COMPUTER SCIENCE

Human-level few-shot concept induction through minimax entropy learning

Chi Zhang[1]*, Baoxiong Jia[1], Yixin Zhu[2]*, Song-Chun Zhu[1,2]

Humans learn concepts both from labeled supervision and by unsupervised observation of patterns, a process machines are being taught to mimic by training on large annotated datasets—a method quite different from the human pathway, wherein few examples with no supervision suffice to induce an unfamiliar relational concept. We introduce a computational model designed to emulate human inductive reasoning on abstract reasoning tasks, such as those in IQ tests, using a minimax entropy approach. This method combines identifying the most effective constraints on data via minimum entropy with determining the best combination of them via maximum entropy. Our model, which applies this unsupervised technique, induces concepts from just one instance, reaching human-level performance on tasks of Raven's Progressive Matrices (RPM), Machine Number Sense (MNS), and Odd-One-Out (O^3). These results demonstrate the potential of minimax entropy learning for enabling machines to learn relational concepts efficiently with minimal input.

INTRODUCTION

Over the past several decades, artificial intelligence and machine learning have seen substantial advancements in perceptual tasks, provided they are supported by extensive, well-curated datasets. Despite these achievements, the domain of relational induction, which is fundamental to human intelligence (*1*–*4*), has not been thoroughly addressed. Humans have the unique ability to not only learn unfamiliar concepts through input-output pairings, akin to object recognition, but also to discern and induce unseen relationships from observing a mere handful of examples. This capability is exemplified by human infants who can recognize patterns, such as the ABB structure in simple phrases like wo fe fe (*5*) and by older children who readily grasp conjunctive and disjunctive relations in causality (*6*). Even individuals from indigenous Amazonian communities, who lack formal education, demonstrate the ability to solve complex geometric concept induction problems similar to those found in IQ tests (*7*). This faculty for relational concept induction, observed not only in humans but also in primates and avians (*8*–*10*), underscores its innate role in intelligence. The proficiency in such induction is often linked to intelligence levels, typically gauged by IQ scores (*1*, *3*, *11*, *12*).

Recent computational efforts have sought to instill machines with inductive reasoning and enhance machine IQ through the examination of Raven's Progressive Matrices (RPM), a widely recognized intelligence test (*13*–*17*). However, prior approaches to these problems by machines markedly diverge from human methodologies: Humans can unsupervisedly infer and induce hidden relational concepts from a sparse array of examples and use this understanding to forecast subsequent outcomes. In contrast, existing machine learning frameworks, particularly those underpinned by deep learning, engage in the learning process supervisedly, correlating given contexts with their outcomes in the hope that the capability for induction will manifest through the training on example pairs. Regrettably, such systems have yet to achieve the level of inductive reasoning routinely demonstrated by humans (*18*–*21*).

A further critical distinction lies in the efficiency of induction between humans and machines. Humans are capable of inducing new concepts from a single contextual instance, while even the most advanced machines necessitate a vast quantity of paired instances for training to "induce" concepts. This discrepancy suggests that machine learning approaches might reduce induction to mere rote memorization, thus undermining the fundamental purpose of intelligence testing.

The core questions for achieving human-like relational concept understanding remain to be answered: How do people induce an unfamiliar concept unsupervisedly by only observing event sequences? Further, how do people generate an unfamiliar concept to explain the observation from so few examples? The problems, when combined, prevent us from reaching human-level few-shot concept induction computationally. For any traditional approaches formulating this problem as classification, feeding a model with extravagant examples not only risks overfitting but also refutes the desirable learning efficiency. Notwithstanding, humans seem to take a unique approach to avoid the aforementioned issues when presented with similar problems, inducing unfamiliar concepts efficiently and unsupervisedly.

To tackle these problems, we present a unified computational framework learned through minimax entropy to achieve relational concept induction and solve challenging abstract reasoning tasks in IQ tests with only a few context examples in a given instance (few-shot). Contrary to the conventional design in formulating concept induction as classification, we take a descriptive route (*22*) in modeling the problem. In our formulation, each instance is independently modeled via a different descriptive model—that is, an energy-based formulation in the Gibbs form as in Filters, RAndom fields and Maximum Entropy (FRAME) (*23*) and DeepFRAME (*24*, *25*). Central to the descriptive route is modeling and learning of the potential function (exponent) in Gibbs' exponential distribution. The potential function is usually realized as the summation of multiple response functions, known as filters. Filter learning strategies have been actively studied in the literature: FRAME (*23*) starts with a finite set of filters and proceeds with a greedy approach in filter selection; Sparse FRAME (*26*) uses a two-stage strategy, which, compared to the original FRAME, first uses matching pursuit to reconstruct observation; Generative Boosting (*27*) recruits a

[1]Beijing Institute for General Artificial Intelligence (BIGAI), Beijing, 10080, China. [2]Institute for Artificial Intelligence, Peking University, Beijing, 10080, China.
*Corresponding author. Email: zhangchi@bigai.ai (C.Z.); yixin.zhu@pku.edu.cn (Y.Z.)

heuristics method in that only filters with sufficiently large gradients will be learned; DeepFRAME (*24*, *25*) takes it to the extreme by modeling a monolithic filter function using deep neural networks and does not explicitly involve filter selection. Unlike prior works, we take a global optimization perspective based on the final objective: We compose our energy model with potential filters initially, dynamically prune and add filters, and then weigh their combination through the minimax entropy principle.

While such a global optimization strategy could theoretically find a better solution than greedy optimization, it also poses a unique challenge: How to handle numerous filters that in the extreme case can cover continuous space? To address this problem, we propose to embed minimax entropy learning under the umbrella of bilevel optimization (*28*–*30*): In the inner-level optimization, we parameterize continuous filter functions and solve for the optimal parameters in each filter family to fit the observation; in the outer-level optimization, we perform minimax entropy learning. In general, the minimum entropy principle, equivalent to minimizing the Kullback-Leibler (KL) divergence between the modeled distribution and the observed distribution, can be thought of as optimizing the best coding scheme with the shortest coding length on average; meanwhile, the maximum entropy principle finds the optimal distribution that matches the observed statistics constrained by the coding system while making the unconstrained dimensions as random as possible.

This idea fits our goal quite effectively: We are only concerned with the underlying "coding scheme"—that is, the hidden concepts to induce which most effectively produce the observation but setting no limits on unrelated patterns. In practice, learning proceeds in three stages in the minimax entropy framework. In the first stage, we solve for the optimal parameters in each continuous filter function and keep them fixed. In the second stage, we perform minimum entropy learning: Here, optimizing the best coding scheme can be reduced to maximizing the log likelihood with respect to filter selection. In the final maximum entropy stage, the best distribution is analytically solved to be the Gibbs form, with responses from selected filters best weighed to maximize the log likelihood. The inner filter learning stage can be implemented independently using any domain-specific solvers, while the outer minimax entropy learning process can be implemented as maximum log likelihood with respect to two sets of parameters: the indicator variables for all filters and the coefficients used to weigh filters' importance. Please refer to Fig. 1 for a graphical illustration of the pipeline.

