

000  
001  
002  
003  
004  
005  
006  
007  
008  
009  
010  
011  
012  
013  
014  
015  
016  
017  
018  
019  
020  
021  
022  
023  
024  
025  
026  
027  
028  
029  
030  
031  
032  
033  
034  
035  
036  
037  
038  
039  
040  
041  
042  
043  
044  
045  
046  
047  
048  
049  
050  
051  
052  
053  

# ARC-AGI WITHOUT PRETRAINING

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Conventional wisdom in the age of LLMs dictates that solving IQ-test-like visual puzzles from the ARC-AGI-1 benchmark requires capabilities derived from massive pretraining. To counter this, we introduce *CompressARC*, a 76K parameter model without any pretraining that solves 20% of evaluation puzzles by minimizing the description length (MDL) of the target puzzle purely during inference time. The MDL endows CompressARC with extreme generalization abilities typically unheard of in deep learning. To our knowledge, CompressARC is the only deep learning method for ARC-AGI where training happens only on a single sample: the target inference puzzle itself, with the final solution information removed. Moreover, CompressARC does not train on the pre-provided ARC-AGI “training set”. Under these extremely data-limited conditions, we do not ordinarily expect any puzzles to be solvable at all. Yet CompressARC still solves a diverse distribution of creative ARC-AGI puzzles, suggesting MDL to be an alternative, highly feasible way to produce intelligence, besides conventional massive pretraining.

## 1 INTRODUCTION

The ARC-AGI-1 benchmark consists of abstract visual reasoning puzzles designed to evaluate a system’s ability to rapidly acquire new skills from minimal input data (Chollet, 2019). Recent progress in LLM-based reasoning has shown impressive skill acquisition capabilities, but these systems still rely on massive amounts of pretraining data. In this paper, we explore how little data is truly required to tackle ARC-AGI by introducing *CompressARC*, a solution method derived from the Minimum Description Length (MDL) principle (Rissanen, 1978). CompressARC performs all of its learning at inference time and achieves 20% accuracy on ARC-AGI-1 evaluation puzzles—using only the puzzle being solved as input data.

The key to CompressARC’s extreme data efficiency is its formulation as a code-golfing problem: to find the shortest possible self-contained program that outputs the entire ARC-AGI dataset, with any unsolved grids filled arbitrarily. By Occam’s razor, the shortest program is expected to contain the “correct” solutions. An overly basic program might store a hard-coded string of the puzzle data (plus arbitrary solutions) for output; but the program will be too long, implying (by Occam’s razor) the outputted solutions will be wrong. On the other hand, finding the optimally shortest program would require exhaustively enumerating many candidate programs, which is computationally infeasible (Solomonoff, 1964; Hutter, 2005). CompressARC strikes a new type of balance by overfitting a neural network to the puzzle data to compress the puzzles into weight matrices; these weights can be hard-coded into the program instead of the puzzles themselves, and then used within the program to recover the memorized puzzles, and also generate solutions. With careful counting of the bit length of the hard-coded weights, this technique converts the combinatorial search of finding the best program into a differentiable optimization problem, allowing us to minimize program length in a reasonable amount of time and still generate good solution predictions.

This framing preserves several attractive properties of the original code-golf formulation, all of which are novel when it comes to neural solutions to ARC-AGI:

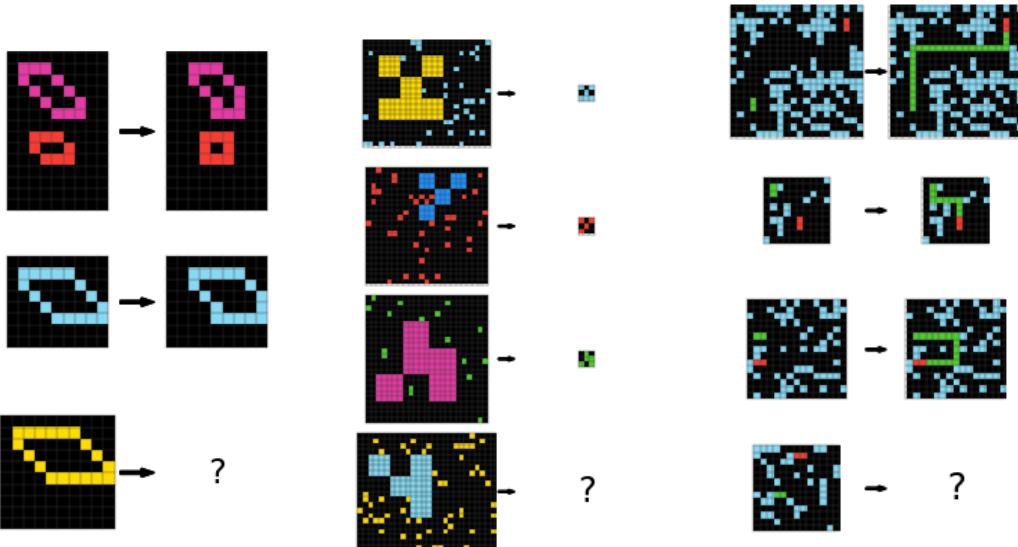
- **No pretraining:** Since we begin with the target puzzle(s) in hand, no training phase is required.
- **Inference-time learning:** Program length is minimized solely during inference by optimizing network weights with respect to the target puzzle(s).
- **Minimal data requirement:** Following Occam’s razor, we assume strong generalization from the shortest program and use only the puzzle(s) themselves—no additional data is loaded into memory.

Despite never using the training set, performing no pretraining, and having only 76K parameters in its network, CompressARC generalizes strongly, solving 20% of evaluation puzzles and 34.75% of training puzzles—performance that would be impossible for traditional deep learning methods under these constraints. [CompressARC’s strong performance in this setting suggests that bringing code-golf to other data-limited contexts beyond ARC-AGI \(e.g., drug discovery, protein design\) may help us extract stronger capabilities in those applications as well.](#)

The remainder of this paper introduces the ARC-AGI benchmark (Section 2), details the problem framing (Section 3), describes CompressARC’s architecture (Section 4), presents empirical results (Section 5), and concludes with a discussion of implications (Section 6).

## 2 BACKGROUND: THE ARC-AGI BENCHMARK

ARC-AGI-1 is an artificial intelligence benchmark designed to test a system’s ability to acquire new skills from minimal examples (Chollet, 2019). Each puzzle in the benchmark consists of a different hidden rule, which the system must apply to an input colored grid to produce a ground truth target colored grid. The hidden rules make use of themes like objectness, goal-directedness, numbers & counting, basic geometry, and topology. Several input-output grid pairs are given as examples to help the system figure out the hidden rule in the puzzle, and no other information is given. [Figure 1 shows three examples of ARC-AGI-1 training puzzles.](#)



(a) **Hidden rule:** Shift every object right by one pixel, except the bottom/right edges of the object.  
 (b) **Hidden rule:** Shrink the big object and set its color to the scattered dots’ color.  
 (c) **Hidden rule:** Extend the green line to meet the red line by turning when hitting a wall.

Figure 1: Three example ARC-AGI-1 puzzles.

While solutions based on LLMs built using internet scale data have scored 87.5% on this benchmark (Chollet, 2024), and neural networks trained on only ARC-AGI data have scored 40.3% (Wang et al., 2025), CompressARC takes the data-minimal perspective to its limits, opting to only train on the test puzzle.

Please refer to Appendix L for more details about the ARC-AGI-1 benchmark. An extended survey of other related work is also included in Appendix A. Note that we will generally refer to ARC-AGI-1 just as ARC-AGI in this paper.

108 3 METHOD  
109

110 Occam’s razor instructs that the shortest possible program whose return value matches the puzzle  
 111 will likely also return the correct solution with it. Thus, we frame ARC-AGI as a code-golfing  
 112 problem: to find the shortest possible program that reproduces the ARC-AGI dataset (Rissanen,  
 113 1978). In this case, the code must be entirely self-contained, receive no inputs, and must print out  
 114 the entire ARC-AGI dataset of puzzles with any solutions filled in. Each puzzle takes the form of  
 115 a tensor of shape  $[n_{\text{example}}, \text{width}, \text{height}, 2]$ , containing color designations for every pixel in the  
 116  $2 \times n_{\text{example}}$  grids. Here,  $n_{\text{example}}$  counts the total number of input/output grid pairs in the  
 117 puzzle, including the test grid pair whose output grid is not known. The shapes listed in this section  
 118 are for explanatory purposes and the actual data format is introduced in Section 4. A naive first try  
 119 at code-golfing the dataset may involve writing a program that hard-codes each puzzle in a giant  
 120 string and prints it out. Improvements can be made by clever ways of de-duplicating structures in the  
 121 printed data (e.g., introducing for loops, etc.); we will detail our own particular strategy below.  
 122

123 3.1 RESTRICTING THE PROGRAM SPACE  
124

125 It is typically infeasible to run an algorithmic search to solve code-golf problems, because we would  
 126 have to search through a huge number of increasingly lengthy programs to find one whose printout  
 127 matches our requirements. Despite this, search can be made more amenable for our ARC-AGI  
 128 code-golf problem if we restrict ourselves to a suitably well-conditioned subspace of programs. A  
 129 formal derivation of CompressARC begins by picking a program subspace consisting of a template  
 130 program (Algorithm 1) to be completed by substituting various hard-coded values into designated  
 131 locations (shown in red). The template program generates each puzzle independently, and performs  
 132 the operations for every puzzle:

- 133 1. it randomly samples a tensor  $z$  of shape  $[n_{\text{example}}, n_{\text{colors}}, \text{width}, \text{height}, 2]$  from a standard  
 134 normal distribution, (line 4)
- 135 2. processes this with a neural network which outputs a  $[n_{\text{example}}, n_{\text{colors}}, \text{width}, \text{height}, 2]$ -shaped logit tensor, (line 6)
- 136 3. and obtains a  $[n_{\text{example}}, \text{width}, \text{height}, 2]$ -shaped puzzle by sampling colors from the probability  
 137 distribution implied by the logit tensor. (line 8)

141  
 142 **Algorithm 1:** Template for a short program that produce completed puzzles  $P_{\text{filled}}$  with solutions  
 143 filled in. Red text is to be substituted in with hard-coded values produced via Algorithm 2.

---

144 1 Define an equivariant\_NN architecture;  
 145  
 146 2 **Set**  $\text{seed\_z} = \langle \text{seed\_z}_1 \rangle$ ; Hardcoded seed from Algo 2, puzzle 1  
 147 4  $z \leftarrow \text{sample}_{\text{seed\_z}}(N(0, 1))$ ; Generate inputs  $z$   
 148 5 **Set**  $\theta = \langle \theta_1 \rangle$ ; Hardcoded weights from Algo 2, puzzle 1  
 149 6  $\text{grid\_logits} \leftarrow \text{equivariant\_NN}_{\theta}(z)$ ; Forward pass  
 150 7 **Set**  $\text{seed\_error} = \langle \text{seed\_error}_1 \rangle$ ; Hardcoded seed from Algo 2, puzzle 1  
 151 8  $P_{\text{filled}} \leftarrow \text{sample}_{\text{seed\_error}}(\text{grid\_logits})$ ; Generate puzzle  
 152 9 **Print**  $P_{\text{filled}}$   
 153  
 154 10 **Set**  $\text{seed\_z} = \langle \text{seed\_z}_2 \rangle$ ; Hardcoded seed from Algo 2, puzzle 2  
 155 11 (...code repeats for all puzzles) ...  
 156

---

157 The template allows for two pseudo-random sampling seeds to be filled in for every puzzle (lines  
 158 3 and 7). The resulting printed puzzles can be guaranteed to match the true puzzles in the dataset  
 159 by manipulating the final seed in line 7, which works after choosing any seed on line 3. With this  
 160 guarantee in place, we can sum up the length of the code for the template, and find that the total  
 161 length varies based on the number of bits/digits required to write down the two seeds. So, in order to  
 search for short programs, we just need to make all the seeds as short as possible.

162 Multiple areas of the program can be adjusted to help minimize the seed length, and we will cover  
 163 each in respective sections: the seeds and the weights on lines 3, 5, and 7 (Section 3.2 below), and  
 164 the architecture on line 1 (Section 4).  
 165

166 **3.2 SEED OPTIMIZATION**  
 167

168 Algorithm 2 presents a **method of** optimizing the seeds and weights in template Algorithm 1 to  
 169 reduce the total seed length. It first tries to manipulate the seed on line 3 of the template to imitate  
 170  $z$  being sampled from a different learned normal distribution (line 8), and then tries to manipulate  
 171 the second seed to guarantee matching puzzle output (line 11). It then performs gradient descent on  
 172 the normal distribution parameters and the neural network weights to minimize the total seed length  
 173 (lines 13-14).  
 174

175 **Algorithm 2:** Minimize Description Length, a.k.a. code-golf. `n_example` denotes the total  
 176 number of input/output pairs in the puzzle, *including the test pair where the output is unknown*.  
 177 Line 11: The smallest possible `seed_error` is picked so that the sampled puzzle  $P_{\text{filled}}$  on line 12  
 178 matches the true puzzle  $P$  on both the input and output grids of demonstration pairs, as well as  
 179 the test inputs, but with no restriction on the test outputs.

---

180 **1 Input:** ARC-AGI dataset;  
 181 2 Define an equivariant\_NN architecture;  
 182 3 **foreach** puzzle  $P$  in ARC-AGI dataset **do**  
 183 4   **Randomly initialize** weights  $\theta$  for equivariant\_NN $_{\theta}$ ;  
 184 5   **Observe** the dimensions `n_example`, `n_colors`, `width`, `height` of puzzle  $P$ ;  
 185   Initialize input distribution:  $\mu$  of shape [`n_example`, `n_colors`, `width`, `height`, 2], diagonal  $\Sigma$ ;  
 186   **foreach** step **do**  
 187     | Set `seed_z`, by manipulating to imitate  $z \sim N(\mu, \Sigma)$ ;  
 188     |  $z \leftarrow \text{sample}_{\text{seed\_z}}(N(0, I), \text{shape}=[\text{n\_example}, \text{n\_colors}, \text{width}, \text{height}, 2])$ ;  
 189     |  $\text{grid\_logits} \leftarrow \text{equivariant\_NN}_{\theta}(z)$ ;  
 190     | Set `seed_error`, by manipulating to obtain desired puzzle  $P$ ;  
 191     |  $P_{\text{filled}} \leftarrow \text{sample}_{\text{seed\_error}}(\text{grid\_logits})$ ;  
 192     |  $L \leftarrow \text{len}(\text{seed\_z}) + \text{len}(\text{seed\_error})$ ;  
 193     |  $\mu, \Sigma, \theta \leftarrow \text{Adam}(\nabla_{\mu} L, \nabla_{\Sigma} L, \nabla_{\theta} L)$ ;  
 194     |  $\approx \text{len}(\text{Algo 1}) + C$   
 195   **end foreach**  
 196   Insert values `seed_z`,  $\theta$ , and `seed_error` into the pseudo-code for Algo 1;  
 197 **end foreach**  
 198 **Return** code for Algo 1

---

199 The idea of “manipulating the sampling seed for  
 200 one distribution to imitate sampling from another  
 201 distribution” tends to be fraught with technicalities  
 202 and subtleties. Under favorable conditions, this  
 203 is cleanly achievable through rejection sampling  
 204 (Forsythe, 1972), which produces a sample from  
 205 the imitated distribution using a seed whose ex-  
 206 pected length is the max log probability ratio be-  
 207 tween the two distributions. This can subsequently  
 208 be improved and extended towards more general  
 209 conditions using Relative Entropy Coding (REC)  
 210 (Harsha et al., 2010; Havasi et al., 2018; Flamich  
 211 et al., 2021), where the expected seed length lowers  
 212 to the KL divergence between the two distributions,  
 213 at the cost of only approximately achieving the de-  
 214 sired sampling distribution. We refer readers to these  
 215 references for details.

