically, complex concepts (e.g., *dog*, *cat*), though consisting of many attributes (e.g., *fur*, *ear*), tend to have a unique attribute of *view as a whole* to distinguish these concepts from others because we can hardly observe them in other concepts. Conversely, simple concepts (e.g., *circle*, *cylinder*) can be observed by many other concepts (e.g., *wheel*, *chimney*) and also have other attributes (e.g., *number of angles*, *smoothness*). Nevertheless, the attribute *shape* is one of the simple attributes to describe these concepts; representation of these concepts emerges iconicity (Guo et al., 2003; Fay et al., 2010; 2013; Qiu et al., 2022).

Meanwhile, for those concepts that are either too simple or too complex (e.g., *watermelon*, *airport*), no unique or simple attribute can distinguish them from others; i.e., $\text{RoA}(z, c)$ is not high. In these cases, we have to describe them with more attributes. Of note, this interpretation is also in line with the principle of rational reference (Frank & Goodman, 2012; Goodman & Frank, 2016).

From similarity to rules Since similarity gradient can be viewed as a partial order defined on a single set (Tenenbaum, 1999), sorting hypotheses requires numerical comparison in the same domain. Hence, similarity judgment in a single attribute space z_i is simply calculating the similarity between concepts c_j and c_k by $d(z_i^{(j)}, z_i^{(k)})$, where $d(\cdot, \cdot)$ can be an arbitrary similarity or distance metric (Ontañón, 2020). As the number of independent attribute spaces increases (i.e., subjective complexity increases), the similarity

becomes subtle as we have to consider multiple independent attributes. Of note, the attribute spaces are those obtained after dimension reduction (Xie et al., 2016); those concept representations are almost distributed uniformly (Blum et al., 2020) unless we assign weights to different attribute spaces by only considering very few attributes. For example, *watermelon* is similar to *tennis* in the attribute space of *shape*, but it becomes similar to *cucumber* in the attribute space of *color*; *airport* is similar to *plane* in the attribute space of *foreground object* and is similar to *land and sky* in the attribute space of *background context*. In this work, we reduce similarity judgment over multiple attribute spaces to rules defining relations over two concepts: At least one shared attribute space bridges the two concepts.

From rules to similarities As the number of independent attribute spaces (*i.e.*, subjective complexity) decreases, rules are moved back to similarity. For example, we have the rule relating *dalmatian* to *spotted tabby* by *fur texture*, and can generalize it to *samoyed* to *angora cat*. However, when the concepts are more complex (*e.g.*, *dalmatian* and *samoyed* fall in *dog*, or *spotted tabby* and *angora cat* belong to *cat*), rules are difficult on these concepts; instead, we directly apply similarity judgment.

Concept complexities and hierarchy When visual complexity moves from low to high, we have visual concepts move from *simple and universal* to *complex and unique*. We argue that these two ends consist of *superordinate* concepts (Xu & Tenenbaum, 2007), usually on higher hierarchies. Objects such as *watermelon*, attribute-specified animals such as *samoyed* are subordinate concepts of *ball* and *dog*, respectively; scenes such as *airport* are compositions of subordinate concepts like *plane* and *land and sky*. In a top-down view, we have concepts with increasing subjective complexity and more shared attribute spaces to generalize by rules. In a bottom-up view, the attribute spaces are reduced to the *simple* or *unique* ones, easy for similarity judgment.

4. Conclusion

We have analyzed the complexity of concept generalization in the natural visual world, in Marr’s *representational* and *computational level*. At the representational level, the subjective complexities significantly fall in an inverted-U relation with the increment of visual complexity. At the computational level, the rule-based generalization is significantly positively correlated with the subjective complexity of the representation, while the trend is the opposite in similarity-based generalization. RoA bridges the two levels by unifying the frequentist properties of natural images (sensory-based) and the Bayesian properties of concepts (knowledge-derived) (Bi, 2021). It is easy to obtain, is flexible to an extent, and captures contextual rationality, thus may serve as humans’ *visual common sense* (Zhu et al.,

2020; Fan et al., 2022). Readers may refer to Appendix A for additional remarks about the rationale for some decisions made during the process of this work.

The limitations of this work lead to several future directions:

- We only demonstrated the inverted-U relation and the correlation empirically. Can we provide them theoretically, from the aspect of information theory and statistics?
- Can we further extend the generalization evaluation to a larger scale that helps to quantitatively probe the continuum space between similarity and rules?
- Are our findings consistent with those in other environments where the concepts are represented in other modalities (*e.g.*, auditory, tactile, gustatory, and olfactory)?
- If using only a few attributes with high RoA improves the accuracy of the visual categorization task, as Sec. 3.1 suggests, can we build an algorithm that samples from RoA adaptively for stronger generalization?
- If RoA reflects humans’ visual commonsense, can we model the communications between individuals toward commonsense knowledge as a pursuit of the common grounds on representative attributes for the concepts to be communicated (Tomasello, 2010)?