In practice, we use object-centric representation when representing images, where we extract objects and related attributes (see Materials and Methods for details). For filter design, we consider different linear filter families for RPM and Odd-One-Out (O^3) and compositional

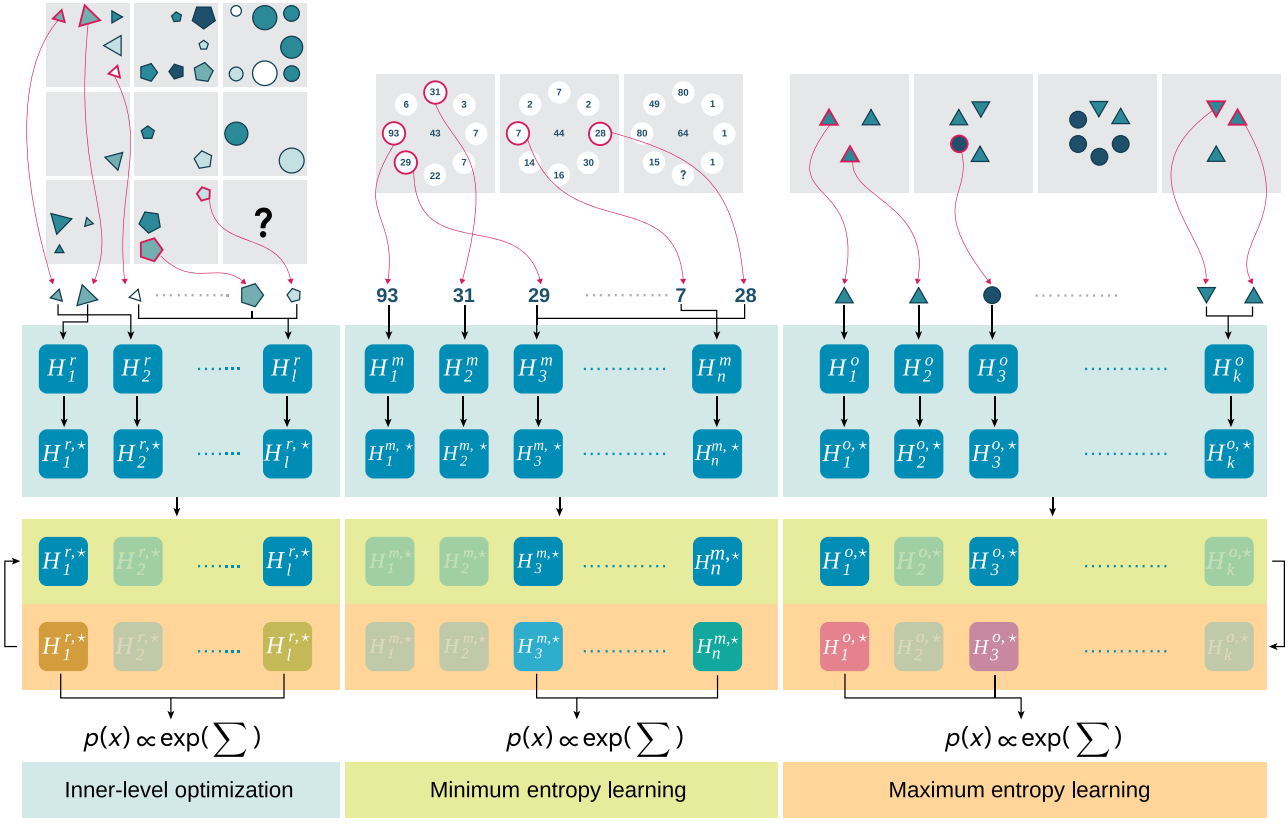

Fig. 1. The minimax entropy learning pipeline for few-shot concept induction. For the three tasks of interest (denoted by *r*, *m*, and *o* respectively), we first perform object extraction and classify their attributes using off-the-shelf methods. We then wire the object representation and perform the inner-level optimization to solve for the best filter in each filter family. Note that the wiring shown is only for illustrative purposes to avoid clutter. After learning the filters, the parameters of the filters are fixed. In the minimum entropy learning stage, we selectively pick filters that can increase the expected log likelihood, equivalent to minimizing the average coding length. Here, the filters grayed out are not chosen. In the maximum entropy stage, we find the best weights (denoted by different colors) for combining the filters. The learned filters and weights are summed and exponentiated to produce the optimal distribution that characterizes the observation.

arithmetic expressions for Machine Number Sense (MNS) (see Materials and Methods for details). We design various forms of filters that cover joint structural inference involving grouping by symmetry, hierarchy, etc.

Tasks

To evaluate the efficacy of the proposed minimax entropy model in few-shot concept induction problems, we compare our model (denoted as ME) with widely used computational approaches as well as human participants on a set of challenging abstract reasoning tasks used in IQ tests: RPM (*14*, *16*, *31*), MNS (*15*), and O^3 (see Fig. 2 for examples of the tasks).

The RPM task is composed of both context panels and choice panels. The context panels are arranged in a three-by-three matrix with the last cell missing; a test taker is asked to pick the best answer from the choice panels to complete the missing cell, such that the hidden relations are satisfied. The task is challenging in the sense that a model must be able to induce the governing relations by observing only two complete rows/columns and apply the induced rule to predict the next frame. While visual perception is simple here, the hidden relational concepts can be complex depending on how relations on different attributes are composed.

The MNS questions are designed in a similar fashion, except that numerical relations must be inferred from a set of two examples satisfying the hidden relations, which vary in geometrical layout and grouping. In addition to the two tasks above, we also design an O^3 task that is widely used in human intelligence evaluation. In this task, a model is presented with four panels and must identify the oddity from the other three with common hidden relations.

All these three tasks are well suited for evaluating few-shot concept induction: They are all designed under the same spirit—that is, one needs to first induce the hidden concepts from a limited number of context samples, either short sequences or static images, before applying it to solve the problem.

We compare our model of ME with traditional relational learning models of WReN (*13*, *32*), CoPINet (*18*), SCL (*33*), and the popular Transformer architecture (*34*, *35*). However, we stress that training our model and that for deep learning models are intrinsically different: Our model is trained unsupervisedly with the context in a single instance every time it is presented with a problem, whereas deep neural models are first trained supervisedly with a massive amount of paired input-output labels and held fixed during inference.