Our main issue with seed manipulation using REC  
 within Algorithm 2 on lines 8 and 11 is that running

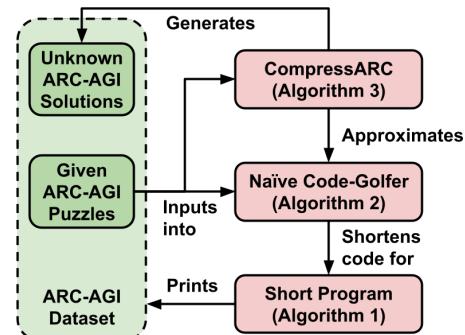


Figure 2: CompressARC approximates a specific code-golfing algorithm that converts the ARC-AGI puzzle dataset into the shortest program that prints it out exactly, along with any solutions. These printed solutions are assumed to be good predictors of the actual solutions, according to Occam’s razor.

REC takes a lot of time when the imitated distribution is far from the sampling distribution (Flamich et al., 2021). So to make it faster, we skip these steps and imitate their expected downstream consequences instead, resulting in CompressARC (Algorithm 3). Namely for line 8,  $z$  is now directly sampled from the imitated distribution, and the seed length from line 13 is replaced by its expected length, which is very close to the KL divergence between the imitated and sampling distributions according to the properties of REC (Flamich et al., 2021). For line 11, the unknown grids of the puzzle are sampled directly, and the seed length from line 13 is approximated closely by the negative log likelihood of sampling the known grids exactly, i.e., the crossentropy (see Appendix B for derivation).

225 **Algorithm 3:** CompressARC. It is the same as Algorithm 2, but with simulated seed manipulation,  
226 and truncated to return solved puzzles instead of description.

```

227 1 Input: ARC-AGI dataset;
228 2 Define an equivariant_NN architecture;
229 3 foreach puzzle  $P$  in ARC-AGI dataset do
230 4   Randomly initialize weights  $\theta$  for equivariant_NN $_{\theta}$ ;
231 5   Observe the dimensions n_example, n_colors, width, height of puzzle  $P$ ;
232 6   Initialize input distribution:  $\mu$  of shape [n_example, n_colors, width, height, 2], diagonal  $\Sigma$ ;
233 7   foreach step do
234 8      $z \leftarrow \text{sample}(N(\mu, \Sigma))$ ;
235 9     grid_logits  $\leftarrow \text{equivariant\_NN}_{\theta}(z)$ ;
236 10     $L \leftarrow \text{KL}(N(\mu, \Sigma) || N(0, I)) + \text{crossentropy}(\text{grid\_logits}, P)$ ;  $\approx \text{len}(\text{Algo 1}) + C$ 
237 11     $\mu, \Sigma, \theta \leftarrow \text{Adam}(\nabla_{\mu} L, \nabla_{\Sigma} L, \nabla_{\theta} L)$ ;
238 12  end foreach
239 13   $P_{\text{filled}} \leftarrow \text{sample}(\text{grid\_logits})$ ;
240 14  Add  $P_{\text{filled}}$  to solved puzzles
241 15 end foreach
242 16 Return solved puzzles

```

Algorithm 3 (CompressARC) is now able to automatically code-golf the ARC-AGI dataset through successive refinement of template Algorithm 1, outputting solutions afterward. The only remaining component to specify is the neural network architecture used within template Algorithm 1, which we will design by hand. Since the architecture definition only appears once in Algorithm 1 while seeds appear repeatedly for every puzzle, using a sophisticated architecture can help us shorten the length of the template Algorithm 1, by trading off architecture description length to allow for shorter seeds. This serves as the primary motivation for us to heavily engineer the neural network architecture.

## 4 ARCHITECTURE

The central idea in designing the neural network architecture is to create a high probability of sampling the ARC-AGI puzzles, consequently reducing the length of the seeds and by extension the length of template Algorithm 1. According to the template structure, this means we need the neural network to have good inductive biases for transforming noise into reasonable-seeming ARC-AGI puzzles.<sup>1</sup>

Since ARC-AGI puzzles would be just as likely to appear in any combination of input/output example orderings, colors, orientations, etc., we want our network to assign them all equal probability by default. So, we made our architecture equivariant to example permutations, color permutations, rotations, and flips; (Cohen & Welling, 2016a) guaranteeing that [computations applied to transformed inputs result in equivalently transformed outputs](#). For any asymmetries a puzzle may have, we require Algorithm 3 to manipulate the seed of the random input  $z$ , to bias the outputted puzzle one way or another.

The architecture, shown in Figure 3, consists of a decoding layer functioning like an embedding matrix (details in Appendix D.1), followed by a core with a residual backbone, followed by a linear readout on the channel dimension (see Appendix D.8). In the core, linear projections on the channel dimension read data from the residual into specially designed operations, which write their outputs

<sup>1</sup>The training puzzles play no role in our method other than to boost our efforts to better engineer the inductive biases into our layers.

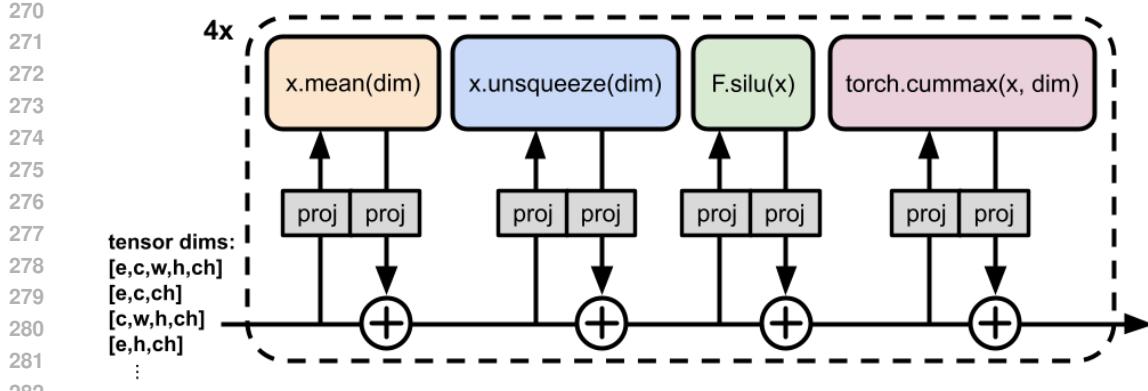


Figure 3: Core structure of CompressARC’s neural network, which operates on multitensor data. Individual operations (colored) read and write to a residual backbone through learned projections (grey) in the channel dimension. The network is equivariant to permutations of indices along the other, non-channel dimensions as a result. Some layers like cummax break certain geometric symmetries, giving the architecture specific geometric abilities listed in Appendix H. Normalization, softmax, shift, and directional layers are not shown.

back into the residual through another projection. Normalization operations are scattered throughout the layers, and then the whole block of core layers is repeated 4 times. This is much like a transformer architecture, (Vaswani et al., 2023) except that the specially designed operations are not the attention and nonlinear operations on sequences, but instead the following operations on puzzle-shaped data:

- summing one tensor along an axis and/or broadcasting the result back up, (see Appendix D.2)
- taking the softmax along one or multiple axes of a tensor, (see Appendix D.3)
- taking a cumulative max along one of the geometric dimensions of a tensor, (see Appendix D.4)
- shifting by one pixel along one of the geometric dimensions of a tensor, (see Appendix D.4)
- elementwise nonlinearity, (see Appendix D.6)
- normalization along the channel dimension, (see Appendix D.7)

along with one more described in Appendix D.5. The operations have no parameters and have their behaviors controlled by their residual read/write weights. All of these read/write projections operate on the channel dimension. We used 16 channels in some parts of the backbone and 8 in others to reduce computational costs. Since these projections take up the majority of the model weights, the entire model only has 76K parameters.

**The actual data format** that the neural network uses for computation is not a single tensor shaped like [n\_example, n\_colors, width, height, channel], but instead a bucket of tensors that each have a different subset of these dimensions, for example a [n\_colors, width, channel] tensor. Both the input  $z$  to the network and the outputted logits, as well as all of the internal activations, take the form of a multitensor. Generally, there is a tensor for every subset of these dimensions for storing information of that shape, which helps to build useful inductive biases. For example, an assignment of grid columns to colors can be stored as a one-hot table in the [color, width, channel]-shaped tensor. More details on multitensors are in Appendix C.

## 5 RESULTS

We gave CompressARC 2000 [inference-time training](#) steps on every puzzle, taking about 20 minutes per puzzle. CompressARC [correctly](#) solved 20% of evaluation set puzzles and 34.75% of training set puzzles [within this budget of inference-time compute](#). Figure 4 shows the performance increase as more inference-time compute is given. Tables 4 and 5 in the Appendix document the numerical solve accuracies with timings.

Table 1: Comparison of solution accuracies of various methods for ARC-AGI-1, sorted by the amount of training data used. Each method makes two solution guesses per puzzle, and a guess is only correct if the grid shape and pixel colors are all correct. The U-Net (Ronneberger et al., 2015) baseline is a supervised model trained during inference time on only the demonstration input/output pairs of grids in the test puzzle to match the constraints of CompressARC; details in Appendix I.

| Method                                  | Trained on:          | Neural   | Acc.       | Dataset split      |
|---|----------------------|----------|------------|--------------------|
| Random guessing                         | Nothing              | ✗        | 0%         | All                |
| Brute force rule search (Kamradt, 2024) | Nothing              | ✗        | 40%        | Private Eval       |
| U-Net baseline                          | Target puzzle        | ✓        | 0.75%      | Public Eval        |
| <b>CompressARC (ours)</b>               | <b>Target puzzle</b> | <b>✓</b> | <b>20%</b> | <b>Public Eval</b> |
| HRM ablation (ARC Prize Team, 2025)     | Test puzzles         | ✓        | 31%        | Public Eval        |
| HRM (Wang et al., 2025)                 | Train+test puzzles   | ✓        | 40.3%      | Public Eval        |
| OpenAI o3 high (Chollet, 2024)          | Internet scale data  | ✓        | 87.5%      | Semi-Priv. Eval    |

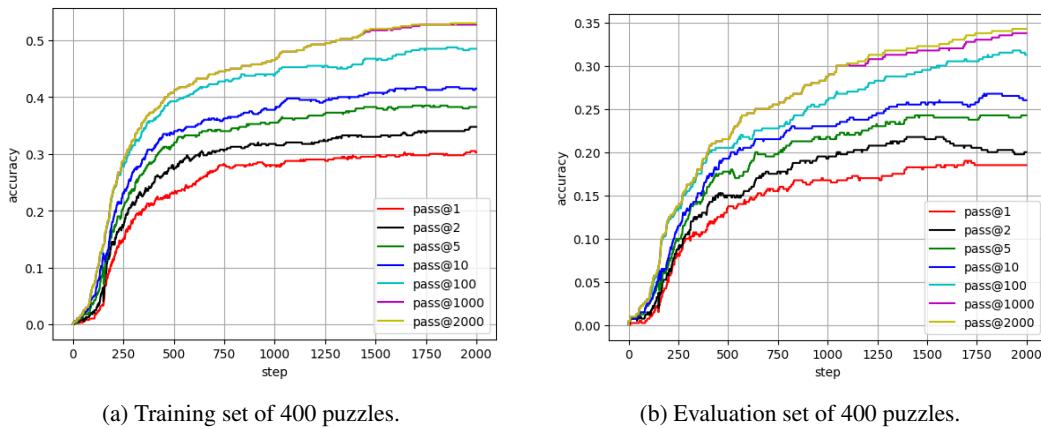


Figure 4: CompressARC’s puzzle solve accuracy rises as inference time learning progresses. Various numbers of allowed solution [guesses](#) (pass@n) for accuracy measurement are shown. The official benchmark is reported with 2 allowed [guesses](#), which is why we report 20% on the evaluation set.

## 5.1 WHAT PUZZLES CAN AND CAN'T WE SOLVE?

**CompressARC tries to use its abilities to figure out as much as it can, until it gets bottlenecked by one of its inabilities.**

For example, puzzle 28e73c20 in the training set requires extension of a pattern from the edge towards the middle, as shown in Figure 12a in the Appendix. Given the layers in its network, CompressARC is generally able to extend patterns for short ranges but not long ranges. So, it does the best that it can, and correctly extends the pattern a short distance before guessing at what happens near the center (Figure 12b, Appendix). Appendix H includes a list of which abilities we have empirically seen CompressARC able to and not able to perform.

## 5.2 CASE STUDY: COLOR THE BOXES

In the puzzle shown (Figure 5), you must color the boxes depending on which side of the grid the box is on. We call this puzzle “Color the Boxes”.

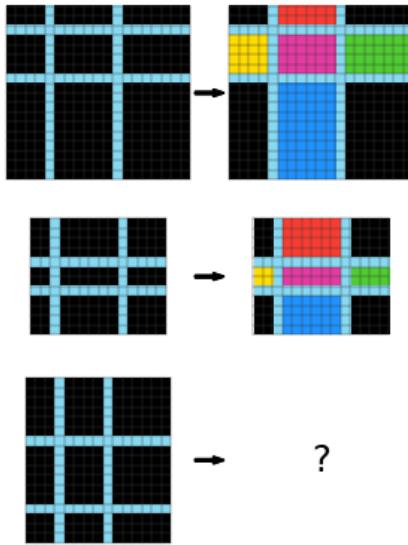
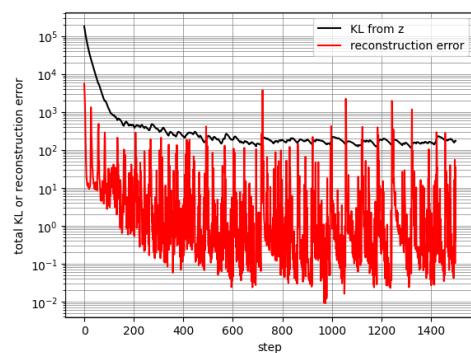


Figure 5: Color the Boxes,  
puzzle 272f95fa.

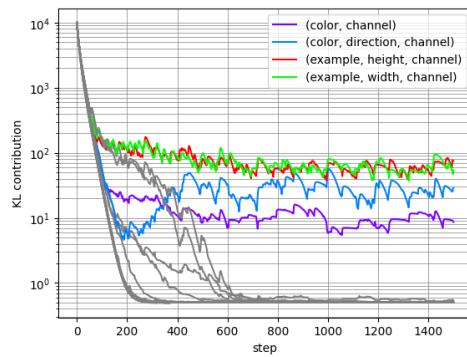
378     **Human Solution:** We first realize that the input is divided into boxes, and the boxes are still there in  
 379     the output, but now they're colored. We then try to figure out which colors go in which boxes. First,  
 380     we notice that the corners are always black. Then, we notice that the middle is always magenta. And  
 381     after that, we notice that the color of the side boxes depends on which direction they are in: red for  
 382     up, blue for down, green for right, and yellow for left. At this point, we copy the input over to the  
 383     answer grid, then we color the middle box magenta, and then color the rest of the boxes according to  
 384     their direction.

385     **CompressARC Solution:** Table 2 shows CompressARC's learning behavior over time. After  
 386     CompressARC is done learning, we can deconstruct its learned  $z$  distribution to find that it codes for  
 387     a color-direction correspondence table and row/column divider positions (Figure 7).

388     During training, the reconstruction error fell extremely quickly. It remained low on average, but  
 389     would spike up every once in a while, causing the KL from  $z$  to bump upwards at these moments, as  
 390     shown in Figure 6a.



405     (a) Relative proportion of the KL and reconstruc-  
 406     tion terms to the loss during training, before taking  
 407     the weighted sum. The KL dominates the loss and  
 408     reconstruction is most often nearly perfect.



409     (b) Breaking down the KL loss during training into  
 410     contributions from each individual shaped tensor  
 411     in the multitensor  $z$ . Four tensors dominate, indi-  
 412     cating they contain information, and the other 14  
 413     fall to zero, indicating their lack of information  
 414     content.