With many questions unanswered, we hope to shed light on future research on Bayesian generalization and offer a different view of generalization (Lake & Baroni, 2018; Ruis et al., 2020; Xie et al., 2021; Li et al., 2022; 2023).

Acknowledgements

The authors thank Miss Chen Zhen (BIGAI) for designing figures and four anonymous reviews for constructive feedback. M.X., S.-C.Z., and Y.Z. are supported in part by the National Key R&D Program of China (2022ZD0114900), M.X. and Y.Z. are supported in part by the Beijing Nova Program, W.H. is supported in part by the startup fund of Beijing Jiaotong University (2023XKRC006), and Y.W. is supported in part by NSF DMS-2015577.

Broader impact

The broader significance of this study lies in its potential to enhance our understanding of concept generalization. By exploring the relationship between concept complexity and the shift between rule- and similarity-based generalization, this research contributes to the development of more advanced cognitive models and artificial intelligence systems that can better mimic human-like learning and reasoning. This, in turn, could lead to the creation of more sophisticated and efficient tools to assist humans in various tasks, such as problem-solving and communication. Ultimately, this research contributes to a deeper understanding of the foundations of human cognition and intelligence, providing insights that could shape the future trajectory of human development and interaction with artificial intelligence.

References

- Abbott, J., Austerweil, J., and Griffiths, T. Constructing a hypothesis space from the web for large-scale bayesian word learning. In *Annual Meeting of the Cognitive Science Society*, 2012. 3
- Abbott, J. T., Heller, K. A., Ghahramani, Z., and Griffiths, T. Testing a bayesian measure of representativeness using a large image database. In *Advances in Neural Information Processing Systems*, 2011. 3
- Allen, C. and Hospedales, T. Analogies explained: Towards understanding word embeddings. In *International Conference on Machine Learning*, 2019. 3
- Bau, D., Zhu, J.-Y., Strobel, H., Lapedriza, A., Zhou, B., and Torralba, A. Understanding the role of individual units in a deep neural network. *Proceedings of the National Academy of Sciences*, 117(48):30071–30078, 2020. 4, 5, 8, A2
- Bi, Y. Dual coding of knowledge in the human brain. *Trends in Cognitive Sciences*, 25(10):883–895, 2021. 3, 9, A1
- Blum, A., Hopcroft, J., and Kannan, R. *Foundations of data science*. Cambridge University Press, 2020. 9
- Bruni, E., Tran, N.-K., and Baroni, M. Multimodal distributional semantics. *Journal of Artificial Intelligence Research*, 49:1–47, 2014. A3
- Chaitin, G. J. Algorithmic information theory. *IBM Journal of Research and Development*, 21(4):350–359, 1977. 3, 8
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In *Conference on Computer Vision and Pattern Recognition*, 2009. 4, 5, A3
- Donderi, D. C. Visual complexity: a review. *Psychological Bulletin*, 132(1):73, 2006. 1, 2
- Edmonds, M., Gao, F., Liu, H., Xie, X., Qi, S., Rothrock, B., Zhu, Y., Wu, Y. N., Lu, H., and Zhu, S.-C. A tale of two explanations: Enhancing human trust by explaining robot behavior. *Science Robotics*, 4(37), 2019. 3
- El Korchi, A. and Ghanou, Y. 2d geometric shapes dataset—for machine learning and pattern recognition. *Data in Brief*, 32: 106090, 2020. 4, 5, A3