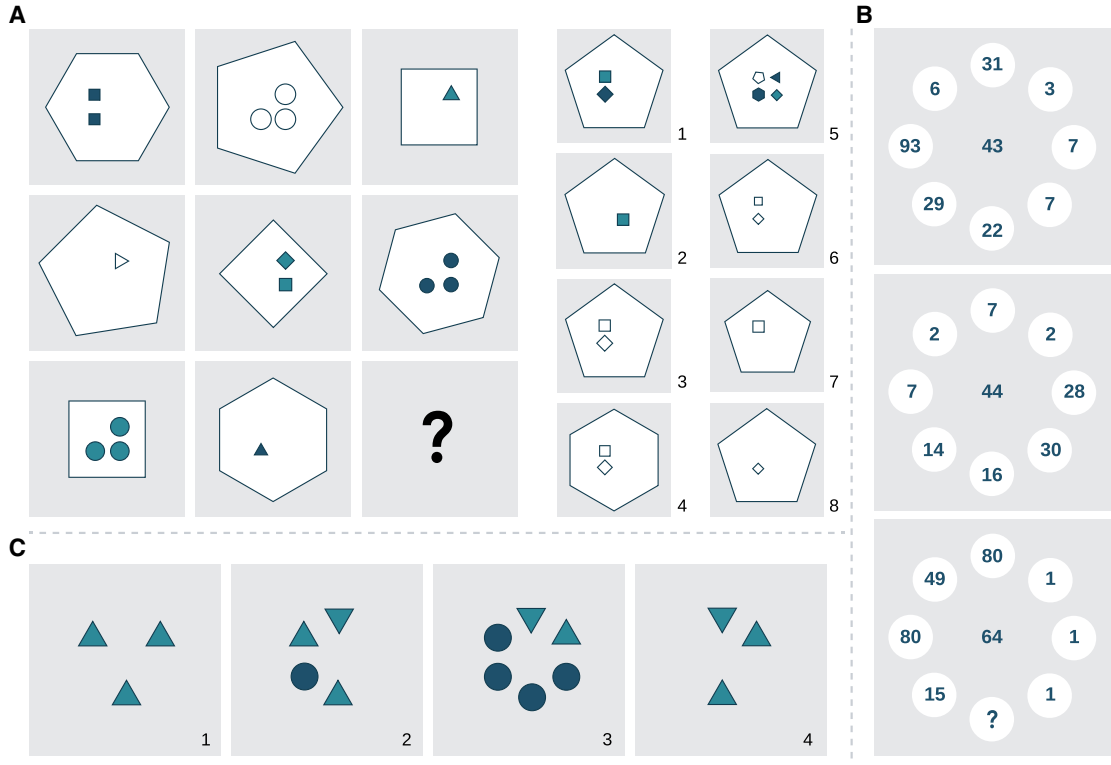

Fig. 2. Examples of RPM, MNS, and O^3. (**A**) In the RPM question, one needs to induce the hidden relations from the first two rows/columns and applies the hidden relations to the last row/column to complete the matrix. In this example, the objects outside in each row are always squares, pentagons, and hexagons and rotating from small to medium to large. Therefore, we expect a large pentagon in the missing panel. Objects inside are always triangles, squares, and circles and of the three different colors. The size of the inside objects remains the same. Besides, the number of objects inside ranges from 1 to 3. Hence, for the missing panel, we are looking for two white squares the same size of the previous inside objects. This stream of reasoning leads us to answer 3. (**B**) In the MNS task, one instead searches for the hidden numerical relations. In this example, we note that $93/31 \times 7 + 22 = 43$ and $6/3 \times 7 + 29 = 43$ in the first panel. The same relation holds in the second panel. Using this numerical expression, we can derive that the missing number should be 63. (**C**) In the O^3 task, there is only one image that is odd against others. Therefore, we can safely induce that the hidden relation should govern the triangles rather than the circles. Therefore, the answer should be 3 as others all have three triangles.

Minimax entropy learning

Instead of formulating concept induction as a discriminative process, minimax entropy learning frames the problem as a descriptive process (*22*). Specifically, for each problem's context set $C = \{x_i\}$, consisting of object-centric representation of either two to three short sequences or static images, our ME model should learn a distribution $p(x)$ on the object-centric representation space that best characterizes the hidden concept while making unconstrained dimensions as random as possible. Formally, assuming the hidden concept can be captured by a set of response functions $\{H_j(\cdot)\}$ or filters, the intuition aforementioned can be formulated as the following maximum entropy principle

$$\begin{aligned} \max_{p} \quad & -\int p(x)\log p(x)\,\mathrm{d}x \\ \text{subject to} \quad & \mathop{\mathbb{E}}_{x\sim p(x)}[H_j(x)] = \mu_j^{\text{obs}}, \forall j \\ & \int p(x)\mathrm{d}x = 1 \end{aligned} \tag{1}$$

where μ_j^{obs} denotes the average filter response on the context panels. This formulation requires the dimensions captured by the filters to match the observed statistics while setting others as unrestrained as possible. The optimization bears an analytical solution

$$p(x) = \frac{1}{Z}\exp\left[-\sum_j \lambda_j H_j(x)\right] \tag{2}$$

where $Z = \int \exp[-\sum_j \lambda_j H_j(x)]\mathrm{d}x$ is the normalizer and λ_j are the optimal Lagrangian multipliers that can be learned through maximum likelihood learning on Eq. 2 (*23*). In the minimum entropy stage, we minimize the entropy of the model on the basis of the maximum entropy results, which is equivalent to minimizing the KL divergence between the true distribution constrained by the hidden concepts $p^\star(x)$ and our approximation $p(x)$ (*23*)

$$\min_p -\int p(x)\log p(x)\mathrm{d}x = \min_p \mathbf{KL}(p^\star, p) = \max_p \mathbb{E}_{p^\star}[\log p(x)] \tag{3}$$

Note that the formulation is also equivalent to maximum likelihood as shown above. This step is implemented as selecting the optimal set of filters $\{H_j(\cdot)\}$ among others to minimize the expected coding length under the coding scheme of the chosen filters (*23*). Unlike the greedy method of feature pursuit (*23*), we explicitly add a set of global indicator variables $\{z_j\}$ to the optimization and alternatively maximize the log likelihood of the distribution, that is

$$\max_{\lambda,\mathbf{z}} \mathbb{E}_{x_i}[\log p(x_i)] = \mathbb{E}_{x_i}\left[-\sum_j \lambda_j z_j H_j(x_i) - \log Z\right] \tag{4}$$

There is still a remaining issue with the minimax entropy learning framework aforementioned: The traditional fixed filter design is limited in expressiveness and cannot adapt to different cases in distinctive scenarios. Adding more filters can potentially mitigate this issue, but a large number of filters will unnecessarily complicate the minimax learning process. Besides, for relational concepts in continuous spaces, capturing a unique one with a finite number of fixed filters is difficult. To address this issue, we further devise to parameterize the filter functions and solve for the optimal one based on the context on-the-fly. Specifically, we embed the minimax entropy learning framework inside a bilevel optimization problem (*28*–*30*): The inner-level optimization works out the optimal parameters corresponding to the best filters in continuous filter families to felicitously describe the hidden concepts, and the outer-level optimization performs the minimax entropy learning steps aforementioned. Formally, during per-instance training, we maximize the log likelihood under the constraints

$$\begin{aligned} \max_{\lambda,\mathbf{z}} \quad & \mathbb{E}_{x_i}[\log p(x_i)] = \mathbb{E}_{x_i}\left[-\sum_j \lambda_j z_j H_j(x_i;\theta_j^\star) - \log Z\right], \\ \text{subject to} \quad & \theta_j^\star = \arg\min \ell_j(\{x_i\}, \theta_j), \forall j \end{aligned} \tag{5}$$

where the optimal filter parameters $\theta_j^\star$ best capture the hidden concepts modeled by the filter family of $H_j(\cdot;\theta_j)$. Such a bilevel design could improve filter expressiveness by finding the best filter in a continuous filter family while avoiding adding unnecessarily more filters.

After training a descriptive model for each instance, we use the model to solve a problem in two ways: Given a candidate set, we can pick the one with the highest/lowest probability as the answer; without one, we can perform maximum A posteriori (MAP) sampling from the distribution to render an answer with a rendering engine.

RESULTS

To evaluate models on induction of different forms of relational concepts from a few examples, the proposed ME model, baseline models, and human performance were compared side by side on the three few-shot concept induction tasks. Apart from testing on the conventional independent and identically distributed (I.I.D.) setup where the data distribution in training and testing are the same, we also evaluate model performance on the out of distribution (O.O.D.) setup for compositional generalization. In particular, we split the datasets into three compositionality regimes based on difficulty: the easy regime, the medium regime, and the hard regime. The easy regime has the least number of either distractors, composed rules, or operands, whereas the hard regime is the most challenging, involving an excessive number of the aforementioned elements. Please refer to the Supplementary Materials for details regarding the separation. For human evaluation, we recruited 600 human participants tested on 6 RPM problems, 8 MNS problems, and 10 O^3 problems. Human participants voluntarily enlisted our tests published on the Credamo platform. We collected 300 valid responses for each set of problems. Please refer to the Supplementary Materials for details regarding human evaluation. All experiments were conducted under an approved Institutional Review Board (IRB) from Peking University.

We present the primary results for the computational models and human participants in Fig. 3 and then analyze the human data in depth in fig. S2. Performance decomposition for all the tasks is reported in the Supplementary Materials. We detail the experimental setup in the Supplementary Materials.