415     Figure 6: Breaking down the loss components during training tells us where and how CompressARC  
 416     prefers to store information relevant to solving a puzzle.

### 5.2.1 SOLUTION ANALYSIS

417     We observe the representations stored in  $z$  to see how CompressARC learns to solve Color the Boxes.

418     Since  $z$  is a multitensor, each of the tensors it contains produces an additive contribution to the total  
 419     KL for  $z$ . By looking at the per-tensor contributions (see Figure 6b), we can determine which tensors  
 420     in  $z$  code for information that is used to represent the puzzle.

421     All the tensors fall to zero information content during training, except for four tensors. In some  
 422     replications of this experiment, we saw one of these four necessary tensors fall to zero information  
 423     content, and CompressARC typically does not recover the correct answer after that. Here we are  
 424     showing a lucky run where the [color, direction, channel] tensor almost falls but gets picked up 200  
 425     steps in, which is right around when the samples from the model begin to show the correct colors in  
 426     the correct boxes.

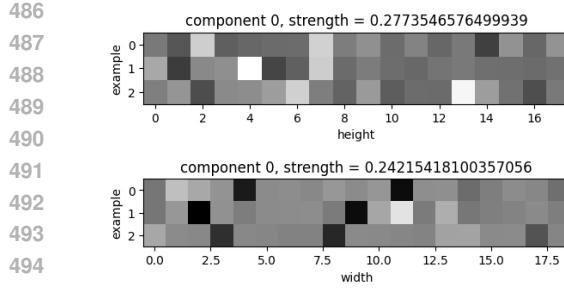
427     We can look at the average output of the decoding layer (explained in Appendix D.1) corresponding  
 428     to individual tensors of  $z$ , to see what information is stored there (see Figure 7). Each tensor contains  
 429     a vector of dimension  $n\_channels$  for various indices of the tensor. Taking the PCA of these vectors  
 430     reveals some number of activated components, telling us how many pieces of information are coded  
 431     by the tensor.

432  
433  
434  
435  
436

Table 2: CompressARC learning the solution for Color the Boxes, over time.

437  
438  
439  
440  
441  
442  
443  
444  
445  
446  
447  
448  
449  
450  
451  
452  
453  
454  
455  
456  
457  
458  
459  
460  
461  
462  
463  
464  
465  
466  
467  
468  
469  
470  
471  
472  
473  
474  
475  
476  
477  
478  
479  
480  
481  
482  
483  
484  
485

| Learning steps | What is CompressARC doing?   | Sampled solution guess |                |
|----------------|--|------------------------|----------------|
| 50             | CompressARC's network outputs an answer grid (sample) with light blue rows/columns wherever the input has the same. It has noticed that all the other input-output pairs in the puzzle exhibit this correspondence. It doesn't know how the other output pixels are assigned colors; an exponential moving average of the network output (sample average) shows the network assigning mostly the same average color to non-light-blue pixels.  | sample                 | sample average |
| 150            | The network outputs a grid where nearby pixels have similar colors. It has likely noticed that this is common among all the outputs, and is guessing that it applies to the answer too.  | sample                 | sample average |
| 200            | The network output now shows larger blobs of colors that are cut off by the light blue borders. It has noticed the common usage of borders to demarcate blobs of colors in other outputs, and applies the same idea here. It has also noticed black corner blobs in other given outputs, which the network imitates.   | sample                 | sample average |
| 350            | The network output now shows the correct colors assigned to boxes of the correct direction from the center. It has realized that a single color-to-direction mapping is used to pick the blob colors in the other given outputs, so it imitates this mapping. It is still not the best at coloring within the lines, and it is also confused about the center blob, probably because the middle does not correspond to a direction. Nevertheless, the average network output does show a tinge of the correct magenta color in the middle, meaning the network is catching on. | sample                 | sample average |
| 1500           | The network is as refined as it will ever be. Sometimes it will still make a mistake in the sample it outputs, but this uncommon and filtered out.   | sample                 | sample average |



(a) **(example, height, channel) and (example, width, channel) tensors.** For every example and row/column, there is a vector of dimension  $n_{\text{channels}}$ . Taking the PCA of this set of vectors, the top principal component (>1000 times stronger than the other components for both tensors) visualized as the (example, height) and (example, width) matrices shown above tells us which example/row and example/column combinations are uniquely identified by the stored information. **For every example, the two brightest pixels in the top matrix give positions of the light blue rows in the puzzle grids, and the darkest two pixels in the bottom matrix indicate the columns.**

Figure 7: Breaking down the loss components during training tells us where and how CompressARC prefers to store information relevant to solving a puzzle.

## 6 DISCUSSION

The prevailing reliance of modern deep learning on high-quality data has put the field in a chokehold when applied to problems requiring intelligent behavior that have less data available. This is especially true for the data-limited ARC-AGI benchmark, where LLMs trained on specially augmented, extended, and curated datasets dominate (Knoop, 2024). In the midst of this circumstance, we built CompressARC, which not only uses no training data at all, but forgoes the entire process of pretraining altogether. One should intuitively expect this to fail and solve no puzzles at all, but by applying MDL to the target puzzle during inference time, CompressARC solves a surprisingly large portion of ARC-AGI-1.

CompressARC’s theoretical underpinnings come from minimizing the length of a programmatic description of the target puzzle. While other MDL search strategies have been scarce due to the intractably large search space of possible programs, CompressARC explores a simplified, neural network-based search space through gradient descent. Though CompressARC’s architecture is heavily engineered, its incredible ability to generalize from as low as two demonstration input/output pairs puts it in an entirely new regime of generalization for ARC-AGI.

Efficiency improvement remains a valuable direction for future work on CompressARC. CompressARC makes use of many custom operations (See Appendices C and D), and adding JIT-compiled kernels or fused CUDA kernels would increase the training iteration speed. Improvements will naturally have a larger effect on larger grids since our architecture’s runtime scales with the number of pixels in the puzzle.

We challenge the assumption that intelligence must arise from massive pretraining and data, showing instead that clever use of MDL and compression principles can lead to surprising capabilities. We use CompressARC as a proof of concept to demonstrate that modern deep learning frameworks can be melded with MDL to create a possible alternative, complimentary route to AGI.

Note: Large Language Models (LLMs) were used to polish the writing in this paper, in particular to find the most clear and concise way of introducing our work in Section 1.

540 REFERENCES  
541

542 Ekin Akyürek, Mehul Damani, Adam Zweiger, Linlu Qiu, Han Guo, Jyothish Pari, Yoon Kim, and Jacob Andreas.  
543 The surprising effectiveness of test-time training for few-shot learning. *arXiv preprint arXiv:2411.07279*,  
544 2024.

545 ARC Prize Team. The hidden drivers of hrm’s performance on arc-agi. <https://arcprize.org/blog/hrm-analysis>, August 2025. Accessed: 2025-11-24.

546 Guillermo Barbadillo. Solution summary for arc24. [https://ironbar.github.io/arc24/05\\_Solution\\_Summary/](https://ironbar.github.io/arc24/05_Solution_Summary/), 2024. Accessed: 2025-05-12.

547 Guillermo Barbadillo. Exploring the combination of search and learn for the arc25 challenge,  
548 November 2025. URL [https://ironbar.github.io/arc25/05\\_Solution\\_Summary/#introduction-what-is-arc-and-why-is-it-relevant](https://ironbar.github.io/arc25/05_Solution_Summary/#introduction-what-is-arc-and-why-is-it-relevant). Accessed: 2025-11-22.

549 Jeremy Berman. How i got a record 53.6% on arc-agi, December 2024. URL <https://jeremyberman.substack.com/p/how-i-got-a-record-536-on-arc-agi>. Blog post on Substack.

550 Clément Bonnet and Matthew V Macfarlane. Searching latent program spaces, 2024. URL <https://arxiv.org/abs/2411.08706>.

551 J. Bretagnolle and C. Huber. Estimation des densités : Risque minimax. In C. Dellacherie, P. A. Meyer,  
552 and M. Weil (eds.), *Séminaire de Probabilités XII*, pp. 342–363, Berlin, Heidelberg, 1978. Springer Berlin  
553 Heidelberg. ISBN 978-3-540-35856-5.

554 Vinod Kumar Chauhan, Jiandong Zhou, Ping Lu, Soheila Molaei, and David A. Clifton. A brief review of  
555 hypernetworks in deep learning. *Artificial Intelligence Review*, 57(9), August 2024. ISSN 1573-7462. doi:  
556 10.1007/s10462-024-10862-8. URL <http://dx.doi.org/10.1007/s10462-024-10862-8>.

557 François Chollet. On the measure of intelligence, 2019. URL <https://arxiv.org/abs/1911.01547>.

558 François Chollet. Openai o3 breakthrough high score on arc-agi-pub. <https://arcprize.org/blog/oai-o3-pub-breakthrough>, 2024. Accessed: 2025-05-12.

559 Taco Cohen and Max Welling. Group equivariant convolutional networks. In Maria Florina Balcan and Kilian Q.  
560 Weinberger (eds.), *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of  
561 *Proceedings of Machine Learning Research*, pp. 2990–2999, New York, New York, USA, 20–22 Jun 2016a.  
562 PMLR. URL <https://proceedings.mlr.press/v48/cohen16.html>.

563 Taco S. Cohen and Max Welling. Group equivariant convolutional networks, 2016b. URL <https://arxiv.org/abs/1602.07576>.

564 Jack Cole and Mohamed Osman. Don’t throw the baby out with the bathwater: How and why deep learning for  
565 arc. *arXiv preprint arXiv:2506.14276*, 2025.

566 Shiqing Fan, Liu Liying, and Ye Luo. An alternative practice of tropical convolution to traditional convolutional  
567 neural networks, 2021. URL <https://arxiv.org/abs/2103.02096>.

568 Gergely Flamich, Marton Havasi, and José Miguel Hernández-Lobato. Compressing images by encoding  
569 their latent representations with relative entropy coding, 2021. URL <https://arxiv.org/abs/2010.01185>.

570 George E. Forsythe. Von neumann’s comparison method for random sampling from the normal and other  
571 distributions. *Mathematics of Computation*, 26(120):817–826, 1972. ISSN 00255718, 10886842. URL  
572 <http://www.jstor.org/stable/2005864>.

573 Alex Graves, Greg Wayne, and Ivo Danihelka. Neural turing machines, 2014. URL <https://arxiv.org/abs/1410.5401>.

574 Ryan Greenblatt. Getting 50% (sota) on arc-agi with gpt-4o. <https://redwoodresearch.substack.com/p/getting-50-sota-on-arc-agi-with-gpt>, 2024. Accessed: 2025-05-12.

575 Prahladh Harsha, Rahul Jain, David McAllester, and Jaikumar Radhakrishnan. The communication complexity  
576 of correlation. *IEEE Transactions on Information Theory*, 56(1):438–449, 2010. doi: 10.1109/TIT.2009.  
577 2034824.

578 Marton Havasi, Robert Peharz, and José Miguel Hernández-Lobato. Minimal random code learning: Getting  
579 bits back from compressed model parameters. *arXiv preprint arXiv:1810.00440*, 2018.

594 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition, 2015.  
 595 URL <https://arxiv.org/abs/1512.03385>.

596

597 Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus), 2023. URL <https://arxiv.org/abs/1606.08415>.

598

599 Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed,  
 600 and Alexander Lerchner. beta-VAE: Learning basic visual concepts with a constrained variational framework.  
 601 In *International Conference on Learning Representations*, 2017. URL <https://openreview.net/forum?id=Sy2fzU9g1>.

602

603 Michael Hodel. Domain specific language for the abstraction and reasoning corpus. [https://github.com/michaelhodel/arc-dsl/blob/main/arc\\_dsl\\_writeup.pdf](https://github.com/michaelhodel/arc-dsl/blob/main/arc_dsl_writeup.pdf), 2024. Accessed: 2025-05-12.

604

605 Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu  
 606 Chen. Lora: Low-rank adaptation of large language models, 2021. URL <https://arxiv.org/abs/2106.09685>.

607

608 Marcus Hutter. *The Universal Algorithmic Agent AIXI*, pp. 141–183. Springer Berlin Heidelberg, Berlin,  
 609 Heidelberg, 2005. ISBN 978-3-540-26877-2. doi: 10.1007/3-540-26877-4\_5. URL [https://doi.org/10.1007/3-540-26877-4\\_5](https://doi.org/10.1007/3-540-26877-4_5).

609

610

611 Marcus Hutter. Hutter prize for lossless compression of human knowledge. <https://prize.hutter1.net/>, 2006. Accessed: 2025-05-12.

612

613

614 Greg Kamradt. Arc prize 2024 solution: 4th place score 40. Kaggle Note-  
 615 book, 2024. URL <https://www.kaggle.com/code/gregkamradt/arc-prize-2024-solution-4th-place-score-40-811b72>. Accessed: 2025-11-22.

615

616

617 Diederik P Kingma and Max Welling. Auto-encoding variational bayes, 2022. URL <https://arxiv.org/abs/1312.6114>.

618

619 Mike Knoop. ARC Prize 2024 Winners & Technical Report Published — arcprize.org. <https://arcprize.org/blog/arc-prize-2024-winners-technical-report>, 2024. [Accessed 12-05-2025].

620

621

622 A.N. Kolmogorov. On tables of random numbers. *Theoretical Computer Science*, 207(2):387–395,  
 623 1998. ISSN 0304-3975. doi: [https://doi.org/10.1016/S0304-3975\(98\)00075-9](https://doi.org/10.1016/S0304-3975(98)00075-9). URL <https://www.sciencedirect.com/science/article/pii/S0304397598000759>.

623

624

625 G. G. Langdon. An introduction to arithmetic coding. *IBM Journal of Research and Development*, 28(2):  
 135–149, 1984. doi: 10.1147/rd.282.0135.

626

627 Solim LeGris, Wai Keen Vong, Brenden M. Lake, and Todd M. Gureckis. H-arc: A robust estimate of human  
 628 performance on the abstraction and reasoning corpus benchmark, 2024. URL <https://arxiv.org/abs/2409.01374>.

629

630 Wen-Ding Li, Keya Hu, Carter Larsen, Yuqing Wu, Simon Alford, Caleb Woo, Spencer M. Dunn, Hao Tang,  
 631 Michelangelo Naim, Dat Nguyen, Wei-Long Zheng, Zenna Tavares, Yewen Pu, and Kevin Ellis. Combining  
 632 induction and transduction for abstract reasoning, 2024a. URL <https://arxiv.org/abs/2411.02272>.

633

634 Wen-Ding Li, Keya Hu, Carter Larsen, Yuqing Wu, Simon Alford, Caleb Woo, Spencer M Dunn, Hao Tang,  
 635 Michelangelo Naim, Dat Nguyen, et al. Combining induction and transduction for abstract reasoning. *arXiv*  
 636 preprint arXiv:2411.02272, 2024b.

637

638 Victor Vikram Odouard. Arc-solution\_documentation. [https://github.com/victorvikram/ARC-icecuber/blob/master/ARC-solution\\_documentation.pdf](https://github.com/victorvikram/ARC-icecuber/blob/master/ARC-solution_documentation.pdf), 2024. Accessed: 2025-  
 639 05-12.

640

641 J. Rissanen. Modeling by shortest data description. *Automatica*, 14(5):465–471, 1978. ISSN 0005-1098.  
 642 doi: [https://doi.org/10.1016/0005-1098\(78\)90005-5](https://doi.org/10.1016/0005-1098(78)90005-5). URL <https://www.sciencedirect.com/science/article/pii/0005109878900055>.

643

644 Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image  
 645 segmentation. In *International Conference on Medical image computing and computer-assisted intervention*,  
 646 pp. 234–241. Springer, 2015.

647

C. Shannon. The zero error capacity of a noisy channel. *IRE Transactions on Information Theory*, 2(3):8–19,  
 1956. doi: 10.1109/TIT.1956.1056798.

648 C. E. Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27(3):379–423,  
 649 1948. doi: 10.1002/j.1538-7305.1948.tb01338.x.