- Ellis, K., Wong, C., Nye, M., Sable-Meyer, M., Cary, L., Morales, L., Hewitt, L., Solar-Lezama, A., and Tenenbaum, J. B. Dreamcoder: Growing generalizable, interpretable knowledge with wake-sleep bayesian program learning. *arXiv preprint arXiv:2006.08381*, 2020. 1
- Fan, L., Xu, M., Cao, Z., Zhu, Y., and Zhu, S.-C. Artificial social intelligence: A comparative and holistic view. *CAAI Artificial Intelligence Research*, 1(2):144–160, 2022. 9
- Fay, N., Garrod, S., Roberts, L., and Swoboda, N. The interactive evolution of human communication systems. *Cognitive Science*, 34(3):351–386, 2010. 3, 8, A1
- Fay, N., Arbib, M., and Garrod, S. How to bootstrap a human communication system. *Cognitive Science*, 37(7):1356–1367, 2013. 3, 8, A1
- Fay, N., Ellison, M., and Garrod, S. Iconicity: From sign to system in human communication and language. *Pragmatics & Cognition*, 22(2):244–263, 2014. 3, A1
- Frank, M. C. and Goodman, N. D. Predicting pragmatic reasoning in language games. *Science*, 336(6084):998–998, 2012. 8
- Gentner, D. and Markman, A. B. Structure mapping in analogy and similarity. *American Psychologist*, 52(1):45, 1997. 3
- Gershman, S. J., Horvitz, E. J., and Tenenbaum, J. B. Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. *Science*, 349(6245):273–278, 2015. A1
- Gittens, A., Achlioptas, D., and Mahoney, M. W. Skip-gram-zipf+ uniform= vector additivity. In *Annual Meeting of the Association for Computational Linguistics*, 2017. 3
- Goodman, N. D. and Frank, M. C. Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences*, 20(11):818–829, 2016. 8
- Griffiths, T. and Tenenbaum, J. From algorithmic to subjective randomness. In *Advances in Neural Information Processing Systems*, 2003. 1, 2
- Griffiths, T. L., Abbott, J. T., and Hsu, A. S. Exploring human cognition using large image databases. *Topics in Cognitive Science*, 8(3):569–588, 2016. 3
- Guo, C.-e., Zhu, S.-C., and Wu, Y. N. Towards a mathematical theory of primal sketch and sketchability. In *International Conference on Computer Vision*, 2003. 8
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Conference on Computer Vision and Pattern Recognition*, 2016. A1
- Holyoak, K. J., Ichien, N., and Lu, H. From semantic vectors to analogical mapping. *Current Directions in Psychological Science*, 31(4):355–361, 2022. 3
- Jiang, G., Xu, M., Xin, S., Liang, W., Peng, Y., Zhang, C., and Zhu, Y. Mewl: Few-shot multimodal word learning with referential uncertainty. In *International Conference on Machine Learning*, 2023. 2, 3
- Johnson, J., Hariharan, B., Van Der Maaten, L., Fei-Fei, L., Lawrence Zitnick, C., and Girshick, R. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Conference on Computer Vision and Pattern Recognition*, 2017. 4, A3
- Julesz, B. Visual pattern discrimination. *IRE Transactions on Information Theory*, 8(2):84–92, 1962. 3
- Kersten, D., Mamassian, P., and Yuille, A. Object perception as bayesian inference. *Annual Review of Psychology*, 55:271–304, 2004. 1, A1
- Krasin, I., Duerig, T., Alldrin, N., Ferrari, V., Abu-El-Haija, S., Kuznetsova, A., Rom, H., Uijlings, J., Popov, S., Kamali, S., Mallocci, M., Pont-Tuset, J., Veit, A., Belongie, S., Gomes, V., Gupta, A., Sun, C., Chechik, G., Cai, D., Feng, Z., Narayanan, D., and Murphy, K. Openimages: A public dataset for large-scale multi-label and multi-class image classification. *Dataset available from <https://storage.googleapis.com/openimages/web/index.html>*, 2017. A3