For the RPM task, we primarily use the RAVEN-FAIR dataset (*31*) for evaluation and defer results on RAVEN (*14*) and I-RAVEN (*16*) to the Supplementary Materials. The RPM task asks for an answer from

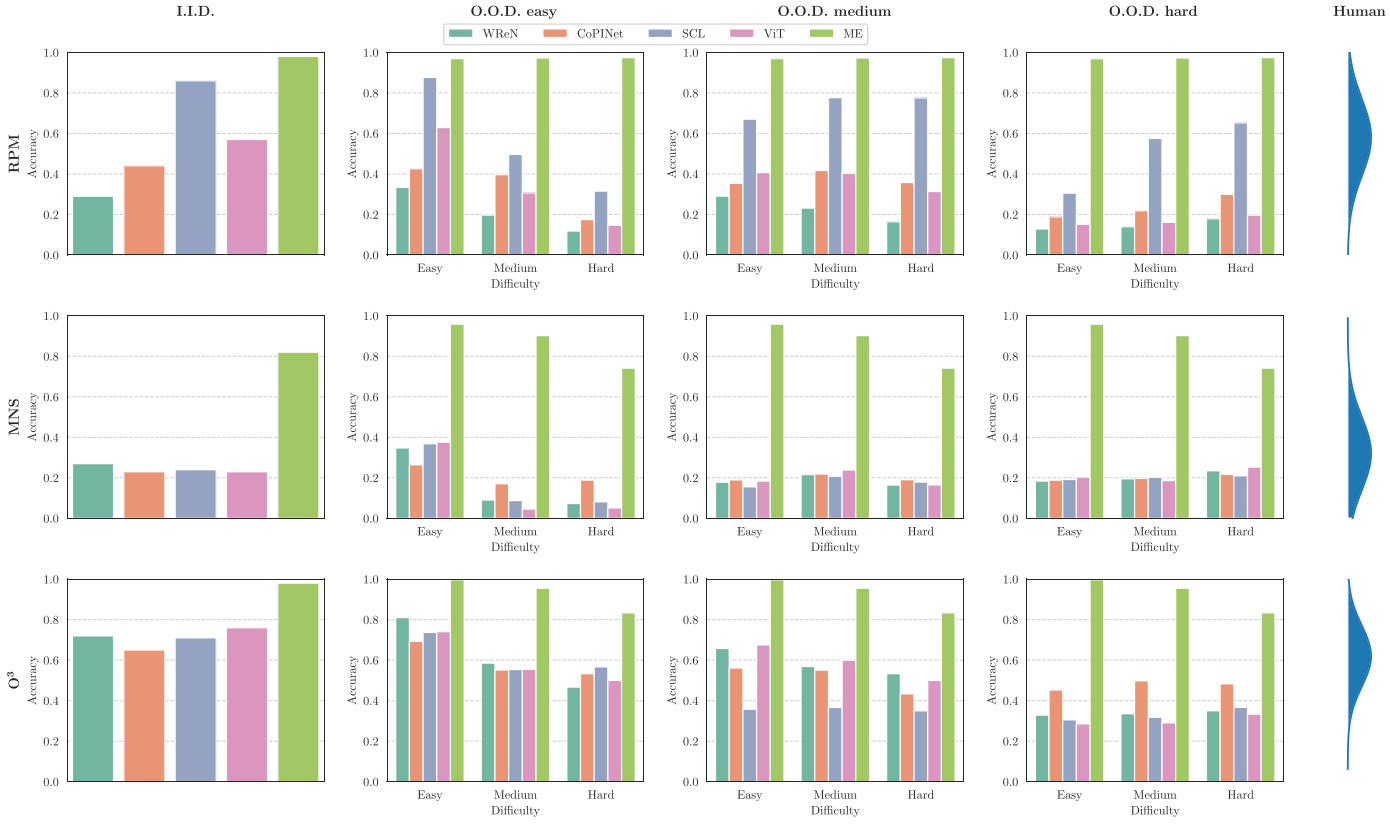

Fig. 3. Performance of different models compared to humans on the task of RPM, MNS, and O^3. The rows denote the three tasks. The first column shows the model performance on the I.I.D. setup and the following three on the O.O.D. setup. In each of the O.O.D. setup, the models are trained on the regime denoted in the title and tested on all three regimes. Note that ME is learned per-instance, and hence, the notion of generalization does not apply. We show its performance on O.O.D. regimes for comparison. The right-most axis shows the Gaussian distribution modeled from the human population.

the choice set that can complete the missing panel in the three-by-three matrix. As shown in Fig. 3, the WReN model can only correctly solve 29% of all the questions, while the SCL learner successfully addresses 86% of them. At first sight, both the neural baseline models and the proposed ME model can already work out most of the problems in RPM. The results in O.O.D. tests clearly show that neural models tend to overfit the dataset without mastering how to perform induction: In the second to the fourth columns, neural models are trained on one regime and tested on the other two. While an ideal human learner should be able to generalize to higher compositionality when learning regimes of lower compositionality (or vice versa), we note that neural models do not have this capability and when tested on other regimes, they are far worse than tested in domain. In contrast, our proposed ME model remains performative either I.I.D. or O.O.D. Note that both our ME model and human participants conduct per-instance learning; they are not trained with a specific training set but rather directly tested on the test set. As shown in Fig. 3, ME notably improves over existing baselines and outperforms a majority of human participants. Looking into the details of model performance on ME and other deep learning models in the Supplementary Materials, we find that ME is much more stable across datasets. The candidates in RAVEN are constructed by randomly perturbing an attribute in the answer panel, and therefore, a deep learning model could potentially leverage the statistics in the choice set, picking the panel whose attributes appear the most in all panels, to game the evaluation procedure. I-RAVEN and RAVEN-FAIR mitigate this issue by generating the choice set from a graph-based perturbation process (*16*, *31*). This minor change considerably affects the performance of CoPINet, whose accuracy cliffs and falls much below superhuman performance. The recent work of Zhang *et al.* (*36*) proposes an induction system, relieves the strong reliance on external knowledge from Zhang *et al.* (*21*), and greatly improves on the systematic generalization capability. However, their system is specifically designed for the RPM task and not adaptable for others like MNS and O^3. Of note, our system is comparable with their ALANS-V model and slightly improves on ALANS-V's performance on both RAVEN (94.4%) and I-RAVEN (93.5%).

Incapability of deep learning models on MNS is evident in the sense that they can reach about 20% accuracy only. In the MNS task (*15*), one needs to find the characteristic numerical structure underlying the context examples and apply it to fill the blank in the final panel. The inability of baseline models could be related to the nature of the task: A model needs to perform structural search to resolve a question. However, the potential solution space is exponentially large, especially when the number of elements increases. Deep learning models, no matter specially designed for relational learning or as efficient classifiers, simply cannot learn to perform structural search but rather just leverage statistical correlations in this task. The observation is even noticeable in the O.O.D. setup of the easy regime: Neural models can

barely solve any problems with more numbers. In general, humans also struggle on this task. ME model's performance decreases with the exponentially increasing size of the search space. By combining bilevel optimization to search for the best filter in each group and weighing the composition to describe the hidden relations, the ME model can still attain reasonably high accuracy in the hardest regime of all, surpassing humans in this task.

The task of O^3 is relatively easier compared to others. In this task, one is required to pick an odd image from the presented set of images. We note that among the deep learning models, the relational learning models of WReN, CoPINet, and SCL fare worse than the simple classifier of ViT in the I.I.D. setup. However, none of them becomes dominant in the O.O.D. regimes. The proposed ME model performs this task by learning only from the four figures, optimizing the best filter to capture the common features that can tell apart the odd sample. The model achieves the best performance on both the I.I.D. setup and the O.O.D. setup.

ME model's additional advantage points to the benefits of modeling the problem using a descriptive form: Unlike the discriminative models, descriptive models can not only be turned into discriminative but also generative. Specifically, one can directly sample from the learned distribution in each problem and render an answer to the problem without being provided the choices. Figure 4 shows the generated samples for RPM and MNS. As for O^3, the answer does not follow the hidden relations and cannot be directly generated. In general, we perform MAP sampling on the conditional distribution that is fixed on the context panels. Since the generated variables are continuous, we round them to the nearest integers. We then feed the variables into rendering engines that render the final images.