650 R.J. Solomonoff. A formal theory of inductive inference. part i. *Information and Control*, 7(1):1–22,  
 651 1964. ISSN 0019-9958. doi: [https://doi.org/10.1016/S0019-9958\(64\)90223-2](https://doi.org/10.1016/S0019-9958(64)90223-2). URL <https://www.sciencedirect.com/science/article/pii/S0019995864902232>.

653 Yu Sun, Xiaolong Wang, Zhuang Liu, John Miller, Alexei A. Efros, and Moritz Hardt. Test-time training with  
 654 self-supervision for generalization under distribution shifts, 2020. URL <https://arxiv.org/abs/1909.13231>.

656 Casper Kaae Sønderby, Tapani Raiko, Lars Maaløe, Søren Kaae Sønderby, and Ole Winther. Ladder variational  
 657 autoencoders, 2016. URL <https://arxiv.org/abs/1602.02282>.

659 Alexandre B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer Publishing Company, Incorporated,  
 660 1st edition, 2008. ISBN 0387790519.

661 Arash Vahdat and Jan Kautz. Nvae: A deep hierarchical variational autoencoder, 2021. URL <https://arxiv.org/abs/2007.03898>.

664 Aaron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural discrete representation learning, 2018.  
 665 URL <https://arxiv.org/abs/1711.00937>.

666 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser,  
 667 and Illia Polosukhin. Attention is all you need, 2023. URL <https://arxiv.org/abs/1706.03762>.

668 Guan Wang, Jin Li, Yuhao Sun, Xing Chen, Changling Liu, Yue Wu, Meng Lu, Sen Song, and Yasin Abbasi  
 669 Yadkori. Hierarchical reasoning model. *arXiv preprint arXiv:2506.21734*, 2025.

671 Ruibin Xiong, Yunchang Yang, Di He, Kai Zheng, Shuxin Zheng, Chen Xing, Huishuai Zhang, Yanyan  
 672 Lan, Liwei Wang, and Tie-Yan Liu. On layer normalization in the transformer architecture, 2020. URL  
 673 <https://arxiv.org/abs/2002.04745>.

674 Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Ruslan Salakhutdinov, and Alexander  
 675 Smola. Deep sets, 2018. URL <https://arxiv.org/abs/1703.06114>.

676

677

678

679

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702 A RELATED WORK  
703704 A.1 EQUIVALENCE OF COMPRESSION AND INTELLIGENCE  
705706 The original inspiration of this work came from the Hutter Prize (Hutter, 2006), which awards a prize  
707 for those who can compress a file of Wikipedia text the most, as a motivation for researchers to build  
708 intelligent systems. It is premised upon the idea that the ability to compress information is equivalent  
709 to intelligence.710 This equivalence between intelligence and compression has a long history. For example, when  
711 talking about intelligent solutions to prediction problems, the ideal predictor implements Solomonoff  
712 Induction, a theoretically best possible but uncomputable prediction algorithm that works universally  
713 for all prediction tasks (Solomonoff, 1964). This prediction algorithm is then equivalent to a best  
714 possible compression algorithm whose compressed code length is the Kolmogorov Complexity of the  
715 data (Kolmogorov, 1998). This prediction algorithm can also be used to decode a description of the  
716 data of minimal length, linking these formulations of intelligence to MDL (Rissanen, 1978). In our  
717 work, we try to approximate this best possible compression algorithm with a neural network.718  
719 A.2 INFORMATION THEORY AND CODING THEORY  
720721 Since we build an information compression system, we make use of many results in information  
722 theory and coding theory. The main result required to motivate our model architecture is the existence  
723 of Relative Entropy Coding (REC) (Flamich et al., 2021). The fact that REC exists means that as  
724 long as a KL divergence can be bounded, the construction of a compression algorithm is always  
725 possible and the issue of realizing the algorithm can be abstracted away. Thus, problems about coding  
726 theory and translating information from Gaussians into binary and back can be ignored, since we can  
727 figure out the binary code length directly from the Gaussians instead. In other words, we only need  
728 to do enough information theory using the Gaussians to get the job done, with no coding theory at  
729 all. While the existence of arithmetic coding (Langdon, 1984) would suffice to abstract the problem  
730 away when distributions are discrete, neural networks operate in a continuous space so we need REC  
731 instead.732 Our architecture sends  $z$  information through an additive white Gaussian noise (AWGN) channel,  
733 so the AWGN channel capacity formula (Gaussian input Gaussian noise) plays a heavy role in the  
734 design of our decoding layer (Shannon, 1948).735 A.3 VARIATIONAL AUTOENCODERS  
736737 The decoder side of the variational autoencoder (Kingma & Welling, 2022) serves as our decompre-  
738 sition algorithm. While we would use something that has more general capabilities like a neural Turing  
739 machine (Graves et al., 2014) instead, neural Turing machines are not very amenable to gradient  
740 descent-based optimization so we stuck with the VAE.741 VAEs have a long history of developments that are relevant to our work. At one point, we tried using  
742 multiple decoding layers to make a hierarchical VAE decoder (Sønderby et al., 2016) instead. This  
743 does not affect the KL calculation because a channel capacity with feedback is equal to the channel  
744 capacity without feedback (Shannon, 1956). But, we found empirically that the first decoding layer  
745 would absorb all of the KL contribution, making the later decoding layers useless. Thus, we only  
746 used one decoding layer at the beginning.747 The beta-VAE (Higgins et al., 2017) introduces a reweighting of the reconstruction loss to be stronger  
748 than the KL loss, and we found that to work well in our case. The NVAE applies a non-constant  
749 weighting to loss components (Vahdat & Kautz, 2021). A rudimentary form of scheduled loss  
750 recombination is used in CompressARC.751  
752 A.4 OTHER ARC-AGI METHODS  
753754 Top-scoring methods to solve ARC-AGI rely on converting puzzle grids into text and then feeding  
755 them into a pretrained large language model which is prompted to find the solution. The predominant  
techniques involve either using the LLM to output the solution grid directly (Li et al., 2024b; Cole

& Osman, 2025; Akyürek et al., 2024), or output a program that can be run to manipulate the grids instead(Li et al., 2024b; Greenblatt, 2024; Barbadillo, 2025; Berman, 2024). Oftentimes, these methods employ several tricks to improve performance:

- Fine-tuning on training puzzle data
  - Applying data augmentation to increase the effective number of puzzles to fine-tune on (Akyürek et al., 2024)
  - Fine-tuning on synthetic data (Li et al., 2024a; Akyürek et al., 2024)
- Employing inference-time training approaches
  - Fine-tuning an individual model specific to each test puzzle, during test time (Akyürek et al., 2024)
  - Test-time training (TTT) techniques (Sun et al., 2020; Barbadillo, 2024)
- Sampling many model outputs or random augmentations of the test puzzle, for ensembling (Cole & Osman, 2025; Greenblatt, 2024)
- LLM reasoning (Chollet, 2024)

Such methods have managed to score up to 87.5% on the semi-private split of ARC-AGI, at a cost of over \$200 equivalent of inference-time compute per puzzle (Chollet, 2024). These approaches all make use of language models that were pretrained on the entire internet, which is in contrast to CompressARC, whose only training data is the test puzzle. Their commonalities with CompressARC are mainly in the emphasis on training individual models on individual test puzzles and the use of ensembling to improve solution predictions.

Aside from these methods, several other methods have been studied:

- An older class of methods consists of hard-coded, large-scale searches through program spaces in hand-written domain-specific languages designed specifically for ARC (Hodel, 2024; Odouard, 2024). **While these methods do not use neural networks to solve puzzles and are less instructive towards the field of machine learning, they share the commonality of using heavily engineered components designed specifically for ARC-AGI.**
- (Bonnet & Macfarlane, 2024) introduced a VAE-based method for searching through a latent space of programs. This is the most similar work to ours that we found due to their VAE setup.

## A.5 DEEP LEARNING ARCHITECTURES

We designed our own neural network architecture from scratch, but not without borrowing crucial design principles from many others.

Our architecture is fundamentally structured like a transformer, consisting of a residual stream where representations are stored and operated upon, followed by a linear head (Vaswani et al., 2023; He et al., 2015). Pre-and post-norms with linear up- and down-projections allow layers to read and write to the residual stream (Xiong et al., 2020). The SiLU-based nonlinear layer is especially similar to a transformer’s (Hendrycks & Gimpel, 2023).

Our equivariance structures are inspired by permutation-invariant neural networks, which are a type of equivariant neural network (Zaheer et al., 2018; Cohen & Welling, 2016b). Equivariance transformations are taken from common augmentations to ARC-AGI puzzles.

## B SEED LENGTH ESTIMATION BY KL AND CROSSENTROPY

In Section 3.1, we estimate  $\text{len}(\text{seed\_z})$  in line 12 of template Algorithm 2 as  $\text{KL}(N(\mu, \Sigma) || N(0, I))$ , and  $\text{len}(\text{seed\_error})$  as  $\text{crossentropy}(\text{grid\_logits}, P)$ . In this section, we will argue for the reasonability of this approximation. Readers may also refer to Flamich et al. (2021), which introduces a better seed manipulation method as “Relative Entropy Coding” (REC). Flamich et al. (2021) shows that seed communication is effectively the most bit-efficient way for an encoder and decoder to communicate samples from a distribution  $Q$  if there is a shared source of randomness  $P$ . This uses nearly  $\text{KL}(P || Q)$  expected bits per sample communicated. We urge readers to refer to Flamich

et al. (2021) for details regarding the manipulation procedure, runtime and memory analysis, and approximation strength. Below, we follow with our own effort at reasoning through why.

To recap, we original procedure in Algorithm 2 manipulates the seed for sampling  $z \sim N(0, I)$  to simulate as though  $z \sim N(\mu, \Sigma)$ , and we would like to show that we can closely approximate this sampling using an expected number of seed bits close to  $\text{KL}(N(\mu, \Sigma) || N(0, I))$ .

For sake of illustration, suppose for instance that Algorithm 2 implements something similar to rejection sampling, (Forsythe, 1972) iterating through seeds one by one and accepting the sample with probability  $\min(1, cw(z))$  for some  $c \ll 1$ , where  $w(z)$  is the probability ratio

$$w(z) = \frac{N(z; \mu, \Sigma)}{N(z; 0, I)}$$

When we pick a small enough  $c$ , the sampling distribution becomes arbitrarily close to  $N(\mu, \Sigma)$  as desired. With this  $c$ , we would like to show that the expected number of rejections leads us to end up with a seed length close to the KL.

We would first like to lower bound the probability  $P_{\text{accept}}$  of accepting at each step, which is

$$\begin{aligned} P_{\text{accept}} &= \int N(z; 0, I) \min(1, cw(z)) dz \\ &= \int N(z; 0, I) \min\left(1, \frac{cN(z; \mu, \Sigma)}{N(z; 0, I)}\right) dz \\ &= \int \min(N(z; 0, I), cN(z; \mu, \Sigma)) dz \end{aligned}$$

We will follow a modified version of a derivation of the Bretagnolle–Huber inequality (Bretagnolle & Huber, 1978) by Tsybakov (2008) to derive a bound on the KL:

$$\begin{aligned} (1 + c)P_{\text{accept}} &\geq (1 + c - P_{\text{accept}})P_{\text{accept}} \\ &= \left( \int \max(N(z; 0, I), cN(z; \mu, \Sigma)) dz \right) \left( \int \min(N(z; 0, I), cN(z; \mu, \Sigma)) dz \right) \end{aligned}$$

where applying the Cauchy-Schwarz inequality with a function space inner product,

$$\begin{aligned} &\geq \left( \int \sqrt{\max(N(z; 0, I), cN(z; \mu, \Sigma)) \min(N(z; 0, I), cN(z; \mu, \Sigma))} dz \right)^2 \\ &= \left( \int \sqrt{cN(z; 0, I)N(z; \mu, \Sigma)} dz \right)^2 \\ &= c \exp\left(2 \ln \int \sqrt{N(z; 0, I)N(z; \mu, \Sigma)} dz\right) \\ &= c \exp\left(2 \ln \int N(z; \mu, \Sigma) \sqrt{\frac{N(z; 0, I)}{N(z; \mu, \Sigma)}} dz\right) \\ &= c \exp\left(2 \ln \mathbb{E}_{z \sim N(z; \mu, \Sigma)} \left[ \sqrt{\frac{N(z; 0, I)}{N(z; \mu, \Sigma)}} \right]\right) \end{aligned}$$

and following with Jensen's inequality,

$$\begin{aligned} &\geq c \exp\left(\mathbb{E}_{z \sim N(z; \mu, \Sigma)} \left[ \ln \frac{N(z; 0, I)}{N(z; \mu, \Sigma)} \right]\right) \\ &= c \exp(-\text{KL}(N(\mu, \Sigma) || N(0, I))) \end{aligned}$$

leads to an acceptance probability of at least

$$P_{\text{accept}} \geq \frac{c}{1 + c} \exp(-\text{KL}(N(\mu, \Sigma) || N(0, I)))$$

864 Therefore, according to the rejection sampling procedure, the number of samples proposed (i.e. the  
 865 expected seed) is at most the inverse of this acceptance probability,  
 866

$$867 \text{seed\_z} \leq \frac{1+c}{c} \exp(\text{KL}(N(\mu, \Sigma) || N(0, 1)))$$

869 so the expected seed length is at most around the logarithm,

$$870 \text{len}(\text{seed\_z}) \leq \text{KL}(N(\mu, \Sigma) || N(0, 1)) + \log(1+c) - \log c$$

871 matching up with our stated **KL** approximation of the seed length.  
 872

873 For the seed\_error term, Algorithm 2 manipulates the seed to sample a puzzle  $P$  from a  
 874 distribution implied by some logits. This is effectively the same as sampling  $\text{grid\_logits} \sim$   
 875  $\text{Categorical\_distribution}(\text{logits})$  and manipulating the seed to try to get  $\text{grid\_logits} \sim$   
 876  $\text{Delta\_distribution}(P)$ . Then, the same **KL**-based bound on required seed length can be used once  
 877 again. The expected seed\_error length is at most

$$878 \text{KL}(\text{Delta\_distribution}(P) || \text{Categorical\_distribution}(\text{logits})) + \log(1+c) - \log c$$

879 which simplifies as  
 880

$$881 \mathbb{E}_{x \sim \text{Delta\_distribution}(P)} \left[ \log \frac{\delta(x = P)}{\text{Categorical\_probability}(x; \text{logits})} \right] + \log(1+c) - \log c$$

$$882 = \log \frac{\delta(P = P)}{\text{Categorical\_probability}(P; \text{logits})} + \log(1+c) - \log c$$

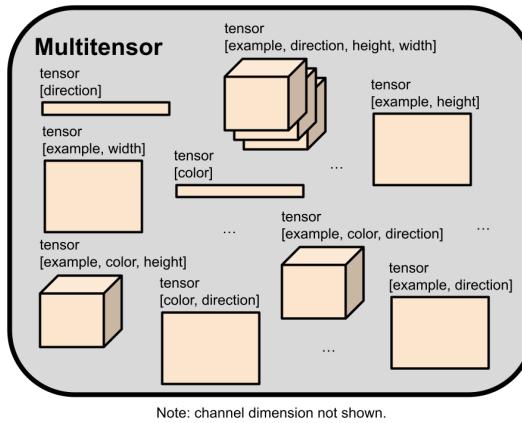
$$883 = -\log \text{Categorical\_probability}(P; \text{logits}) + \log(1+c) - \log c$$

$$884 = \text{cross\_entropy}(\text{logits}, P) + \log(1+c) - \log c$$

885 where the  $\delta$  is 1 when the statement within is true, and 0 otherwise.  
 886

## 887 C MULTITENSORS

888 The actual data ( $z$ , hidden activations, and puzzles) passing through our layers comes in a format that  
 889 we call a “**multitensor**”, which is just a bucket of tensors of various shapes, as shown in Figure 8.  
 890 All the equivariances we use can be described in terms of how they change a multitensor.  
 891



900 Figure 8: Our neural network’s internal representations come in the form of a “multitensor”, a bucket of tensors of different shapes. One of the tensors is shaped like  
 901  $[\text{example}, \text{color}, \text{height}, \text{width}, \text{channel}]$ , an adequate shape for storing a whole ARC-AGI puzzle.  
 902

903 Most common classes of machine learning architectures operate on a single type of tensor with  
 904 constant rank. LLMs operate on rank-3 tensors of shape  $[\text{n\_batch}, \text{n\_tokens}, \text{n\_channels}]$ ,  
 905 and Convolutional Neural Networks (CNNs) operate on rank-4 tensors of shape  
 906  $[\text{n\_batch}, \text{n\_channels}, \text{height}, \text{width}]$ . Our multitensors are a set of varying-rank ten-  
 907 sors of unique type, whose dimensions are a subset of a rank-6 tensor of shape

[n\_example, n\_colors, n\_directions, height, width, n\_channels], as illustrated in Figure 8. We always keep the channel dimension, so there are at most 32 tensors in each multitensor. We also maintain several rules (see Appendix E.1) that determine whether a tensor shape is “legal” or not, which reduces the number of tensors in a multitensor to 18.

| Dimension | Description  |
|-----------|--|
| Example   | Number of examples in the ARC-AGI puzzle, including the one with held-out answer   |
| Color     | Number of unique colors in the ARC-AGI puzzle, not including black, see Appendix F.2   |
| Direction | 8  |
| Height    | Determined when preprocessing the puzzle, see Appendix F.1   |
| Width     | Determined when preprocessing the puzzle, see Appendix F.1   |
| Channel   | In the residual connections, the size is 8 if the direction dimension is included, else 16. Within layers it is layer-dependent. |

Table 3: Size conventions for multitensor dimensions.