- Lake, B. and Baroni, M. Generalization without systematicity: On the compositional skills of sequence-to-sequence recurrent networks. In *International Conference on Machine Learning*, 2018. 9
- Lake, B. M., Salakhutdinov, R., and Tenenbaum, J. B. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015. 1, A1
- Lake, B. M., Ullman, T. D., Tenenbaum, J. B., and Gershman, S. J. Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40, 2017. A1
- Li, M., Vitányi, P., et al. *An introduction to Kolmogorov complexity and its applications*. Springer, 2008. 3, 4
- Li, Q., Zhu, Y., Liang, Y., Wu, Y. N., Zhu, S.-C., and Huang, S. Neural-symbolic recursive machine for systematic generalization. *arXiv preprint arXiv:2210.01603*, 2022. 9
- Li, Q., Huang, S., Hong, Y., Zhu, Y., Wu, Y. N., and Zhu, S.-C. A minimalist dataset for systematic generalization of perception, syntax, and semantics. In *International Conference on Learning Representations*, 2023. 9
- Lieder, F. and Griffiths, T. L. Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources. *Behavioral and Brain Sciences*, 43, 2020. A1
- Lind, J. T. and Mehlum, H. With or without u? the appropriate test for a u-shaped relationship. *Oxford bulletin of economics and statistics*, 72(1):109–118, 2010. A2
- Marr, D. *Vision*. W. H. Freeman and Company, 1982. 2, A1
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., and Dean, J. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems*, 2013. 3, 6
- Miller, G. A. *WordNet: An electronic lexical database*. MIT Press, 1998. 3
- Ontañón, S. An overview of distance and similarity functions for structured data. *Artificial Intelligence Review*, 53(7):5309–5351, 2020. 8
- Pennington, J., Socher, R., and Manning, C. D. Glove: Global vectors for word representation. In *Annual Conference on Empirical Methods in Natural Language Processing*, pp. 1532–1543, 2014. 6
- Qiu, S., Xie, S., Fan, L., Gao, T., Zhu, S.-C., and Zhu, Y. Emergent graphical conventions in a visual communication game. In *Advances in Neural Information Processing Systems*, 2022. 3, 8
- Ruis, L., Andreas, J., Baroni, M., Bouchacourt, D., and Lake, B. M. A benchmark for systematic generalization in grounded language understanding. In *Advances in Neural Information Processing Systems*, 2020. 9
- Salakhutdinov, R., Tenenbaum, J., and Torralba, A. One-shot learning with a hierarchical nonparametric bayesian model. In *ICML Workshop on Unsupervised and Transfer Learning*, 2012. 2
- Shannon, C. E. A mathematical theory of communication. *The Bell System Technical Journal*, 27(3):379–423, 1948. 4
- Shepard, R. N. Toward a universal law of generalization for psychological science. *Science*, 237(4820):1317–1323, 1987. 1
- Simonsohn, U. Two lines: A valid alternative to the invalid testing of u-shaped relationships with quadratic regressions. *Advances in Methods and Practices in Psychological Science*, 1(4):538–555, 2018. 6
- Sloman, S. A. and Rips, L. J. *Similarity and symbols in human thinking*. MIT Press, 1998. 1
- Spearman, C. *The proof and measurement of association between two things*. Appleton-Century-Crofts, 1961. 6
- Sun, Z. and Firestone, C. Seeing and speaking: How verbal “description length” encodes visual complexity. *Journal of Experimental Psychology: General*, 2021. 1, 2
- Tatman, R. The lego parts, sets, colors, and inventories of every official lego set, 2017. 4, 5, A3
- Tenenbaum, J. B. Bayesian modeling of human concept learning. In *Advances in Neural Information Processing Systems*, 1998. 1
- Tenenbaum, J. B. Rules and similarity in concept learning. In *Advances in Neural Information Processing Systems*, 1999. 1, 4, 8, A1
- Tenenbaum, J. B. and Griffiths, T. L. Generalization, similarity, and bayesian inference. *Behavioral and Brain Sciences*, 24(4): 629–640, 2001a. 1, A1
- Tenenbaum, J. B. and Griffiths, T. L. The rational basis of representativeness. In *Annual Meeting of the Cognitive Science Society*, 2001b. 3, 4
- Tenenbaum, J. B., Kemp, C., Griffiths, T. L., and Goodman, N. D. How to grow a mind: Statistics, structure, and abstraction. *Science*, 331(6022):1279–1285, 2011. 1
- Tishby, N. and Zaslavsky, N. Deep learning and the information bottleneck principle. In *2015 IEEE Information Theory Workshop (ITW)*, 2015. 4
- Tomasello, M. *Origins of human communication*. MIT press, 2010. 9
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. A3
- Vendrov, I., Kiros, R., Fidler, S., and Urtasun, R. Order-embeddings of images and language. *arXiv preprint arXiv:1511.06361*, 2015. 6
- Wolfram, S. et al. *A new kind of science*. Wolfram Media Champaign, 2002. 2
- Wu, Y. N., Guo, C.-E., and Zhu, S.-C. From information scaling of natural images to regimes of statistical models. *Quarterly of Applied Mathematics*, pp. 81–122, 2008. 3, 4, A1
- Wu, Y. N., Si, Z., Gong, H., and Zhu, S.-C. Learning active basis model for object detection and recognition. *International Journal of Computer Vision*, 90(2):198–235, 2010. 3