In the following, we conduct ablation study to see how much each stage of the learning process contributes to the final performance improvement. Effectively, the first stage carries out filter learning; the second stage performs minimum entropy learning, selecting the right composition of filters; the third stage weighs different filters via maximum entropy learning. As shown in Table 1, with only filter families learned and filter selection and weighing randomly initialized, our model can only achieve performance better than or comparable to chance selection. However, the minimum entropy learning stage for filter selection tremendously boosts the model accuracy, indicating that correctly picking the filters to constrain the concept space plays the most crucial role in capturing the hidden relations in the problem context. Compared to learning filter families and filter selection only, the last stage of maximum entropy learning, when incorporated,

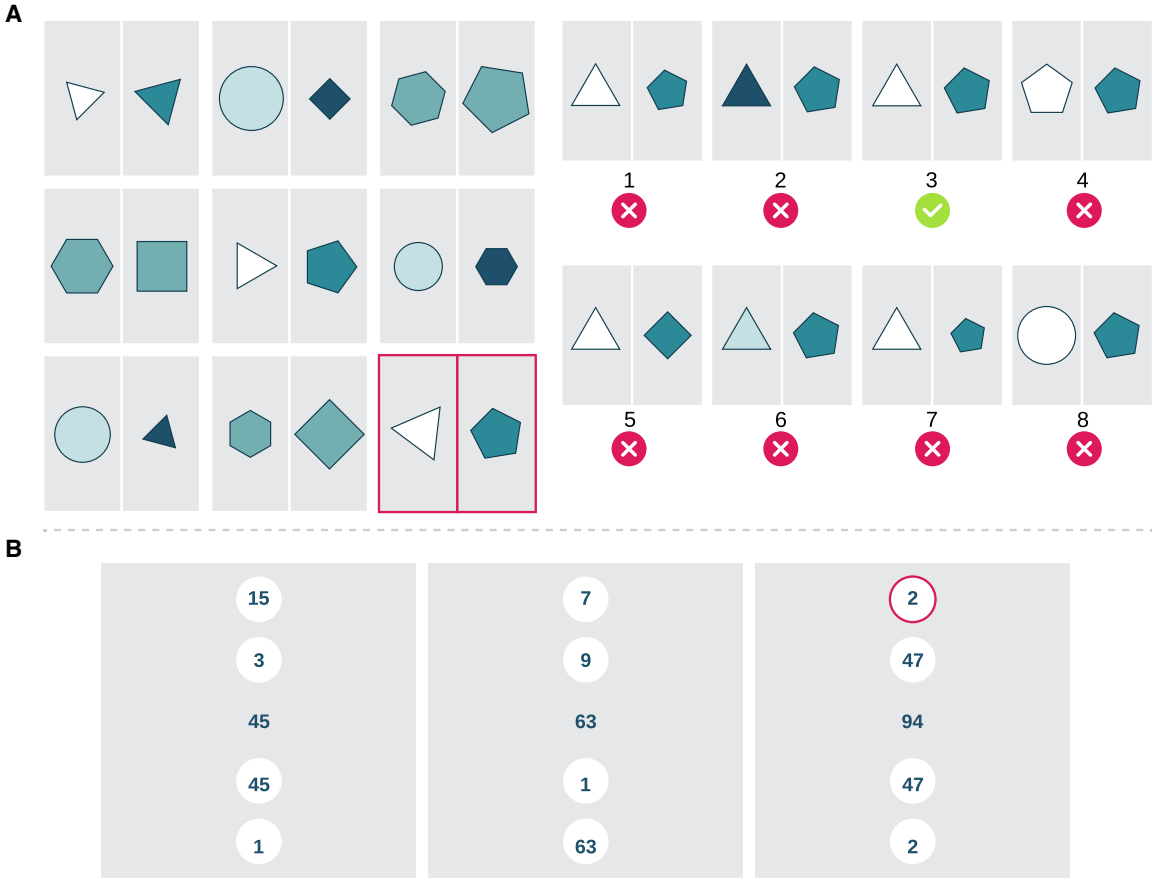

Fig. 4. Examples where the answers are generated by the ME model. (**A**) In the RPM example, we note that in the first two rows, the left part consists of triangles, hexagons, and circles and of the three different colors. Therefore, we know the left part of the answer should be a white triangle. Similarly, the right part should be a pentagon. As objects in the right part rotate from small to medium to large. The answer panel should have a medium-sized pentagon. Hence the answer. Note that the generated panel has random rotations and looks different from the candidate but still should be considered correct. (**B**) In the MNS example, the two numbers on each side multiply to the center number. Therefore, the missing slot (red circle) should be filled with 2. Our ME model generates the exact solution.

Table 1. Ablation study for the proposed ME model on RPM, MNS, and O^3. The performance is measured by accuracy on the test sets. ME FL only denotes the performance of our ME model with only the first stage of filter learning, whereas ME FL + *z* with both the first stage of filter learning and the second stage of filter selection via minimum entropy learning. The full model with the first stage of filter learning, the second stage of filter selection, and the third stage of maximum entropy learning is our final ME model. Since on O^3, we only use one filter family, and the second stage and the third stage can be combined into one. Please refer to the Supplementary Materials for detailed model performance.

Method	RPM	MNS	O^3
ME FL only	24.2%	30.2%	24.8%
ME FL + *z*	89.3%	76.3%	–
ME	97.6%	81.9%	98.2%

squeezes out the last bit of performance by properly weighing different components to compose the final distribution.

To evaluate our model of ME in a more traditional setup, we have also compared ME with other models on the task of Synthetic Visual Reasoning Test (*37*), commonly used for few-shot concept learning assessment (see the Supplementary Materials). The experimental results show that ME remains performant and data-efficient compared to pure deep learning models like ResNet (*38*) and ViT (*35*) and that ME can be transformed for the conventional supervised learning setup.

DISCUSSION

In this work, we present the ME model to solve the few-shot concept induction problems in RPM, MNS, and O^3. The model uses the descriptive minimax entropy learning principle, where the minimum entropy principle finds the tightest set of constraints that can reproduce the observed distribution and the maximum entropy principle sets the unconstrained dimensions as free as possible. The minimum entropy learning process is improved from the original greedy feature pursuit algorithm to the global optimization procedure. The model is further combined with bilevel optimization to address existing problems. By parameterizing the filter functions, we can not only expand the search space but also avoid the complexity introduced by an excessive number of filters. In the learning phase, we first optimize for the best filter in each filter family and then perform minimax entropy learning. During inference, the descriptive model of ME can be turned into both discriminative and generative. In discrimitive evaluation, our ME model markedly outperforms existing deep learning–based models and is on par with, if not better than, human participants. In addition, the ME model can also generate the final answers via conditional MAP sampling, provided with rendering engines.

We are constantly witnessing the changing landscapes in artificial intelligence: With the scale of models and datasets exponentially increasing, machines nowadays can perform some of the everyday tasks extremely well (*38*–*40*). However, humans are still much better than machines at inducing unfamiliar relational concepts: We only need a limited number of samples without supervision and can explain an unfamiliar phenomenon with the induced concept in different forms (*1*, *3*, *41*, *42*). Our work provides a piece of evidence that by properly combining bilevel optimization and minimax entropy learning, the gap between humans and machines in elementary concept induction can be well reduced. Machine learning and cognitive reasoning research have only begun to investigate the few-shot concept induction task under a big-data-small-task paradigm (*16*, *18*–*20*, *33*), whereas our model has already managed to perform at human level.

Although successful on these tasks, the ME model relies on accurate perception on visual objects. The object-centric representation is simple in these tasks as well. Nevertheless, in real-world scenarios, people can see far richer components in pictures, like changing part-whole hierarchies, symmetry, stability, aesthetics, and even motion (*4*). Such structures are common in everyday life but out of the scope in this study. Besides, our filters are mostly linear and therefore cannot well capture complex physical concepts or social dynamics.