To give an idea of how a multitensor stores data, an ARC-AGI puzzle can be represented by using the [example, color, height, width, channel] tensor, by using the channel dimension to select either the input or output grid, and the height/width dimensions for pixel location, a one hot vector in the color dimension, specifying what color that pixel is. The [example, height, channel] and [example, width, channel] tensors can similarly be used to store masks representing grid shapes for every example for every input/output grid. All those tensors are included in a single multitensor that is computed by the network just before the final linear head (described in Appendix D.8).

When we apply an operation on a multitensor, we by default assume that all non-channel dimensions are treated identically as batch dimensions by default. The operation is copied across the indices of dimensions unless specified. This ensures that we keep all our symmetries intact until we use a specific layer meant to break a specific symmetry.

A final note on the channel dimension: usually when talking about a tensor’s shape, we will not even mention the channel dimension as it is included by default.

## D LAYERS IN THE ARCHITECTURE

### D.1 DECODING LAYER

This layer’s job is to sample a multitensor  $z$  and bound its information content, before it is passed to the next layer. This layer and outputs the KL divergence between the learned  $z$  distribution and  $N(0, I)$ . Penalizing the KL prevents CompressARC from learning a distribution for  $z$  that memorizes the ARC-AGI puzzle in an uncompressed fashion, and forces CompressARC to represent the puzzle more succinctly. Specifically, it forces the network to spend more bits on the KL whenever it uses  $z$  to break a symmetry, and the larger the symmetry group broken, the more bits it spends.

This layer takes as input:

- A learned target multiscalar, called the “target capacity”.<sup>2</sup> The decoding layer will output  $z$  whose information content per tensor is close to the target capacity,<sup>3</sup>
- learned per-element means for  $z$ ,<sup>4</sup>
- learned per-element capacity adjustments for  $z$ .

<sup>2</sup>Target capacities are exponentially parameterized and rescaled by 10x to increase sensitivity to learning, initialized at a constant  $10^4$  nats per tensor, and forced to be above a minimum value of half a nat.

<sup>3</sup>The actual information content, which the layer computes later on, will be slightly different because of the per-element capacity adjustments.

<sup>4</sup>Means are initialized using normal distribution of variance  $10^{-4}$ .

We begin by normalizing the learned per-element means for  $z$ .<sup>5</sup> Then, we figure out how much Gaussian noise we must add into every tensor to make the AWGN channel capacity (Shannon, 1948) equal to the target capacity for every tensor (including per-element capacity adjustments). We apply the noise to sample  $z$ , keeping unit variance of  $z$  by rescaling.<sup>6</sup>

We compute the information content of  $z$  as the KL divergence between the distribution of this sample and  $N(0, 1)$ .

Finally, we postprocess the noisy  $z$  by scaling it by the sigmoid of the signal-to-noise ratio.<sup>7</sup> This ensures that  $z$  is kept as-is when its variance consists mostly of useful information and it is nearly zero when its variance consists mostly of noise. All this is done 4 times to make a channel dimension of 4. Then we apply a projection (with different weights per tensor in the multitensor, i.e., per-tensor projections) mapping the channel dimension up to the dimension of the residual stream.

## D.2 MULTITENSOR COMMUNICATION LAYER

This layer allows different tensors in a multitensor to interact with each other.

First, the input from the residual stream passes through per-tensor projections to a fixed size (8 for downwards communication and 16 for upwards communication). Then a message is sent to every other tensor that has at least the same dimensions for upwards communication, or at most the same dimensions for downwards communication. This message is created by either taking means along dimensions to remove them, or unsqueezing+broadcasting dimensions to add them, as in Figure 9. All the messages received by every tensor are summed together and normalization is applied. This result gets up-projected back and then added to the residual stream.

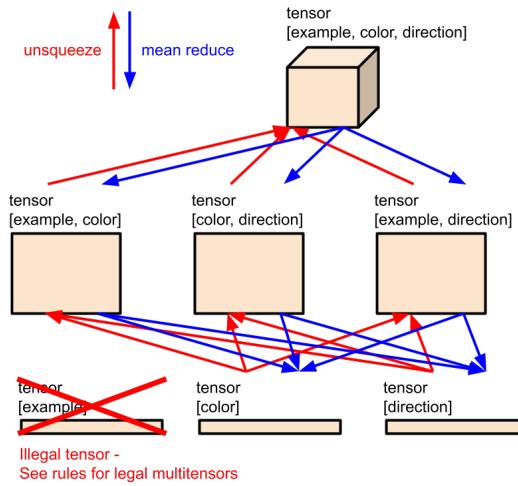


Figure 9: Multitensor communication layer. Higher rank tensors shown at the top, lower rank at the bottom. Tensors transform between ranks by mean reduction and unsqueezing dimensions.

## D.3 SOFTMAX LAYER

This layer allows the network to work with internal one-hot representations, by giving it the tools to denoise and sharpen noisy one-hot vectors. For every tensor in the input multitensor, this layer lists out all the possible subsets of dimensions of the tensor to take a softmax over,<sup>8</sup> takes the softmax

<sup>5</sup>Means and variances for normalization are computed along all non-channel dimensions.

<sup>6</sup>There are many caveats with the way this is implemented and how it works; please refer to the code (see Appendix O) for more details.

<sup>7</sup>We are careful not to let the postprocessing operation, which contains unbounded amounts of information via the signal-to-noise ratios, to leak lots of information across the layer. We only let a bit of it leak by averaging the signal-to-noise ratios across individual tensors in the multitensor.

<sup>8</sup>One exception: we always include the example dimension in the subset of dimensions.

1026 over these subsets of dimensions, and concatenates all the softmaxed results together in the channel  
 1027 dimension. The output dimension varies across different tensors in the multitensor, depending on  
 1028 their tensor rank. A pre-norm is applied, and per-tensor projections map to and from the residual  
 1029 stream. The layer has input channel dimension of 2.  
 1030

#### 1031 D.4 DIRECTIONAL CUMMAX/SHIFT LAYER 1032

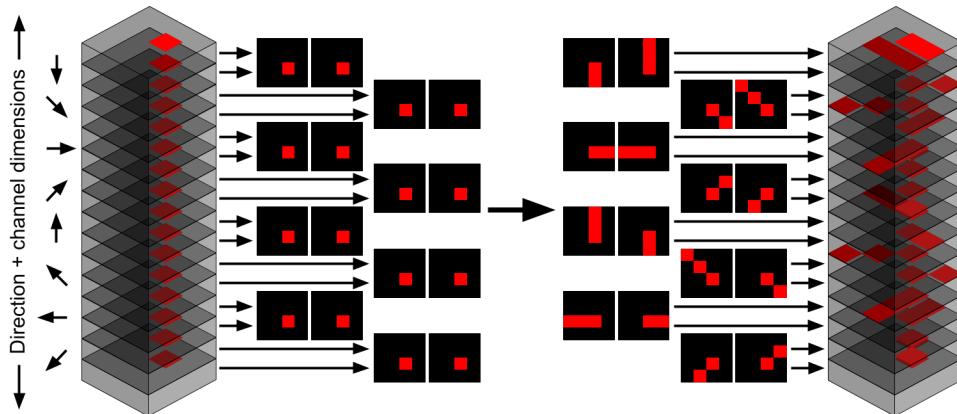
1033 The directional cummax and shift layers allow the network to perform the non-equivariant cummax  
 1034 and shift operations in an equivariant way, namely by applying the operations once per direction, and  
 1035 only letting the output be influenced by the results once the directions are aggregated back together  
 1036 (by the multitensor communication layer). These layers are the sole reason we included the direction  
 1037 dimension when defining a multitensor: to store the results of directional layers and operate on  
 1038 each individually. Of course, this means when we apply a spatial equivariance transformation, we  
 1039 must also permute the indices of the direction dimension accordingly, which can get complicated  
 1040 sometimes.  
 1041

1041 The directional cummax layer takes the eight indices of the direction dimension, treats each slice as  
 1042 corresponding to one direction (4 cardinal, 4 diagonal), performs a cumulative max in the respective  
 1043 direction for each slice, does it in the opposite direction for half the channels, and stacks the slices  
 1044 back together in the direction dimension. An illustration is in Figure 10. The slices are rescaled to  
 1045 have min  $-1$  and max  $1$  before applying the cumulative max.  
 1046

1046 The directional shift layer does the same thing, but for shifting the grid by one pixel instead of  
 1047 applying the cumulative max, and without the rescaling.  
 1048

1048 Some details:  
 1049

- 1050 • Per-tensor projections map to and from the residual stream, with pre-norm.
- 1051 • Input channel dimension is 4.
- 1052 • These layers are only applied to the [example, color, direction, height, width, channel] and  
 1053 [example, direction, height, width, channel] tensors in the input multitensor.  
 1054



1070 Figure 10: The directional cummax layer takes a directional tensor, splits it along the direction axis,  
 1071 and applies a cumulative max in a different direction for each direction slice. This operation helps  
 1072 CompressARC transport information across long distances in the puzzle grid.  
 1073

#### 1075 D.5 DIRECTIONAL COMMUNICATION LAYER 1076

1077 By default, the network is equivariant to permutations of the eight directions, but we only want  
 1078 symmetry up to rotations and flips. So, this layer provides a way to send information between two  
 1079 slices in the direction dimension, depending on the angular difference in the two directions. This  
 layer defines a separate linear map to be used for each of the 64 possible combinations of angles,  
 1080

1080 but the weights of the linear maps are minimally tied such that the directional communication layer  
 1081 is equivariant to reflections and rotations. This gets complicated really fast, since the direction  
 1082 dimension's indices also permute when equivariance transformations are applied. Every direction  
 1083 slice in a tensor accumulates its 8 messages, and adds the results together.<sup>9</sup>

1084 For this layer, there are per-tensor projections to and from the residual stream with pre-norm. The  
 1085 input channel dimension is 2.

1087 **D.6 NONLINEAR LAYER**

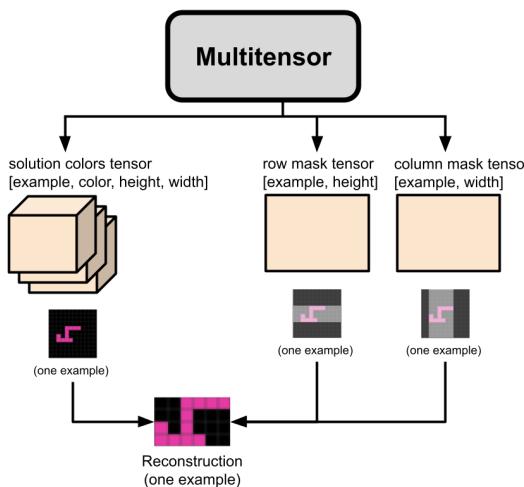
1089 We use a SiLU nonlinearity with channel dimension 16, surrounded by per-tensor projections with  
 1090 pre-norm.

1092 **D.7 NORMALIZATION LAYER**

1094 We normalize all the tensors in the multitensor, using means and variances computed across all  
 1095 dimensions except the channel dimension. Normalization as used within other layers also generally  
 1096 operates this way.

1097 **D.8 LINEAR HEADS**

1099 We must take the final multitensor, and convert it to the format of an ARC-AGI puzzle. More  
 1100 specifically, we must convert the multitensor into a distribution over ARC-AGI puzzles, so that we  
 1101 can compute the log-likelihood of the observed grids in the puzzle.



1119 Figure 11: The linear head layer takes the final multitensor of the residual stream and  
 1120 reads a [example, color, height, width, channel] tensor to be interpreted as color logits, and a  
 1121 [example, height, channel] tensor and a [example, width, channel] tensor to serve as shape masks.

1123 The colors of every pixel for every example for both input and output, have logits defined by the  
 1124 [example, color, height, width, channel] tensor, with the channel dimension linearly mapped down to  
 1125 a size of 2, representing the input and output grids.<sup>10</sup> The log-likelihood is given by the crossentropy,  
 1126 with sum reduction across all the grids.

1127 For grids of non-constant shape, the [example, height, channel] and [example, width, channel] tensors  
 1128 are used to create distributions over possible contiguous rectangular slices of each grid of colors,

1129 <sup>9</sup>We also multiply the results by coefficients depending on the angle: 1 for 0 degrees and 180 degrees, 0.2 for  
 1130 45 degrees and 135 degrees, and 0.4 for 90 degrees.

1131 <sup>10</sup>The linear map is initialized to be identical for both the input and output grid, but isn't fixed this way during  
 1132 learning. Sometimes this empirically helps with problems of inconsistent input vs output grid shapes. The  
 1133 bias on this linear map is multiplied by 100 before usage, otherwise it doesn't seem to be learned fast enough  
 empirically. This isn't done for the shape tensors described by the following paragraph though.

1134 as shown in Figure 11. Again, the channel dimension is mapped down to a size of 2 for input and  
 1135 output grids. For every grid, we have a vector of size [width] and a vector of size [height]. The log  
 1136 likelihood of every slice of the vector is taken to be the sum of the values within the slice, minus  
 1137 the values outside the slice. The log likelihoods for all the possible slices are then normalized to  
 1138 have total probability one, and the colors for every slice are given by the color logits defined in the  
 1139 previous paragraph.

1140 With the puzzle distribution now defined, we can now evaluate the log-likelihood of the observed  
 1141 target puzzle, to use as the reconstruction error.<sup>11</sup>  
 1142

## 1143 E OTHER ARCHITECTURAL DETAILS

### 1144 E.1 RULES FOR LEGAL MULTITENSORS

- 1147 1. At least one non-example dimension must be included. Examples are not special for any reason  
 not having to do with colors, directions, rows, and columns.
- 1148 2. If the width or height dimension is included, the example dimension should also be included.  
 1149 Positions are intrinsic to grids, which are indexed by the example dimension. Without a grid it  
 1150 doesn't make as much sense to talk about positions.

### 1153 E.2 WEIGHT TYING FOR REFLECTION/ROTATION SYMMETRY

1155 When applying a different linear layer to every tensor in a multitensor, we have a linear layer for  
 1156 tensors having a width but not height dimension, and another linear layer for tensors having a height  
 1157 but not width dimension. Whenever this is the case, we tie the weights together in order to preserve  
 1158 the whole network's equivariance to diagonal reflections and 90 degree rotations, which swap the  
 1159 width and height dimensions.