- Wu, Y. N., Gao, R., Han, T., and Zhu, S.-C. A tale of three probabilistic families: Discriminative, descriptive, and generative models. *Quarterly of Applied Mathematics*, 77(2):423–465, 2019. 3
- Xian, Y., Lampert, C. H., Schiele, B., and Akata, Z. Zero-shot learning—a comprehensive evaluation of the good, the bad and the ugly. *Transactions on Pattern Analysis and Machine Intelligence*, 41(9):2251–2265, 2018. 4, 5, A3
- Xie, J., Lu, Y., Zhu, S.-C., and Wu, Y. N. A theory of generative convnet. In *International Conference on Machine Learning*, 2016. 3, 9
- Xie, J., Gao, R., Nijkamp, E., Zhu, S.-C., and Wu, Y. N. Representation learning: A statistical perspective. *Annual Review of Statistics and Its Application*, 7:303–335, 2020. 3, 4
- Xie, S., Ma, X., Yu, P., Zhu, Y., Wu, Y. N., and Zhu, S.-C. Halma: Humanlike abstraction learning meets affordance in rapid problem solving. *arXiv preprint arXiv:2102.11344*, 2021. 9
- Xu, F. and Tenenbaum, J. B. Word learning as bayesian inference. *Psychological Review*, 114(2):245, 2007. 1, 2, 9, A6
- Zhang, C., Gao, F., Jia, B., Zhu, Y., and Zhu, S.-C. Raven: A dataset for relational and analogical visual reasoning. In *Conference on Computer Vision and Pattern Recognition*, 2019a. 2
- Zhang, C., Jia, B., Gao, F., Zhu, Y., Lu, H., and Zhu, S.-C. Learning perceptual inference by contrasting. In *Advances in Neural Information Processing Systems*, 2019b. 3
- Zhang, C., Jia, B., Edmonds, M., Zhu, S.-C., and Zhu, Y. Acre: Abstract causal reasoning beyond covariation. In *Conference on Computer Vision and Pattern Recognition*, 2021a. 4, 5, A3
- Zhang, C., Jia, B., Zhu, S.-C., and Zhu, Y. Abstract spatial-temporal reasoning via probabilistic abduction and execution. In *Conference on Computer Vision and Pattern Recognition*, 2021b. 3
- Zhang, C., Xie, S., Jia, B., Wu, Y. N., Zhu, S.-C., and Zhu, Y. Learning algebraic representation for systematic generalization in abstract reasoning. In *European Conference on Computer Vision*, 2022. 3
- Zhou, B., Lapedriza, A., Khosla, A., Oliva, A., and Torralba, A. Places: A 10 million image database for scene recognition. *Transactions on Pattern Analysis and Machine Intelligence*, 40(6):1452–1464, 2017. 4, 5, A3
- Zhu, S.-C., Wu, Y. N., and Mumford, D. Minimax entropy principle and its application to texture modeling. *Neural Computation*, 9(8):1627–1660, 1997. 3
- Zhu, S.-C., Wu, Y. N., and Mumford, D. Filters, random fields and maximum entropy (frame): Towards a unified theory for texture modeling. *International Journal of Computer Vision*, 27(2):107–126, 1998. 3
- Zhu, S.-C., Guo, C.-E., Wang, Y., and Xu, Z. What are textons? *International Journal of Computer Vision*, 62(1):121–143, 2005. 3
- Zhu, Y., Gao, T., Fan, L., Huang, S., Edmonds, M., Liu, H., Gao, F., Zhang, C., Qi, S., Wu, Y. N., et al. Dark, beyond deep: A paradigm shift to cognitive ai with humanlike common sense. *Engineering*, 6(3):310–345, 2020. 9

A. Additional remarks

A.1. The uniqueness of the natural visual world

Why do we only use the modality of vision to investigate the complexity of Bayesian generalization? Vision is unique for the diverse complexities both in the natural visual world and in the semantic space (Kersten et al., 2004), which relates vision to the discussion of levels of abstraction (Wu et al., 2008). To some extent, vision serves as the bridge between abstract language-derived knowledge and perceptual sensory-derived knowledge (Bi, 2021). The two ends of the continuum of Bayesian generalization touch the functional essences of rule-based symbolic signals and similarity-based perceptual signals (Tenenbaum, 1999). In particular, the very final development of symbolic signals leads to the emergence of language, based on the compositionality of symbols and rules as the basic feature of language. The psychology literature supports the hypothesis that language is emerged from visual communications by abstracting visual concepts toward hieroglyphs through their iconicity (Fay et al., 2010; 2013; 2014). Both simple and universal visual concepts, such as geometric shapes, and complex and unique visual concepts, such as animals and artificial objects, are all related to corresponding abstract concepts by iconicity. By contrast, those concepts that are neither simple nor unique are unlikely to be abstracted by iconicity since they are described by multiple representative attributes—though each attribute can be generalized through iconicity respectively, putting different attribute spaces together is not making sense—by contrast, those concepts naturally satisfy the compositionality of language, thus are appropriate for rule-based generalization. In this sense, vision is not only a modality of data but is the hallmark of human intelligence, evolving perceptual senses toward language for communications. Hence, vision is meaningful and sufficient for investigating the complexity of Bayesian generalization.