Capturing the fascinating ability of humans should be of paramount importance to building artificial general intelligence. And we argue that understanding few-shot concept induction in humans requires us to answer developmental questions. What kinds of inductive bias do children acquire to solve similar tasks so effectively? As such tests are usually recruited to measure human intelligence, what is the factor in telling smarter people from average population? How do people quickly identify the hidden relations in a large but loosely constrained space? Is a mechanism similar to minimax entropy learning applied?

With many questions remaining to be answered, we hope this work may shed light on the small-data-big-task paradigm: Ultimately, we are striving for a learning machine that can efficiently use a general mechanism to induce the hidden concepts in various reasoning scenarios.

MATERIALS AND METHODS

Datasets

The RAVEN series of datasets (*14*, *16*) are commonly used for evaluating model performance on the RPM task. Each RAVEN instance consists of a context set and a choice set. The context set is composed of eight panels, arranged in a 3 × 3 matrix with the last one missing. To finish the task, a model needs to pick one from the eight choices to complete the missing panel. The RAVEN series of datasets are divided into seven configurations: center, 2 × 2 grid, 3 × 3 grid, left-right, up-down, in-center-out-center, and in-grid-out-center [see (*14*) for details]. For each configuration, the dataset has 10,000 data points, which are further organized into sixfold for training, twofold for validation, and twofold for testing. Each panel includes several objects with attributes of type, size, and color. On the panel-level, one can also summarize the attributes of number and object position. A predefined set of relations can govern each of the object-level attributes and the panel-level attributes: constant, progression, arithmetic, and distribute of three. Please refer to (*14*) for detailed descriptions

of the relations. The original RAVEN dataset generates the choices by perturbing only one attribute of the ground-truth answer and can cause models to learn to leverage statistical shortcuts hidden in the choice set. I-RAVEN (*16*) mitigates this issue by modeling the perturbation using a tree structure, such that models cannot use context-blind information to solve a problem. RAVEN-FAIR (*31*) further adds to the challenge by using a graph structure for answer generation.

The MNS dataset (*15*) was developed to evaluate MNS by testing how accurately models could infer the hidden numerical relations exemplified in two panels and use the relations to work out the missing number in the last one. The problems differ in problem types, configurations, layouts, and interpretation methods, totaling in 28 distinctive styles. Each style has also 10,000 data points, adding up to 280,000 samples. Train-val-test split is done as in RAVEN, with sixfold for training, twofold for validation, and twofold for testing.

The O^3 dataset introduced is created under the same framework of the RAVEN series datasets. In each instance, four panels are generated where only one is odd from others. Each panel comprises of distinctive objects, and the attributes are defined the same as in RAVEN. To characterize oddity, we use the number of a specific object: The common set of panels all have the same number of a specific object while the odd one does not. We also make sure that there is not a second interpretation to tell the difference by performing a check on every three compositions based on the criterion. We devise three different configurations and generate 10,000 data points for each one. Train-val-test split is done the same as mentioned above.

Problem representation

For representing each problem instance in RPM, MNS, and O^3, we adopt an object-centric scheme (*43*). Specifically, we detect objects and extract object attributes: For RPM and O^3, we obtain object position, type, size, and color and encode them in one-hot representation, whereas for MNS, we directly use the corresponding integer numbers.

As RPM involves spatial-temporal information, we perform minimax entropy learning on each row: Each x_i in the context panel set C is composed of three panels in a row. Using the one-hot encoding scheme for RPM, the input to our minimax entropy learning framework is a tensor of $3 \times n \times m$, where n represents the maximum number of objects among the three panels and m is the concatenated attribute space. For MNS and O^3 that only involve temporally independent frames, we directly learn from the two-dimensional object-centric representation: integer matrix for MNS and one-hot matrix for O^3.

Computing μ_j^{obs}

For RPM and MNS, the design principle is the same: We are looking for the shared relations in all given context panels. Therefore, we can directly use the mean filter response on the context panels. Formally

$$\mu_j^{\text{obs}} = \frac{1}{|C|} \sum_i H_j(x_i; \theta_j^\star) \tag{6}$$

where C is the context panel set.

O^3 is slightly different, as the odd one is among the given context panels but does not follow the relations hidden among others. Therefore, we propose the following way to estimate the mean filter response

$$\mu_j^{\text{obs}} = \frac{1}{|C| - 1} \sum_{i \in C_{-k}} H_j(x_i; \theta_j^\star) \tag{7}$$

where C_{-k} denotes the set without x_k. For the selection of k, we use the Fisher's criterion to measure the distance between two distributions (*44*, *45*). Specifically, we set the subset without x_k and the singleton set of $\{x_k\}$ as two classes and compute the ratio of between-class variance and within-class variance. The k with maximum ratio is selected. This formulation identifies the panel that produces the largest gap in the filter response and only uses the mean filter response of the common set as the statistics to match.

Filter and loss design

For RPM problems, we consider three continuous filter families corresponding to unary relations, binary relations, and ternary relations. We use linear families and minimize ℓ_2 loss to solve for the best one in each family. Specifically, for unary relations, we use

$$\frac{1}{2}(\|Ax_{i,1} + b - x_{i,2}\|_2^2 + \|Ax_{i,2} + b - x_{i,3}\|_2^2) \tag{8}$$

for binary relations, we use

$$\left\| A \begin{bmatrix} x_{i,1} \\ x_{i,2} \end{bmatrix} + b - x_{i,3} \right\|_2^2 \tag{9}$$

for ternary relations, we use

$$\left\| A \begin{bmatrix} x_{i,1} \\ x_{i,2} \\ x_{i,3} \end{bmatrix} + b \right\|_2^2 \tag{10}$$

Here, A and b are the parameters we need to solve in the filter functions. Note that here we only use linear functions and ℓ_2 loss because linear regressions admit closed form solutions and do not require an extensive amount of data points to be learned, satisfying our need of few-shot concept induction.

For MNS problems, we use a combinatorial space of basic arithmetic expressions. Specifically, we consider filter families of arithmetic expressions of length from two to six and use the ℓ_1 loss to search the best arithmetic expression in each filter family.

For O^3 problems, we also use a linear form of filter response

$$Ax_i + b \tag{11}$$

where A and b are the parameters. Combining the filter family and the statistics computation criterion, we can directly turn it into Fisher's discriminant analysis and use it for our inner-level optimization. In practice, we note that this formulation does not yield the best results due to the lack of sparsity constraints. Therefore, we reformulate Fisher's discriminant analysis as linear regression based on their equivalence (*45*), add Lasso sparsity constraints, and solve using a decent method. We then solve for the best model in all four possibilities of k, find the one with the largest gap, and pick it as our final filter.

Optimization

Optimizing Eq. 5 is challenging because gradients with respect to λ are stochastic and $\mathbf{z}$ are discrete variables. Specifically, for λ, deterministic gradients are hard to obtain: Gradients with respect to λ are derived as (*23*–*25*)

$$\nabla_{\lambda}\left(\mathbb{E}_{x_i}\left[\sum_j \lambda_j z_j H_j(x_i;\theta_j^{\star})\right]-\mathbb{E}_{x\sim p}\left[\sum_j \lambda_j z_j H_j(x;\theta_j^{\star})\right]\right) \tag{12}$$

where x follows the distribution of the learned $p(\cdot)$. To compute the second term, we perform Markov Chain Monte Carlo sampling to sample x and use the samples to approximate the mean. For optimizing $\mathbf{z}$, we could use anther stochastic optimization trick: We instantiate each z_j as a random variable following a Bernoulli distribution parameterized by ϕ_j and add another layer of averaging over the maximum likelihood objective, that is

$$\underset{\lambda,\phi}{\text{maximize}} \underset{\mathbf{z}\sim\text{Bernoulli}(\phi)}{\mathbb{E}} \mathbb{E}_{x_i}[\log p(x_i)] \tag{13}$$

In this way, optimizing $\mathbf{z}$ could be understood as reinforcement learning (*46*), where the parameterized policy outputs an action distribution and can be jointly optimized by REINFORCE (*47*). In practice, this optimization is extremely unstable as the combinatorial action space is large and computing the "reward," or the exact log likelihood, involves estimating the normalizer Z. Therefore, we use a simplified heuristics: We only select z_j that has a sufficiently small loss value. This simple strategy globally improves on the trivial zero initialization and can be understood as pruning early filters that do not look promising to increase likelihood.