1160 The softmax layer is not completely symmetrized because different indices of the output correspond  
 1161 to different combinations of dimension to softmax over. Tying the weights properly would be a bit  
 1162 complicated and time consuming for the performance improvement we expect, so we did not do this.  
 1163

### 1164 E.3 TRAINING/INITIALIZATION

1166 We train for 2000 iterations using Adam, with learning rate 0.01,  $\beta_1$  of 0.5, and  $\beta_2$  of 0.9. **Weights**  
 1167 **are essentially all initialized with Xavier normal initialization.**

## 1169 F PREPROCESSING

### 1171 F.1 OUTPUT SHAPE DETERMINATION

1173 The raw data consists of grids of various shapes, while the neural network operates on grids of  
 1174 constant shape. Most of the preprocessing that we do is aimed towards this shape inconsistency  
 1175 problem.

1176 Before doing any training, we determine whether the given ARC-AGI puzzle follows three possible  
 1177 shape consistency rules:

- 1179 1. The outputs in a given ARC-AGI puzzle are always the same shape as corresponding inputs.
- 1180 2. All the inputs in the given ARC-AGI puzzle are the same shape.
- 1181 3. All the outputs in the given ARC-AGI puzzle are the same shape.

1183 <sup>11</sup>There are multiple slices of the same shape that result in the correct puzzle to be decoded. We sum together  
 1184 the probabilities of getting any of the slices by applying a logsumexp to the log probabilities. But, we found  
 1185 empirically that training prematurely collapses onto one particular slice. So, we pre-multiply and post-divide  
 1186 the log probabilities by a coefficient when applying the logsumexp. The coefficient starts at 0.1 and increases  
 1187 exponentially to 1 over the first 100 iterations of training. We also pre-multiply the masks by the square of this  
 1188 coefficient as well, to ensure they are not able to strongly concentrate on one slice too early in training.

1188 Based on rules 1 and 3, we try to predict the shape of held-out outputs, prioritizing rule 1 over rule  
 1189 3. If either rule holds, we force the postprocessing step to only consider the predicted shape by  
 1190 overwriting the masks produced by the linear head layer. If neither rule holds, we make a temporary  
 1191 prediction of the largest width and height out of the grids in the given ARC-AGI puzzle, and we allow  
 1192 the masks to predict shapes that are smaller than that.

1193 The largest width and height that is given or predicted, are used as the size of the multitensor’s width  
 1194 and height dimensions.

1195 The predicted shapes are also used as masks when performing the multitensor communication,  
 1196 directional communication and directional cummax/shift layers. We did not apply masks for the  
 1197 other layers because of time constraints and because we do not believe it will provide for much of a  
 1198 performance improvement.<sup>12</sup>

## 1200 F.2 NUMBER OF COLORS

1202 We notice that in almost all ARC-AGI puzzles, colors that are not present in the puzzle are not present  
 1203 in the true answers. Hence, any colors that do not appear in the puzzle are not given an index in the  
 1204 color dimension of the multitensor.

1205 In addition, black is treated as a special color that is never included in the multitensor, since it  
 1206 normally represents the background in many puzzles. When performing color classification, a tensor  
 1207 of zeros is appended to the color dimension after applying the linear head, to represent logits for the  
 1208 black color.

## 1210 G POSTPROCESSING

1212 Since the generated answer grid is stochastic from randomness in  $z$ , we save the answer grids  
 1213 throughout training, and roughly speaking, we choose the most frequently occurring one as our  
 1214 denoised final prediction. This is complicated by the variable shape grids present in some puzzles.

1216 Generally, when we sample answers from the network by taking the logits of the  
 1217 [example, color, height, width, channel] tensor and argmaxing over the color dimension, we find that  
 1218 the grids are noisy and will often have the wrong colors for several random pixels. We developed  
 1219 several methods for removing this noise:

- 1221 1. Find the most commonly sampled answer.
- 1222 2. Construct an exponential moving average of the output color logits before taking the softmax to  
 1223 produce probabilities. Also construct an exponential moving average of the masks.
- 1224 3. Construct an exponential moving average of the output color probabilities after taking the softmax.  
 1225 Also construct an exponential moving average of the masks.

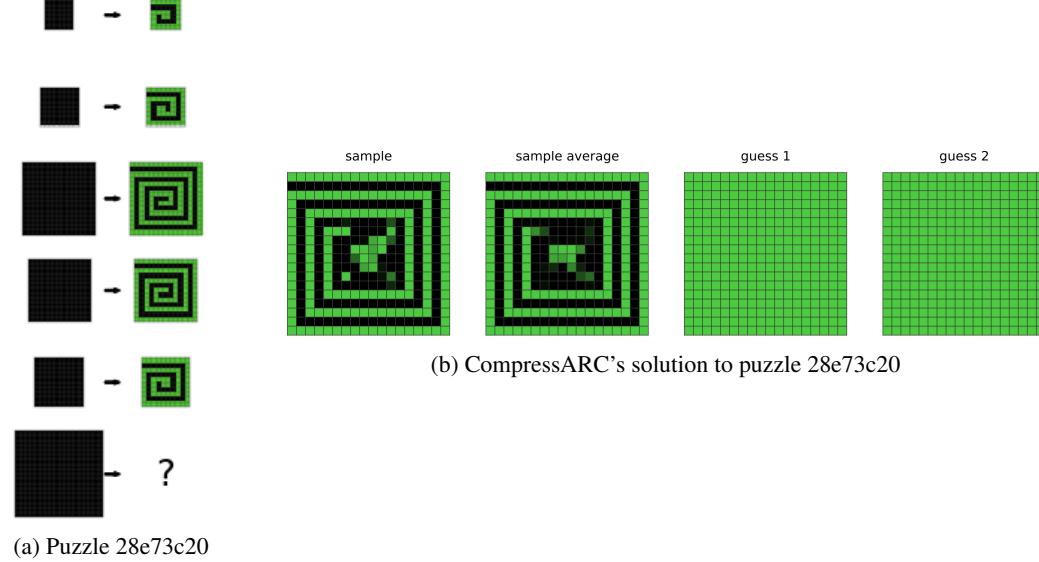
1227 When applying these techniques, we always take the slice of highest probability given the mask, and  
 1228 then we take the colors of highest probability afterwards.

1229 We explored several different rules for when to select which method, and arrived at a combination of  
 1230 1 and 2 with a few modifications:

- 1232 • At every iteration, count up the sampled answer, as well as the exponential moving average answer  
 1233 (decay = 0.97).
- 1234 • If before 150 iterations of training, then downweight the answer by a factor of  $e^{-10}$ . (Effectively,  
 1235 don’t count the answer.)
- 1236 • If the answer is from the exponential moving average as opposed to the sample, then downweight  
 1237 the answer by a factor of  $e^{-4}$ .
- 1239 • Downweight the answer by a factor of  $e^{-10 * \text{uncertainty}}$ , where uncertainty is the average (across  
 1240 pixels) negative log probability assigned to the top color of every pixel.

1241 <sup>12</sup>The two masks for the input and output are combined together to make one mask for use in these operations,  
 1242 since the channel dimension in these operations don’t necessarily correspond to the input and output grids.

1242  
 1243 **H EMPIRICALLY OBSERVED ABILITIES AND DISABILITIES OF**  
 1244 **COMPRESSARC**



1265 Figure 12: Puzzle 28e73c20, and CompressARC’s solution to it.  
 1266

1267 A short list of abilities that **can** be performed by CompressARC includes:

1268

- 1269 • Assigning individual colors to individual procedures (see puzzle 0ca9ddb6)
- 1270 • Infilling (see puzzle 0dfd9992)
- 1271 • Cropping (see puzzle 1c786137)
- 1272 • Connecting dots with lines, including 45 degree diagonal lines (see puzzle 1f876c06)
- 1273 • Same color detection (see puzzle 1f876c06)
- 1274 • Identifying pixel adjacencies (see puzzle 42a50994)
- 1275 • Assigning individual colors to individual examples (see puzzle 3bd67248)
- 1276 • Identifying parts of a shape (see puzzle 025d127b)
- 1277 • Translation by short distances (see puzzle 025d127b)

1282 We believe these abilities to be individually endowed by select layers in the architecture, which we  
 1283 designed specifically for the purpose of conferring those abilities to CompressARC.  
 1284

1285 A short list of abilities that **cannot** be performed by CompressARC includes:

1286

- 1287 • Assigning two colors to each other (see puzzle 0d3d703e)
- 1288 • Repeating an operation in series many times (see puzzle 0a938d79)
- 1289 • Counting/numbers (see puzzle ce9e57f2)
- 1290 • Translation, rotation, reflections, rescaling, image duplication (see puzzles 0e206a2e, 5ad4f10b,  
 1291 and 2bcee788)
- 1292 • Detecting topological properties such as connectivity (see puzzle 7b6016b9)
- 1293 • Planning, simulating the behavior of an agent (see puzzle 2dd70a9a)
- 1294 • Long range extensions of patterns (see puzzle 28e73c20 above)

1296 **I BASELINES**  
12971298 The U-Net baseline in Table 1 was created to observe the performance of a more standard approach  
1299 when subject to the same constraints of CompressARC, namely the avoidance of any training before  
1300 inference time, and the sole use of the test puzzle as training data during inference time.  
13011302 The training algorithm for the baseline consists of feeding each input grid into the U-Net and using  
1303 the U-Net output to classify the pixel color of the output grid. Puzzles where the input grid and output  
1304 grid did not match shape were skipped and assumed to receive a score of zero. The most common  
1305 two output grids occurring in the second half of the 10000 steps of training were used as the two  
1306 solution guesses.  
13071308 We did not change the width or height of the grids in order to fit the ARC-AGI grids into the U-Net.  
1309 The U-Net’s BatchNorm was replaced with a GroupNorm, and the middle pooling/upsampling layers  
1310 were skipped if the activation grids were too small to be pooled anymore.  
13111312 We experimented with applying a random augmentation transformation to the input grids and reversing  
1313 the transformation on the output before computing the loss and/or ensembling predictions, but we  
1314 discarded this idea due to worse performance on the training set (2.5% without augmentations, 1%  
1315 with augmentations).  
13161317 **J PUZZLE SOLVE ACCURACY TABLES**  
13181319 See Tables 4 and 5 for numerically reported puzzle solve accuracies on the whole dataset.  
13201321 Table 4: CompressARC’s puzzle solve accuracy on the training set as a function of the number of  
1322 steps of inference time learning it is given, for various numbers of allowed **guesses** (pass@n). The  
1323 official benchmark is reported with 2 allowed **guesses**, which is why we report 20% on the evaluation  
1324 set. Total training set solve time is reported for an NVIDIA RTX 4070 GPU by solving one puzzle at  
1325 a time in a sequence.  
1326

| Training Iteration | Time  | Pass@1 | Pass@2 | Pass@5 | Pass@10 | Pass@100 | Pass@1000 |
|--------------------|-------|--------|--------|--------|---------|----------|-----------|
| 100                | 6 h   | 1.00%  | 2.25%  | 3.50%  | 4.75%   | 6.75%    | 6.75%     |
| 200                | 13 h  | 11.50% | 14.25% | 16.50% | 18.25%  | 23.25%   | 23.50%    |
| 300                | 19 h  | 18.50% | 21.25% | 23.50% | 26.75%  | 31.50%   | 32.50%    |
| 400                | 26 h  | 21.00% | 25.00% | 28.75% | 31.00%  | 36.00%   | 37.50%    |
| 500                | 32 h  | 23.00% | 27.50% | 31.50% | 33.50%  | 39.25%   | 40.75%    |
| 750                | 49 h  | 28.00% | 30.50% | 34.00% | 36.25%  | 42.75%   | 44.50%    |
| 1000               | 65 h  | 28.00% | 31.75% | 35.50% | 37.75%  | 43.75%   | 46.50%    |
| 1250               | 81 h  | 29.00% | 32.25% | 37.00% | 39.25%  | 45.50%   | 49.25%    |
| 1500               | 97 h  | 29.50% | 33.00% | 38.25% | 40.75%  | 46.75%   | 51.75%    |
| 2000               | 130 h | 30.25% | 34.75% | 38.25% | 41.50%  | 48.50%   | 52.75%    |

1337 **K HOW TO IMPROVE OUR WORK**  
13381339 At the time of release of CompressARC, there were several ideas which we thought of trying or  
1340 attempted at some point, but didn’t manage to get working for one reason or another. Some ideas we  
1341 still believe in, but didn’t use, are listed below.  
13421343 **K.1 JOINT COMPRESSION VIA WEIGHT SHARING BETWEEN PUZZLES**  
13441345 Template Algorithm 1 includes a hard-coded value of  $\theta$  for every single puzzle. We might be able to  
1346 further shorten the template program length by sharing a single  $\theta$  between all the puzzles, knowing  
1347 that Occam’s razor says a shorter program corresponds to more correct puzzle solutions. Algorithm 2  
1348 would have to be changed accordingly.  
1349

To implement this, we would most likely explore strategies like:

1350 Table 5: CompressARC’s puzzle solve accuracy on the evaluation set, reported the same way as in  
 1351 Table 4.

| 1353 <b>Training Iteration</b> | 1354 <b>Time</b> | 1355 <b>Pass@1</b> | 1356 <b>Pass@2</b> | 1357 <b>Pass@5</b> | 1358 <b>Pass@10</b> | 1359 <b>Pass@100</b> | 1360 <b>Pass@1000</b> |
|--------------------------------|------------------|--------------------|--------------------|--------------------|---------------------|----------------------|-----------------------|
| 1361 100                       | 1362 7 h         | 1363 0.75%         | 1364 1.25%         | 1365 2.25%         | 1366 2.50%          | 1367 3.00%           | 1368 3.00%            |
| 1369 200                       | 1370 14 h        | 1371 5.00%         | 1372 6.00%         | 1373 7.00%         | 1374 7.75%          | 1375 12.00%          | 1376 12.25%           |
| 1377 300                       | 1378 21 h        | 1379 10.00%        | 1380 10.75%        | 1381 12.25%        | 1382 13.25%         | 1383 15.50%          | 1384 16.25%           |
| 1385 400                       | 1386 28 h        | 1387 11.75%        | 1388 13.75%        | 1389 16.00%        | 1390 17.00%         | 1391 19.75%          | 1392 20.00%           |
| 1393 500                       | 1394 34 h        | 1395 13.50%        | 1396 15.00%        | 1397 17.75%        | 1398 19.25%         | 1399 20.50%          | 1400 21.50%           |
| 1401 750                       | 1402 52 h        | 1403 15.50%        | 1404 17.75%        | 1405 19.75%        | 1406 21.50%         | 1407 22.75%          | 1408 25.50%           |
| 1409 1000                      | 1410 69 h        | 1411 16.75%        | 1412 19.25%        | 1413 21.75%        | 1414 23.00%         | 1415 26.00%          | 1416 28.75%           |
| 1417 1250                      | 1418 86 h        | 1419 17.00%        | 1420 20.75%        | 1421 23.00%        | 1422 24.50%         | 1423 28.25%          | 1424 30.75%           |
| 1425 1500                      | 1426 103 h       | 1427 18.25%        | 1428 21.50%        | 1429 24.25%        | 1430 25.50%         | 1431 29.50%          | 1432 31.75%           |
| 1433 2000                      | 1434 138 h       | 1435 18.50%        | 1436 20.00%        | 1437 24.25%        | 1438 26.00%         | 1439 31.25%          | 1440 33.75%           |

- 1365 • Using the same network weights for all puzzles, and training for puzzles in parallel. Each puzzle  
 1366 gets assigned some perturbation to the weights, that is constrained in some way, e.g., LORA (Hu  
 1367 et al., 2021).
- 1368 • Learning a "puzzle embedding" for every puzzle that is a high dimensional vector (more than 16  
 1369 dim, less than 256 dim), and learning a linear mapping from puzzle embeddings to weights for our  
 1370 network. This mapping serves as a basic hypernetwork, i.e., a neural network that outputs weights  
 1371 for another neural network (Chauhan et al., 2024).