Consider other modalities, say audio, the second common resource of sensory input. Although we could define audio complexity and try to correlate it with subjective complexity, audio is only a perceptual sensory—abstraction of raw audio is not related to any semantic meaning, thus does not provide much insight into human intelligence; also the diversity of audio complexity is far less than its visual counterpart. Hence, generalizing the experiments to audio data may be a bonus but never provides us insights as deep as that provided by visual data.

A.2. The appropriateness of the computational modeling

Thanks to Marr’s paradigm (Marr, 1982), we could separate the computational-level problem and the representational-level problem, where we study computation problems regardless of their algorithmic representation or physical implementation in either humans or machines (Lake et al., 2015). Hence, under the same computation problem, whether the algorithm is neural networks or brain circuits is not the problem in the scope.

Since the two parts of our computation problem—Bayesian generalization (Tenenbaum & Griffiths, 2001a) and subjective complexity (Lake et al., 2017)—have established solid backgrounds in human cognition, we have a sufficient prerequisite for studying the complexities in the natural visual world. Though there may be infinite interpretations of human cognitive models (Lieder & Griffiths, 2020), constrained by previous theories and the principle of resource-rational analysis (Lieder & Griffiths, 2020; Gershman et al., 2015), we can make assumptions about the Bayesian derivations.

B. Implementation details

B.1. Implementing basic discriminative models

The basic discriminative models are employed from ResNet (He et al., 2016), thus the feature space is spanned by a 512-d or 2048-d feature vector (dimensions are different by the different depths of ResNet architecture). All models are trained on eight NVIDIA A100 80GB GPUs. All images are resized to a fixed size of 224 * 224, and data augmentation techniques such as random cropping and horizontal flipping are applied to increase the variability of the training data. Mean subtraction and standard deviation normalization are applied. During training, we used a cross-entropy loss function to optimize the model for classification performance. We also employed early stopping based on the validation accuracy. For both ImageNet and Places365 datasets, we utilized the official pre-trained models from PyTorch and proceeded to finetune them over the course of 20 epochs. For other datasets, we train our model from scratch.

B.2. Implementing RoA

In general, the RoA computes a score for each attribute z_i over each concept c . The output of RoA is a matrix where the column space is the context of all the concepts in the natural visual world, and the row space is all the attributes.

Assume we have three samples $\{x^{(1)}, x^{(2)}, x^{(3)}\} \in X$ of concept c , then the output of f provides the attribute vectors $\mathbf{z}^{(1)}, \mathbf{z}^{(2)}, \mathbf{z}^{(3)} \in \mathbb{R}^{H \times W \times D}$ respectively. We then adaptively pool each feature map $\mathbf{z}_i^{(1)}, \mathbf{z}_i^{(2)}, \mathbf{z}_i^{(3)} \in \mathbb{R}^{H \times W \times d}$ in each dimension of the attribute vector to a scalar $z_i^{(1)}, z_i^{(2)}, z_i^{(3)}$, thus $\mathbf{z}^{(1)}, \mathbf{z}^{(2)}, \mathbf{z}^{(3)} \in \mathbb{R}^d$. $P(z_i|c)$ is calculated by normalizing over the dimensions of the centroid vector of all \mathbf{z} given the set of samples of concept c , e.g., $\widetilde{\mathbf{z}}^{(k)}, k = 1, 2, 3$.

B.3. Implementing subjective complexity measurement

Since the calculation of the absolute value of $L(\hat{c})$ may encounter multiple solutions, we employ accuracy gain (Bau et al., 2020) to compute a relative $L(\hat{c})$ specifically for \hat{c} . The accuracy gain approach considers the categorization accuracy difference for a single concept before and after removing the effect of a specific neuron, defined as:

$$\begin{aligned} \Delta Acc_K(\hat{c}) &= P(\hat{c} = c | c = \arg_c \max P(c|z_1, \dots, z_K; \phi)) \\ &\quad - P(\hat{c} = c | c = \arg_c \max P(c|z_1, \dots, z_{K-1}; \phi)), \end{aligned} \tag{A1}$$

where $K \geq 2$ and $\Delta Acc_1(\hat{c}) = P(\hat{c} = c | c = \arg_c \max P(c|z_1; \phi))$. Hence, the relative $L(\hat{c})$ is exactly computed by:

$$L_{relative}(\hat{c}) = \min_K \max_K Acc_K(c), \tag{A2}$$

which serves as the heuristic to search for the minimum K to calculate absolute $L(\hat{c})$.