Generating answers

We perform MAP sampling on the conditional distribution given the context for generation. Take the RPM problem in Fig. 4 as an example. Formally, the conditional distribution can be represented as

$$\begin{aligned} p(x_{3,3}|x_{3,1},x_{3,2}) &\propto p(x_{3,1},x_{3,2},x_{3,3}) \\ &= \frac{1}{Z}\exp\left[-\sum_j \lambda_j z_j H_j(x_{3,1},x_{3,2},x_{3,3})\right] \end{aligned} \tag{14}$$

where $x_{3,k}$ denotes the kth image in the last row. Remember that for RPM, the distribution is learned for each row. Then, to perform MAP sampling, one can directly optimize the energy term, that is

$$\begin{aligned} x_{3,3} &= \arg\max_{x_{3,3}} p(x_{3,1},x_{3,2},x_{3,3}) \\ &= \arg\min_{x_{3,3}} \sum_j \lambda_j z_j H_j(x_{3,1},x_{3,2},x_{3,3}) \end{aligned} \tag{15}$$

In practice, we use simple gradient descent to optimize the energy term. After getting $x_{3,3}$, we round the variables to the nearest integers, as it is one-hot object-centric representation. We then feed the object-centric representation into a dedicated rendering engine to generate the final image. The process for MNS is similar, except that the distribution is modeled for each panel and that we only need existing numbers in the last panel to solve for the missing one.

Supplementary Materials

This PDF file includes:

Supplementary Text
Figs. S1 and S2
Tables S1 to S5
References

REFERENCES AND NOTES

1. P. A. Carpenter, M. A. Just, P. Shell, What one intelligence test measures: A theoretical account of the processing in the raven progressive matrices test. *Psychol. Rev.* **97**, 404–431 (1990).
2. R. E. Snow, P. C. Kyllonen, B. Marshalek, The topography of ability and learning correlations. *Adv. Psychol. Hum. Intell.* **2**, 103 (1984).
3. S. M. Jaeggi, M. Buschkuehl, J. Jonides, W. J. Perrig, Improving fluid intelligence with training on working memory. *Proc. Natl. Acad. Sci. U.S.A.* **105**, 6829–6833 (2008).
4. Y. Zhu, T. Gao, L. Fan, S. Huang, M. Edmonds, H. Liu, F. Gao, C. Zhang, S. Qi, Y. N. Wu, J. Tenenbaum, S.-C. Zhu, Dark, beyond deep: A paradigm shift to cognitive ai with humanlike common sense. *Engineering* **6**, 310–345 (2020).
5. S. P. Johnson, K. J. Fernandes, M. C. Frank, N. Kirkham, G. Marcus, H. Rabagliati, J. A. Slemmer, Abstract rule learning for visual sequences in 8- and 11-month-olds. *Inf. Dent.* **14**, 2–18 (2009).
6. C. G. Lucas, S. Bridgers, T. L. Griffiths, A. Gopnik, When children are better (or at least more open-minded) learners than adults: Developmental differences in learning the forms of causal relationships. *Cognition* **131**, 284–299 (2014).
7. S. Dehaene, V. Izard, P. Pica, E. Spelke, Core knowledge of geometry in an amazonian indigene group. *Science* **311**, 381–384 (2006).
8. A. A. Wright, J. S. Katz, Mechanisms of same/different concept learning in primates and avians. *Behav. Processes* **72**, 234–254 (2006).
9. J. S. Katz, A. A. Wright, Same/different abstract-concept learning by pigeons. *J. Exp. Psychol. Anim. Behav. Processes* **32**, 80–86 (2006).
10. V. Truppa, E. Piano Mortari, D. Garofoli, S. Privitera, E. Visalberghi, Same/different concept learning by capuchin monkeys in matching-to-sample tasks. *PLOS ONE* **6**, e23809 (2011).
11. L. L. Thurstone, T. G. Thurstone, Factorial studies of intelligence. *Psychol. Monogr.* **2**, 94 (1941).
12. J. C. Raven, J. H. Court, *Raven's Progressive Matrices and Vocabulary Scales* (Oxford, UK: Oxford Pyschologists Press, 1998).
13. D. Barrett, F. Hill, A. Santoro, A. Morcos, T. Lillicrap, Measuring abstract reasoning in neural networks. in *International Conference on Machine Learning (ICML)*, July 2018 (Stockholm, Sweden2018).
14. C. Zhang, F. Gao, B. Jia, Y. Zhu, S.-C. Zhu, Raven: relational and analogical visual reasoning. in Conference on Computer Vision and Pattern Recognition (CVPR), June 2019 (Long Beach, CA, US2019).
15. W. Zhang, C. Zhang, Y. Zhu, S.-C. Zhu, Machine number sense: A dataset of visual arithmetic problems for abstract and relational reasoning. in AAAI Conference on Artificial Intelligence (AAAI), (New York, NY, US), February 2020.
16. S. Hu, Y. Ma, X. Liu, Y. Wei, S. Bai, Stratified rule-aware network for abstract visual reasoning. in AAAI Conference on Artificial Intelligence (AAAI), (Virtual), February 2021.
17. C. Zhang, B. Jia, M. Edmonds, S.-C. Zhu, Y. Zhu, Acre: Abstract causal reasoning beyond covariation. in Conference on Computer Vision and Pattern Recognition (CVPR), (Virtual), June 2021.
18. C. Zhang, B. Jia, F. Gao, Y. Zhu, H. Lu, S.-C. Zhu, Learning perceptual inference by contrasting. in Advances in Neural Information Processing Systems (NeurIPS), (Vancouver, Canada), December 2019.
19. K. Zheng, Z.-J. Zha, W. Wei, Abstract reasoning with distracting features. in Advances in Neural Information Processing Systems (NeurIPS), (Vancouver, Canada), December 2019.
20. D. Wang, M. Jamnik, P. Lio, Abstract diagrammatic reasoning with multiplex graph networks. in International Conference on Learning Representations (ICLR), (New Orleans, LA, US), May 2019.
21. C. Zhang, B. Jia, S.-C. Zhu, Y. Zhu, Abstract spatial-temporal reasoning via probabilistic abduction and execution. in Conference on Computer Vision and Pattern Recognition (CVPR), (Virtual), June 2021.
22. Y. N. Wu, R. Gao, T. Han, S.-C. Zhu, A tale of three probabilistic families: Discriminative, descriptive, generative models. *Q. Appl. Math.* **77**, 423–465 (2019).
23. S.-C. Zhu, Y. N. Wu, D. Mumford, Minimax entropy principle and its application to texture modeling. *Neural Comput.* **9**, 1627–1660 (1997).
24. Y. Lu, S.-C. Zhu, Y. Wu, Learning frame models using cnn filters. in AAAI Conference on Artificial Intelligence (AAAI), (Phoenix, AZ, US), February 2016.
25. J. Xie, Y. Lu, S.-C. Zhu, Y. Wu, A theory of generative convnet. in International Conference on Machine Learning (ICML), (New York, NY, US), June 2016.
26. J. Xie, W. Hu, S.-C. Zhu, Y. N. Wu, Learning sparse frame models for natural image patterns. *Int. J. Comput. Vis.* **114**, 91–112 (2015).