1373 Unfortunately, testing this would require changing CompressARC (Algorithm 3) to run all puzzles  
 1374 in parallel rather than one at a time in series. This would slow down the research iteration process,  
 1375 which is why we did not explore this option.

## 1378 K.2 CONVOLUTION-LIKE LAYERS FOR SHAPE COPYING TASKS

1379 This improvement is more ARC-AGI-specific and may have less to do with AGI in our view. Many  
 1380 ARC-AGI-1 puzzles can be seen to involve copying shapes from one place to another, and our  
 1381 network has no inductive biases for such an operation. An operation which is capable of copying  
 1382 shapes onto multiple locations is the convolution. With one grid storing the shape and another with  
 1383 pixels activated at locations to copy to, convolving the two grids will produce another grid with the  
 1384 shape copied to the designated locations.

1385 There are several issues with introducing a convolutional operation for the network to use. Ideally,  
 1386 we would read two grids via projection from the residual stream, convolve them, and write it back in  
 1387 via another projection, with norms in the right places and such. Ignoring the fact that the grid size  
 1388 changes during convolution (can be solved with two parallel networks using different grid sizes), the  
 1389 bigger problem is that convolutions tend to amplify noise in the grids much more than the sparse  
 1390 signals, so their inductive bias is not good for shape copying. We can try to apply a softmax to one  
 1391 or both of the grids to reduce the noise (and to draw an interesting connection to attention), but we  
 1392 didn’t find any success.

1393 The last idea that we were tried before discarding the idea was to modify the functional form of the  
 1394 convolution:

$$1395 (f * g)(x) = \sum_y f(x - y)g(y)$$

1396 to a tropical convolution (Fan et al., 2021), which we found to work well on toy puzzles, but not well  
 1397 enough for ARC-AGI-1 training puzzles (which is why we discarded this idea):

$$1398 (f * g)(x) = \max_y f(x - y) + g(y)$$

1404 Convolutions, when repeated with some grids flipped by 180 degrees, tend to create high activations  
 1405 at the center pixel, so sometimes it is important to zero out the center pixel to preserve the signal.  
 1406

### 1407 K.3 KL FLOOR FOR POSTERIOR COLLAPSE

1409 We noticed during testing that crucial posterior tensors whose KL fell to zero during learning would  
 1410 never make a recovery and play their role in the encoding, just as in the phenomenon of mode collapse  
 1411 in variational autoencoders (van den Oord et al., 2018). We believe that the KL divergence may upper  
 1412 bound the information content of the gradient training signal for parts of the network that process the  
 1413 encoded information. Thus, when a tensor in  $z$  falls to zero KL, the network stops learning to use its  
 1414 encoded information, and the KL is no longer incentivized to recover. If we artificially hold the KL  
 1415 above zero for an extended period of training, then the network may learn to make use of the tensor's  
 1416 information, incentivizing the KL to stay above zero when released again.  
 1417

1418 We implemented a mechanism to keep the KL above a minimum threshold so that the network always  
 1419 learns to use that information, but we do not believe the network learns fast enough for this to be  
 1420 useful, as we have never seen a tensor recover before. Therefore, it might be useful to explore  
 1421 different ways to schedule this KL floor to start high and decay to zero, to allow learning when the  
 1422 KL is forced to be high, and to leave the KL unaffected later on in learning. This might cause training  
 1423 results to be more consistent across runs.  
 1424

### 1425 K.4 REGULARIZATION

1426 In template Algorithm 1, we do not code-golf  $\theta$  to reduce the number of bits it takes up. If we were  
 1427 to code-golf  $\theta$  as well, this would produce an extra KL term to the loss in CompressARC (Algorithm  
 1428 3), and the KL term would simplify to an L2 regularization on  $\theta$  under certain reasonable limits. It  
 1429 is somewhat reckless for us to neglect code-golfing  $\theta$  in our work due to the sheer number of bits  $\theta$   
 1430 contributes, and making this change may improve our results.  
 1431

## 1432 L ADDITIONAL DETAILS ABOUT THE ARC-AGI BENCHMARK

1433 **Hidden Rules:** For every puzzle, there is a hidden rule that maps each input grid to each output  
 1434 grid. There are 400 training puzzles and they are easier to solve than the 400 evaluation puzzles. The  
 1435 training set is intended to help teach your system the following general themes which underlie the  
 1436 hidden rules in the evaluation set:  
 1437

- 1438 • **Objectness:** Objects persist and cannot appear or disappear without reason. Objects can interact or  
 1439 not depending on the circumstances.
- 1440 • **Goal-directedness:** Objects can be animate or inanimate. Some objects are “agents” - they have  
 1441 intentions and they pursue goals.
- 1442 • **Numbers & counting:** Objects can be counted or sorted by their shape, appearance, or movement  
 1443 using basic mathematics like addition, subtraction, and comparison.
- 1444 • **Basic geometry & topology:** Objects can be shapes like rectangles, triangles, and circles which  
 1445 can be mirrored, rotated, translated, deformed, combined, repeated, etc. Differences in distances  
 1446 can be detected.

1447 The puzzles are designed so that **humans can reasonably find the answer, but machines should**  
 1448 **have more difficulty**. The average human can solve 76.2% of the training set, and a human expert  
 1449 can solve 98.5% (LeGris et al., 2024).  
 1450

1451 **Scoring:** You are given some number of examples of input-to-output mappings, and you get **two**  
 1452 **guesses** to guess the output grid(s) for a given input grid, without being told the hidden rule. **One**  
 1453 **guess consists of guessing the width and height of the output grid(s), as well as all the pixel colors**  
 1454 **within, and if each of these is correct, then that guess succeeds. If either guess succeeds**, then you  
 1455 score 1 for that puzzle, else you score 0. Some puzzles have more than one input/output pair that you  
 1456 have to guess, in which case the score for that puzzle may be in between.  
 1457

**Scoring Environment:** The competitions launched by the ARC Prize Foundation have been restricted  
 1458 to 12 hours of compute per solution submission, in a constrained environment with no internet access.  
 1459

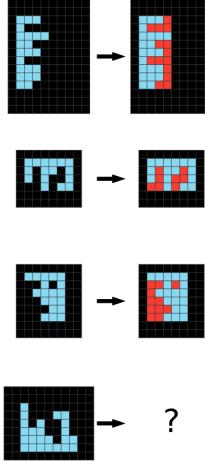
1458 This is where a hidden semi-private evaluation set is used to score solutions. The scores we report are  
 1459 on the public evaluation set, which is of the same difficulty as the semi-private evaluation set, which  
 1460 we had no access to when we performed this work.  
 1461

## 1462 M ADDITIONAL CASE STUDIES

1463 Below, we show two additional puzzles and a dissection of CompressARC’s solution to them.  
 1464

### 1465 M.1 CASE STUDY: BOUNDING BOX

1466 Puzzle 6d75e8bb is part of the training split, see Figure 13.



1467 Figure 13: Bounding Box: Puzzle 6d75e8bb from the training split.  
 1468

#### 1469 M.1.1 WATCHING THE NETWORK LEARN: BOUNDING BOX

1470 **Human Solution:** We first realize that the input is red and black, and the output is also red and black,  
 1471 but some of the black pixels are replaced by light blue pixels. We see that the red shape remains  
 1472 unaffected. We notice that the light blue box surrounds the red shape, and finally that it is the smallest  
 1473 possible surrounding box that contains the red shape. At this point, we copy the input over to the  
 1474 answer grid, then we figure out the horizontal and vertical extent of the red shape, and color all of the  
 1475 non-red pixels within that extent as light blue.  
 1476

1477 **CompressARC Solution: See Table 6**

#### 1478 M.1.2 SOLUTION ANALYSIS: BOUNDING BOX

1479 Figure 14 shows the amount of contained information in every tensor within  $z$ .  
 1480

1481 All the tensors in  $z$  fall to zero information content during training, except for three tensors. From 600-  
 1482 1000 steps, we see the (example, height, width, channel) tensor suffer a massive drop in information  
 1483 content, with no change in the outputted answer. We believe it was being used to identify the light blue  
 1484 pixels in the input, but this information then got memorized by the nonlinear portions of the network,  
 1485 using the (example, height, channel) and (example, width, channel) as positional encodings.  
 1486

1487 Figure 15 shows the average output of the decoding layer for these tensors to see what information is  
 1488 stored there.  
 1489

## 1490 M.2 CASE STUDY: CENTER CROSS

1491 Puzzle 41e4d17e is part of the training split, see Figure 16a.  
 1492

1512

1513

1514

1515

1516

1517

Table 6: CompressARC learning the solution for Bounding Box, over time.

1518

1519

1520

1521

1522

1523

1524

1525

1526

1527

1528

1529

1530

1531

1532

1533

1534

1535

1536

1537

1538

1539

1540

1541

1542

1543

1544

1545

1546

1547

1548

1549

1550

1551

1552

1553

1554

1555

1556

1557

1558

1559

1560

1561

1562

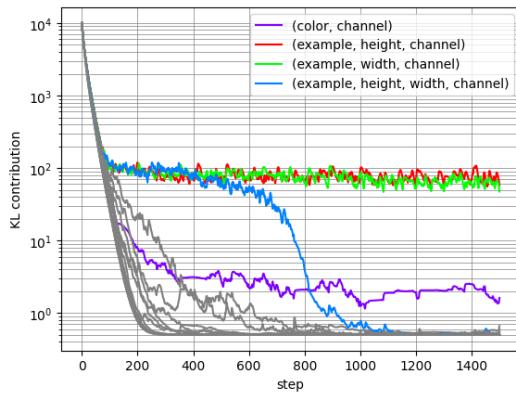
1563

1564

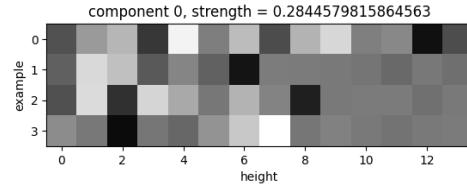
1565

| Learning steps | What is CompressARC doing?   | Sampled solution guess |                |         |         |
|----------------|--|------------------------|----------------|---------|---------|
|                |  | sample                 | sample average | guess 1 | guess 2 |
| 50             | The average of sampled outputs shows that light blue pixels in the input are generally preserved in the output. However, black pixels in the input are haphazardly and randomly colored light blue and red. CompressARC does not seem to know that the colored input/output pixels lie within some kind of bounding box, or that the bounding box is the same for the input and output grids.  |                        |                |         |         |
| 100            | The average of sampled outputs shows red pixels confined to an imaginary rectangle surrounding the light blue pixels. CompressARC seems to have perceived that other examples use a common bounding box for the input and output pixels, but is not completely sure about where the boundary lies and what colors go inside the box in the output. Nevertheless, guess 2 (the second most frequently sampled output) shows that the correct answer is already being sampled quite often now. |                        |                |         |         |
| 150            | The average of sampled outputs shows almost all of the pixels in the imaginary bounding box to be colored red. CompressARC has figured out the answer, and further training only refines the answer.   |                        |                |         |         |

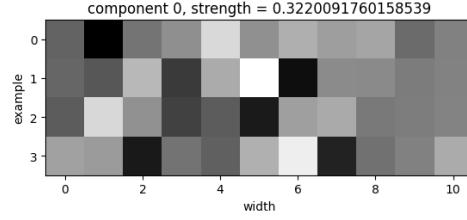
1566  
1567  
1568  
1569  
1570  
1571  
1572  
1573  
1574  
1575  
1576  
1577  
1578



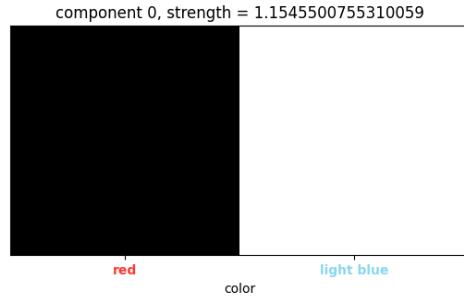
1579  
1580 Figure 14: Breaking down the KL loss during training into contributions from each individual shaped  
1581 tensor in the multitensor  $z$ .  
1582  
1583



1590  
1591 (a) **(example, height, channel) tensor.** The first  
1592 principal component is 771 times stronger than  
1593 the second principal component. **A brighter pixel**  
1594 **indicates a row with more light blue pixels.** It  
1595 is unclear how CompressARC knows where the  
1596 borders of the bounding box are.  
1597  
1598  
1599  
1600  
1601  
1602  
1603  
1604  
1605



1606  
1607 (b) **(example, width, channel) tensor.** The first  
1608 principal component is 550 times stronger than  
1609 the second principal component. **A darker pixel**  
1610 **indicates a column with more light blue pixels.**  
1611 It is unclear how CompressARC knows where the  
1612 borders of the bounding box are.  
1613  
1614  
1615  
1616  
1617  
1618  
1619



1616  
1617 (c) **(color, channel) tensor.** This tensor serves to  
1618 distinguish the roles of the two colors apart.  
1619