C. Method appropriateness checking

C.1. Checking the assumptions of the two-lines test

The ‘‘two-lines test’’ requires a weaker assumption than the mostly used quadratic regression test for testing U-shapes (Lind & Mehlum, 2010), hence the former is employed instead of the latter. Let $y = f(x)$ be the ground-truth function, the U-shape assumes only a sign flip effect in discrete data, where there exists x_c such that $f'(x), x \leq x_c$ and $f'(x), x \geq x_c$ has opposite signs (Lind & Mehlum, 2010). To note, since the data is originally discrete, there is no need to check the existence of $f'(x)$ because it is estimated based on the discrete data points. Hence, the basic hypothesis of the U-shape is that at least one such x_c exists, and the null hypothesis is that no such x_c exists. The null hypothesis is rejected by estimating many x_c values and running two separate linear regressions for $x \leq x_c$ and $x \geq x_c$, respectively. The fact that two regression lines are of opposite signs rejects the null hypothesis. By contrast, the quadratic regression test assumes that the first-order derivative function $f'(x)$ is continuous in the domain. Hence, there is no need to employ the quadratic regression test.

C.2. Checking the assumptions of the linear regression test

The assumptions of the test are (i) linearity of the data, (ii) x values are statistically independent, and (iii) the errors are homoscedastic and normally distributed. We did test the applicability of the linear regression test: (i) the two relations between rank correlation and subjective complexity are intuitively in lines $([(0.1, 7.8), (1.28, 79.1), (2.91, 46.7), (3.08, 99.5)])$ and $([(0.1, 17.1), (1.28, 33.2), (2.91, 15.8), (3.08, 10.2)])$; (ii) all the evaluations are run separately with different random seeds, thus the predictors are statistically independent; (iii) the errors are homoscedastic since the only independent variable is the dataset, which is unlikely to be the source of constant variance of the errors. Consider the null hypothesis of the linear regression test that the coefficient β_1 is zero, which leads to a trivial solution. However, the p-values of both the positive and negative relations are less than 0.05, rejecting the null hypothesis.

C.3. The correctness for combining representation and computation

As illustrated in Fig. 5, we integrated the results in representation *vs.* complexity into this plot to use these plots to demonstrate the computation-mode-shift—the two U-shapes come with opposite trends intuitively show the landscape for concept complexity *vs.* the computation mode, that similarity-based generalization tends to emerge in concepts with very low or very high visual complexity (*i.e.*, the concepts with low subjective complexity, on the left and right ends of the visual complexity axis), and rule-based generalization tends to emerge in concepts with neither very low nor very high visual complexity (*i.e.*, the concepts with high subjective complexity, in the middle of the visual complexity axis). This is the exact claim of the paper. The quantitative results on the significant positive relation between rule-based generalization rank correlation and subjective complexity, and the significant negative relation between similarity-based generalization rank correlation and subjective complexity, both support the claim.

D. Dataset construction

D.1. Empirical analysis datasets

Several widely-used image datasets that represent different concept-wise visual complexity are selected: LEGO (Tatman, 2017), 2D-Geo (El Korchy & Ghanou, 2020), ACRE (Zhang et al., 2021a), AWA (Xian et al., 2018), Places (Zhou et al., 2017), and ImageNet (Deng et al., 2009). In particular, we use the ImageNet subset ImageNet-1k and the AWA2 version of AWA. The ACRE dataset is based on the well-known CLEVR universe (Johnson et al., 2017), which can be rendered with a single object in one panel and without the blicket machine according to (Zhang et al., 2021a). Please refer to Fig. A2 for some examples of the datasets we use. We limit images of each concept in all of these datasets to about 1k to ensure a balanced number of learning samples, which may lead to a gap between our models and the SOTA.

All the codes including the dataset construction, training, and analyzing are available at <https://github.com/YuzheSHI/bayesian-generalization-complexity>.

D.2. Definition of the vocabulary

We leverage a fully-connected probabilistic graph model to obtain the representativeness of every attribute for every concept, where each node is a piece of natural language that serves as either a concept or an attribute describing other concepts. We exploit the RoA in language to generate the in-domain and out-of-domain visual datasets for Bayesian generalization.