27. J. Xie, Y. Lu, S.-C. Zhu, Y. N. Wu, Inducing wavelets into random fields via generative boosting. *Appl. Comput. Harmon. Anal.* **41**, 4–25 (2016).
28. L. Franceschi, P. Frasconi, S. Salzo, R. Grazzi, M. Pontil, Bilevel programming for hyperparameter optimization and meta-learning. in International Conference on Machine Learning (ICML), (Stockholm, Sweden), June 2018.
29. B. Colson, P. Marcotte, G. Savard, An overview of bilevel optimization. *Ann. Oper. Res.* **153**, 235–256 (2007).
30. J. F. Bard, *Practical Bilevel Optimization: Algorithms and Applications* (Berlin, Germany: Springer, 2013).
31. Y. Benny, N. Pekar, L. Wolf, Scale-localized abstract reasoning. in Conference on Computer Vision and Pattern Recognition (CVPR), (Virtual), June 2021.
32. A. Santoro, D. Raposo, D. G. Barrett, M. Malinowski, R. Pascanu, P. Battaglia, T. Lillicrap, A simple neural network module for relational reasoning. in Advances in Neural Information Processing Systems (NeurIPS), (Long Beach, CA, US), December 2017.
33. Y. Wu, H. Dong, R. Grosse, J. Ba, The Scattering Compositional Learner: Discovering Objects, Attributes, Relationships in Analogical Reasoning. arXiv. 2007.04212. (2020).
34. A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, I. Polosukhin, Attention is all you need. in Advances in Neural Information Processing Systems (NeurIPS), (Long Beach, CA, US), December 2017.
35. A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, N. Houlsby, An image is worth 16x16 words: Transformers for image recognition at scale. in International Conference on Learning Representations (ICLR), (Virtual), May 2021.
36. C. Zhang, S. Xie, B. Jia, Y. N. Wu, S.-C. Zhu, Y. Zhu, Learning algebraic representation for systematic generalization in abstract reasoning. in European Conference on Computer Vision (ECCV), (Tel Aviv, Israel), October 2022.
37. F. Fleuret, T. Li, C. Dubout, E. K. Wampler, S. Yantis, D. Geman, Comparing machines and humans on a visual categorization test. *Proc. Natl. Acad. Sci. U.S.A.* **108**, 17621–17625 (2011).
38. K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition. in Conference on Computer Vision and Pattern Recognition (CVPR), (Las Vegas, Nevada, US), June 2016.
39. D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, D. Hassabis, Mastering the game of go with deep neural networks and tree search. *Nature* **529**, 484–489 (2016).
40. R. Bommasani, D. A. Hudson, E. Adeli, R. Altman, S. Arora, S. von Arx, M. S. Bernstein, J. Bohg, A. Bosselut, E. Brunskill, E. Brynjolfsson, S. Buch, D. Card, R. Castellon, N. S. Chatterji, A. S. Chen, K. A. Creel, J. Davis, D. Demszky, C. Donahue, M. Doumbouya, E. Durmus, S. Ermon, J. Etchemendy, K. Ethayarajh, L. Fei-Fei, C. Finn, T. Gale, L. E. Gillespie, K. Goel, N. D. Goodman, S. Grossman, N. Guha, T. Hashimoto, P. Henderson, J. Hewitt, D. E. Ho, J. Hong, K. Hsu, J. Huang, T. F. Icard, S. Jain, D. Jurafsky, P. Kalluri, S. Karamcheti, G. Keeling, F. Khani, O. Khattab, P. W. Koh, M. S. Krass, R. Krishna, R. Kuditipudi, A. Kumar, F. Ladhak, M. Lee, T. Lee, J. Leskovec, I. Levent, X. L. Li, X. Li, T. Ma, A. Malik, C. D. Manning, S. P. Mirchandani, E. Mitchell, Z. Munyikwa, S. Nair, A. Narayan, D. Narayanan, B. Newman, A. Nie, J. C. Niebles, H. Nilforoshan, J. F. Nyarko, G. Ogut, L. Orr, I. Papadimitriou, J. S. Park, C. Piech, E. Portelance, C. Potts, A. Raghunathan, R. Reich, H. Ren, F. Rong, Y. H. Roohani, C. Ruiz, J. Ryan, C. R'e, D. Sadigh, S. Sagawa, K. Santhanam, A. Shih, K. P. Srinivasan, A. Tamkin, R. Taori, A. W. Thomas, F. Tramèr, R. E. Wang, W. Wang, B. Wu, J. Wu, Y. Wu, S. M. Xie, M. Yasunaga, J. You, M. A. Zaharia, M. Zhang, T. Zhang, X. Zhang, Y. Zhang, L. Zheng, K. Zhou, P. Liang, On the opportunities and risks of foundation models. arXiv. 21083.07258 (2021).
41. B. M. Lake, R. Salakhutdinov, J. B. Tenenbaum, Human-level concept learning through probabilistic program induction. *Science* **350**, 1332–1338 (2015).
42. B. M. Lake, T. D. Ullman, J. B. Tenenbaum, S. J. Gershman, Building machines that learn and think like people. *Behav. Brain Sci.* **40**, e253 (2017).
43. K. Yi, J. Wu, C. Gan, A. Torralba, P. Kohli, J. Tenenbaum, Neural-symbolic vqa: Disentangling reasoning from vision and language understanding. in Advances in Neural Information Processing Systems (NeurIPS), (Montreal, Canada), December 2018.
44. R. A. Fisher, The use of multiple measurements in taxonomic problems. *Ann. Eugen.* **7**, 179–188 (1936).
45. C. M. Bishop, N. M. Nasrabadi, *Pattern Recognition and Machine Learning* (Berlin, Germany: Springer, 2006).
46. R. S. Sutton, A. G. Barto, *Reinforcement Learning: An Introduction* (Cambridge, MA, US: MIT Press, 2018).
47. R. S. Sutton, D. McAllester, S. Singh, Y. Mansour, Policy gradient methods for reinforcement learning with function approximation. in Advances in Neural Information Processing Systems (NeurIPS), (Denver, CO, US), November 1999.
48. D. P. Kingma and J. Ba, Adam: A method for stochastic optimizon. in International Conference on Learning Representations (ICLR), (San Diego, CA, US), May 2015.
49. M.-K. Hu, Visual pattern recognition by moment invariants. *IRE Trans. Inf. Theory* **8**, 179–187 (1962).
50. Yicai Global, "ranking of chinese cities' business attractiveness 2022." https://yicaiglobal.com/news/ranking-of-chinese-cities-business-attractiveness-2022, June 2022.

Acknowledgments: We thank Y. Peng (Peking University) for helping with the IRB, Z. Chen (BIGAI) for crafting our visualization, M. J. Nitzberg (UC Berkeley), L. Morris (UC Berkeley), and J. Morris (UC Berkeley) for proofreading the draft, and anonymous reviewers and editors for providing helpful guidance for revision. Last, we thank anonymous human participants for participating in the research. **Funding:** All authors were supported in part by the National Science and Technology Major Project (no. 2022ZD0114900). C.Z., B.J., and S.-C.Z. were supported in part by the Beijing Institute for General Artificial Intelligence (BIGAI) (no. KY20220060). Y.Z. was supported in part by the Beijing Nova Program. **Author contributions:** Conceptualization: C.Z., B.J., Y.Z., and S.-C.Z. Data curation: C.Z. and Y.Z. Formal analysis: C.Z. and Y.Z. Funding acquisition: Y.Z. and S.-C.Z. Investigation: C.Z. Methodology: C.Z., B.J., Y.Z., and S.-C.Z. Project administration: C.Z., Y.Z., and S.-C.Z. Software: C.Z. Supervision: Y.Z. and S.-C.Z. Validation: C.Z. and Y.Z. Visualization: C.Z. and B.J. Writing—original draft: C.Z., B.J., Y.Z., and S.-C.Z. Writing—review and editing: C.Z., B.J., and Y.Z. **Competing interests:** The authors declare that they have no competing interests. **Data and materials availability:** All data needed to evaluate the conclusions in the paper are present in the paper and/or the Supplementary Materials.

Submitted 12 December 2022
Accepted 14 March 2024
Published 19 April 2024
10.1126/sciadv.adg2488