Technically, we use the vocabulary from a WordPiece model (e.g., the base version of Bert (Vaswani et al., 2017)), where a word is tokenized into word pieces (also known as subwords) so that each word piece is an element of the dictionary. Non-word-initial units are prefixed with the sign “##” as a continuation symbol. In this way, there is no Out-Of-Vocabulary. This brings the benefit of generalization over all words. Using all these words as attributes or features leads to sufficient coverage. Moreover, some symbols are reserved for unused placeholders, leaving room for features that the language cannot describe. The readers can refer to `vocab.txt` in the supplementary materials for more details about the attribute list.

D.3. Human-in-the-loop dataset validation

We constructed the *similarity-based generalization* and the *rule-based generalization* datasets using both manual approaches and automatic approaches. Details of all datasets are demonstrated in Tab. A1.

For *in-domain similarity-based generalization*, a concept pair with a human-annotated similarity score was first retrieved from MEN dataset (Bruni et al., 2014) and ImageNet dataset (Deng et al., 2009). Next, we used Amazon Mechanical Turk (AMT) to crowd-source the image aligned to the concept. In total, 305 pairs were selected from 500 candidates. One image was aligned with each concept.

For *in-domain rule-based generalization*, we generated the dataset using objects of easy-to-disentangle attributes (e.g., shape and color) (Zhang et al., 2021a; Johnson et al., 2017). Based on these attributes, we constructed the quadruple relation (e.g., blue cube:red cube::blue cylinder:red cylinder). In total, 4800 images and 24 quadruple relations were collected.

For *out-of-domain similarity-based generalization* and *out-of-domain rule-based generalization*, we collected images from an open internet image dataset (Krasin et al., 2017) based on a predefined set of similarity pairs and rule quadruples. Of note, all the selected pairs or rules were uniformly sampled from the dataset instead of manually picked. All the selected images were under human validation.

In the study, AMT workers recruited have acceptance rates higher than or equal to 90% and approved hits more than 500. Each AMT worker was compensated at the rate of 0.01 USD per selection. In total, we have tested 1000 judgments for 500 concept pairs; two judgments per pair. Fig. A1 shows an example of the AMT interface. This human evaluation is approved by the IRB of Peking University.

E. Additional Results

E.1. The convergence of representation vs. generalization

Does the training setting of the representation model affect its generalization ability? Fig. A3 shows the rank correlation on in-domain generalization evaluation w.r.t. the number of training epochs for visual categorization. This result empirically

shows that the generalization ability converges when the representation models are well-trained after 6-10 epochs, and that ability is stable after convergence. The regression line is significantly vertical to the y-axis ($b = .03, a = 69.67, p < 1e - 4$). Hence, we can assume that there are no significant distinctions in generalization ability between representation models being trained to convergence but with different training settings.

E.2. Additional visualization of RoA


Additional visualization results of RoA are illustrated in Fig. A4. Most (21 out of 25) concepts are unknown; high saturation indicates high RoA value.

Fig. A4a shows the concatenation of the 7 confusion matrices where the n -th diagonal indicates the n -top RoA of the concepts.

Fig. A4b shows the concatenation of 120 highest (from the left) and 60 lowest (from the right) attributes with the mean of RoA in the context.

Fig. A4c shows the concatenation of 120 highest (from the left) and 60 lowest (from the right) attributes with the variance of RoA in the context.

Given the concept pair, please evaluate whether the image pair below shows the corresponding concept.



Sun



Sunlight

True

Submit

Figure A1. The AMT interface used to collect human judgments.

LEGO



Plate 1*2



Peg 2M



Technic Lever 3M




Brick 2*2




Brick 1*1


2D-Geo



Triangle




Square



Star



Hexagon



Circle

ACRE




Metal Blue Cube




Metal Purple Cylinder



Rubber Cyan Sphere



Metal Brown Cylinder




Rubber Red Sphere

AWA



Killer Whale



Giant Panda



Fox




Elephant



Giraffe

Places365



Alley



Arch



Bridge




Farm



Ocean

ImageNet



Cock



Stingray



Hen



Jay



Magpie

Figure A2. Examples of datasets used in our work.

Table A1. Details of the datasets for generalization evaluation. Subordinate level indicates the concept being generalized to is a subordinate concept of the known ones, whereas superordinate level indicates the concept being generalized to is a superordinate concept of the known ones. Subordinate level, basic level, and superordinate level are terms introduced in (Xu & Tenenbaum, 2007).

Group	In-domain		Out-of-domain			
Generalization type	Similarity-based	Rule-based	Similarity-based	Rule-based		
Concept hierarchy	basic level	basic level	basic level	basic level	subordinate level	superordinate level
Test-set size	305	24	21	10	10	10




Figure A3. Rank correlation of generalization w.r.t. the number of training epochs for visual categorization.