1620  
1621 Figure 15: Breaking down the loss components during training tells us where and how CompressARC  
1622 prefers to store information relevant to solving a puzzle.  
1623  
1624  
1625  
1626  
1627  
1628  
1629  
1630  
1631  
1632  
1633  
1634  
1635  
1636  
1637  
1638  
1639  
1640  
1641  
1642  
1643  
1644  
1645  
1646  
1647  
1648  
1649  
1650  
1651  
1652  
1653  
1654  
1655  
1656  
1657  
1658  
1659  
1660  
1661  
1662  
1663  
1664  
1665  
1666  
1667  
1668  
1669  
1670  
1671  
1672  
1673  
1674  
1675  
1676  
1677  
1678  
1679  
1680  
1681  
1682  
1683  
1684  
1685  
1686  
1687  
1688  
1689  
1690  
1691  
1692  
1693  
1694  
1695  
1696  
1697  
1698  
1699  
1700  
1701  
1702  
1703  
1704  
1705  
1706  
1707  
1708  
1709  
1710  
1711  
1712  
1713  
1714  
1715  
1716  
1717  
1718  
1719  
1720  
1721  
1722  
1723  
1724  
1725  
1726  
1727  
1728  
1729  
1730  
1731  
1732  
1733  
1734  
1735  
1736  
1737  
1738  
1739  
1740  
1741  
1742  
1743  
1744  
1745  
1746  
1747  
1748  
1749  
1750  
1751  
1752  
1753  
1754  
1755  
1756  
1757  
1758  
1759  
1760  
1761  
1762  
1763  
1764  
1765  
1766  
1767  
1768  
1769  
1770  
1771  
1772  
1773  
1774  
1775  
1776  
1777  
1778  
1779  
1780  
1781  
1782  
1783  
1784  
1785  
1786  
1787  
1788  
1789  
1790  
1791  
1792  
1793  
1794  
1795  
1796  
1797  
1798  
1799  
1800  
1801  
1802  
1803  
1804  
1805  
1806  
1807  
1808  
1809  
1810  
1811  
1812  
1813  
1814  
1815  
1816  
1817  
1818  
1819  
1820  
1821  
1822  
1823  
1824  
1825  
1826  
1827  
1828  
1829  
1830  
1831  
1832  
1833  
1834  
1835  
1836  
1837  
1838  
1839  
1840  
1841  
1842  
1843  
1844  
1845  
1846  
1847  
1848  
1849  
1850  
1851  
1852  
1853  
1854  
1855  
1856  
1857  
1858  
1859  
1860  
1861  
1862  
1863  
1864  
1865  
1866  
1867  
1868  
1869  
1870  
1871  
1872  
1873  
1874  
1875  
1876  
1877  
1878  
1879  
1880  
1881  
1882  
1883  
1884  
1885  
1886  
1887  
1888  
1889  
1890  
1891  
1892  
1893  
1894  
1895  
1896  
1897  
1898  
1899  
1900  
1901  
1902  
1903  
1904  
1905  
1906  
1907  
1908  
1909  
1910  
1911  
1912  
1913  
1914  
1915  
1916  
1917  
1918  
1919  
1920  
1921  
1922  
1923  
1924  
1925  
1926  
1927  
1928  
1929  
1930  
1931  
1932  
1933  
1934  
1935  
1936  
1937  
1938  
1939  
1940  
1941  
1942  
1943  
1944  
1945  
1946  
1947  
1948  
1949  
1950  
1951  
1952  
1953  
1954  
1955  
1956  
1957  
1958  
1959  
1960  
1961  
1962  
1963  
1964  
1965  
1966  
1967  
1968  
1969  
1970  
1971  
1972  
1973  
1974  
1975  
1976  
1977  
1978  
1979  
1980  
1981  
1982  
1983  
1984  
1985  
1986  
1987  
1988  
1989  
1990  
1991  
1992  
1993  
1994  
1995  
1996  
1997  
1998  
1999  
2000  
2001  
2002  
2003  
2004  
2005  
2006  
2007  
2008  
2009  
2010  
2011  
2012  
2013  
2014  
2015  
2016  
2017  
2018  
2019  
2020  
2021  
2022  
2023  
2024  
2025  
2026  
2027  
2028  
2029  
2030  
2031  
2032  
2033  
2034  
2035  
2036  
2037  
2038  
2039  
2040  
2041  
2042  
2043  
2044  
2045  
2046  
2047  
2048  
2049  
2050  
2051  
2052  
2053  
2054  
2055  
2056  
2057  
2058  
2059  
2060  
2061  
2062  
2063  
2064  
2065  
2066  
2067  
2068  
2069  
2070  
2071  
2072  
2073  
2074  
2075  
2076  
2077  
2078  
2079  
2080  
2081  
2082  
2083  
2084  
2085  
2086  
2087  
2088  
2089  
2090  
2091  
2092  
2093  
2094  
2095  
2096  
2097  
2098  
2099  
2100  
2101  
2102  
2103  
2104  
2105  
2106  
2107  
2108  
2109  
2110  
2111  
2112  
2113  
2114  
2115  
2116  
2117  
2118  
2119  
2120  
2121  
2122  
2123  
2124  
2125  
2126  
2127  
2128  
2129  
2130  
2131  
2132  
2133  
2134  
2135  
2136  
2137  
2138  
2139  
2140  
2141  
2142  
2143  
2144  
2145  
2146  
2147  
2148  
2149  
2150  
2151  
2152  
2153  
2154  
2155  
2156  
2157  
2158  
2159  
2160  
2161  
2162  
2163  
2164  
2165  
2166  
2167  
2168  
2169  
2170  
2171  
2172  
2173  
2174  
2175  
2176  
2177  
2178  
2179  
2180  
2181  
2182  
2183  
2184  
2185  
2186  
2187  
2188  
2189  
2190  
2191  
2192  
2193  
2194  
2195  
2196  
2197  
2198  
2199  
2200  
2201  
2202  
2203  
2204  
2205  
2206  
2207  
2208  
2209  
2210  
2211  
2212  
2213  
2214  
2215  
2216  
2217  
2218  
2219  
2220  
2221  
2222  
2223  
2224  
2225  
2226  
2227  
2228  
2229  
2230  
2231  
2232  
2233  
2234  
2235  
2236  
2237  
2238  
2239  
2240  
2241  
2242  
2243  
2244  
2245  
2246  
2247  
2248  
2249  
2250  
2251  
2252  
2253  
2254  
2255  
2256  
2257  
2258  
2259  
2260  
2261  
2262  
2263  
2264  
2265  
2266  
2267  
2268  
2269  
2270  
2271  
2272  
2273  
2274  
2275  
2276  
2277  
2278  
2279  
2280  
2281  
2282  
2283  
2284  
2285  
2286  
2287  
2288  
2289  
2290  
2291  
2292  
2293  
2294  
2295  
2296  
2297  
2298  
2299  
2300  
2301  
2302  
2303  
2304  
2305  
2306  
2307  
2308  
2309  
2310  
2311  
2312  
2313  
2314  
2315  
2316  
2317  
2318  
2319  
2320  
2321  
2322  
2323  
2324  
2325  
2326  
2327  
2328  
2329  
2330  
2331  
2332  
2333  
2334  
2335  
2336  
2337  
2338  
2339  
2340  
2341  
2342  
2343  
2344  
2345  
2346  
2347  
2348  
2349  
2350  
2351  
2352  
2353  
2354  
2355  
2356  
2357  
2358  
2359  
2360  
2361  
2362  
2363  
2364  
2365  
2366  
2367  
2368  
2369  
2370  
2371  
2372  
2373  
2374  
2375  
2376  
2377  
2378  
2379  
2380  
2381  
2382  
2383  
2384  
2385  
2386  
2387  
2388  
2389  
2390  
2391  
2392  
2393  
2394  
2395  
2396  
2397  
2398  
2399  
2400  
2401  
2402  
2403  
2404  
2405  
2406  
2407  
2408  
2409  
2410  
2411  
2412  
2413  
2414  
2415  
2416  
2417  
2418  
2419  
2420  
2421  
2422  
2423  
2424  
2425  
2426  
2427  
2428  
2429  
2430  
2431  
2432  
2433  
2434  
2435  
2436  
2437  
2438  
2439  
2440  
2441  
2442  
2443  
2444  
2445  
2446  
2447  
2448  
2449  
2450  
2451  
2452  
2453  
2454  
2455  
2456  
2457  
2458  
2459  
2460  
2461  
2462  
2463  
2464  
2465  
2466  
2467  
2468  
2469  
2470  
2471  
2472  
2473  
2474  
2475  
2476  
2477  
2478  
2479  
2480  
2481  
2482  
2483  
2484  
2485  
2486  
2487  
2488  
2489  
2490  
2491  
2492  
2493  
2494  
2495  
2496  
2497  
2498  
2499  
2500  
2501  
2502  
2503  
2504  
2505  
2506  
2507  
2508  
2509  
2510  
2511  
2512  
2513  
2514  
2515  
2516  
2517  
2518  
2519  
2520  
2521  
2522  
2523  
2524  
2525  
2526  
2527  
2528  
2529  
2530  
2531  
2532  
2533  
2534  
2535  
2536  
2537  
2538  
2539  
2540  
2541  
2542  
2543  
2544  
2545  
2546  
2547  
2548  
2549  
2550  
2551  
2552  
2553  
2554  
2555  
2556  
2557  
2558  
2559  
2560  
2561  
2562  
2563  
2564  
2565  
2566  
2567  
2568  
2569  
2570  
2571  
2572  
2573  
2574  
2575  
2576  
2577  
2578  
2579  
2580  
2581  
2582  
2583  
2584  
2585  
2586  
2587  
2588  
2589  
2590  
2591  
2592  
2593  
2594  
2595  
2596  
2597  
2598  
2599  
2600  
2601  
2602  
2603  
2604  
2605  
2606  
2607  
2608  
2609  
2610  
2611  
2612  
2613  
2614  
2615  
2616  
2617  
2618  
2619  
2620  
2621  
2622  
2623  
2624  
2625  
2626  
2627  
2628  
2629  
2630  
2631  
2632  
2633  
2634  
2635  
2636  
2637  
2638  
2639  
2640  
2641  
2642  
2643  
2644  
2645  
2646  
2647  
2648  
2649  
2650  
2651  
2652  
2653  
2654  
2655  
2656  
2657  
2658  
2659  
2660  
2661  
2662  
2663  
2664  
2665  
2666  
2667  
2668  
2669  
2670  
2671  
2672  
2673  
2674  
2675  
2676  
2677  
2678  
2679  
2680  
2681  
2682  
2683  
2684  
2685  
2686  
2687  
2688  
2689  
2690  
2691  
2692  
2693  
2694  
2695  
2696  
2697  
2698  
2699  
2700  
2701  
2702  
2703  
2704  
2705  
2706  
2707  
2708  
2709  
2710  
2711  
2712  
2713  
2714  
2715  
2716  
2717  
2718  
2719  
2720  
2721  
2722  
2723  
2724  
2725  
2726  
2727  
2728  
2729  
2730  
2731  
2732  
2733  
2734  
2735  
2736  
2737  
2738  
2739  
2740  
2741  
2742  
2743  
2744  
2745  
2746  
2747  
2748  
2749  
2750  
2751  
2752  
2753  
2754  
2755  
2756  
2757  
2758  
2759  
2760  
2761  
2762  
2763  
2764  
2765  
2766  
2767  
2768  
2769  
2770  
2771  
2772  
2773  
2774  
2775  
2776  
2777  
2778  
2779  
2780  
2781  
2782  
2783  
2784  
2785  
2786  
2787  
2788  
2789  
2790  
2791  
2792  
2793  
2794  
2795  
2796  
2797  
2798  
2799  
2800  
2801  
2802  
2803  
2804  
2805  
2806  
2807  
2808  
2809  
2810  
2811  
2812  
2813  
2814  
2815  
2816  
2817  
2818  
2819  
2820  
2821  
2822  
2823  
2824  
2825  
2826  
2827  
2828  
2829  
2830  
2831  
2832  
2833  
2834  
2835  
2836  
2837  
2838  
2839  
2840  
2841  
2842  
2843  
2844  
2845  
2846  
2847  
2848  
2849  
2850  
2851  
2852  
2853  
2854  
2855  
2856  
2857  
2858  
2859  
2860  
2861  
2862  
2863  
2864  
2865  
2866  
2867  
2868  
2869  
2870  
2871  
2872  
2873  
2874  
2875  
2876  
2877  
2878  
2879  
2880  
2881  
2882  
2883  
2884  
2885  
2886  
2887  
2888  
2889  
2890  
2891  
2892  
2893  
2894  
2895  
2896  
2897  
2898  
2899  
2900  
2901  
2902  
2903  
2904  
2905  
2906  
2907  
2908  
2909  
2910  
2911  
2912  
2913  
2914  
2915  
2916  
2917  
2918  
2919  
2920  
2921  
2922  
2923  
2924  
2925  
2926  
2927  
2928  
2929  
2930  
2931  
2932  
2933  
2934  
2935  
2936  
2937  
2938  
2939  
2940  
2941  
2942  
2943  
2944  
2945  
2946  
2947  
2948  
2949  
2950  
2951  
2952  
2953  
2954  
2955  
2956  
2957  
2958  
2959  
2960  
2961  
2962  
2963  
2964  
2965  
2966  
2967  
2968  
2969  
2970  
2971  
2972  
2973  
2974  
2975  
2976  
2977  
2978  
2979  
2980  
2981  
2982  
2983  
2984  
2985  
2986  
2987  
2988  
2989  
2990  
2991  
2992  
2993  
2994  
2995  
2996  
2997  
2998  
2999  
3000  
3001  
3002  
3003  
3004  
3005  
3006  
3007  
3008  
3009  
3010  
3011  
3012  
3013  
3014  
3015  
3016  
3017  
3018  
3019  
3020  
3021  
3022  
3023  
3024  
3025  
3026  
3027  
3028  
3029  
3030  
3031  
3032  
3033  
3034  
3035  
3036  
3037  
3038  
3039  
3040  
3041  
3042  
3043  
3044  
3045  
3046  
3047  
3048  
3049  
3050  
3051  
3052  
3053  
3054  
3055  
3056  
3057  
3058  
3059  
3060  
3061  
3062  
3063  
3064  
3065  
3066  
3067  
3068  
3069  
3070  
3071  
3072  
3073  
3074  
3075  
3076  
3077  
3078  
3079  
3080  
3081  
3082  
3083  
3084  
3085  
3086  
3087  
3088  
3089  
3090  
3091  
3092  
3093  
3094  
3095  
3096  
3097  
3098  
3099  
3100  
3101  
3102  
3103  
3104  
3105  
3106  
3107  
3108  
3109  
3110  
3111  
3112  
3113  
3114  
3115  
3116  
3117  
3118  
3119  
3120  
3121  
3122  
3123  
3124  
3125  
3126  
3127  
3128  
3129  
3130  
3131  
3132  
3133  
3134  
3135  
3136  
3137  
3138  
3139  
3140  
3141  
3142  
3143  
3144  
3145  
3146  
3147  
3148  
3149  
3150  
3151  
3152  
3153  
3154  
3155  
3156  
3157  
3158  
3159  
3160  
3161  
3162  
3163  
3164  
3165  
3166  
3167  
3168  
3169  
3170  
3171  
3172  
3173  
3174  
3175  
3176  
3177  
3178  
3179  
3180  
3181  
3182  
3183  
3184  
3185  
3186  
3187  
3188  
3189  
3190  
3191  
3192  
3193  
3194  
3195  
3196  
3197  
3198  
3199  
3200  
3201  
3202  
3203  
3204  
3205  
3206  
3207  
3208  
3209  
3210  
3211  
3212  
3213  
3214  
3215  
3216  
3217  
3218  
3219  
3220  
3221  
3222  
3223  
3224  
3225  
3226  
3227  
3228  
3229  
3230  
3231  
3232  
3233  
3234  
3235  
3236  
3237  
3238  
3239  
3240  
3241  
3242  
3243  
3244  
3245  
3246  
3247  
3248  
3249  
3250  
3251  
3252  
3253  
3254  
3255  
3256  
3257  
3258  
3259  
3260  
3261  
3262  
3263  
3264  
3265  
3266  
3267  
3268  
3269  
3270  
3271  
3272  
3273  
3274  
3275  
3276  
3277  
3278  
3279  
3280  
3281  
3282  
3283  
3284  
3285  
3286  
3287  
3288  
3289  
3290  
329

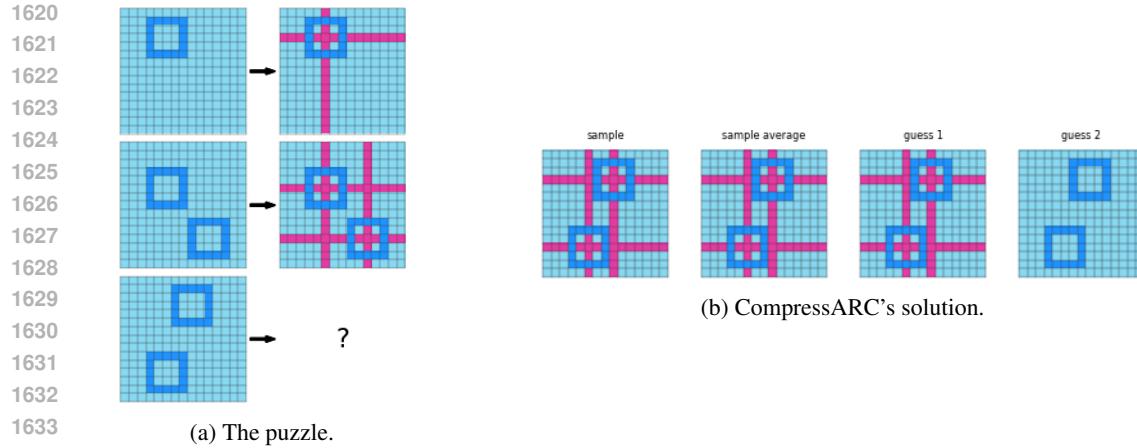


Figure 16: Center Cross: Puzzle 41e4d17e from the training split.

**CompressARC Solution:** We don't show CompressARC's solution evolving over time because we think it is uninteresting; instead will describe. We don't see much change in CompressARC's answer over time during learning. It starts by copying over the input grid, and at some point, magenta rows and columns start to appear, and they slowly settle on the correct positions. At no point does CompressARC mistakenly draw the rays on top of the bubbles; it has always had the order correct.

#### M.2.1 SOLUTION ANALYSIS: CENTER CROSS

Figure 17 shows another plot of the amount of information in every tensor in  $z$ . The only surviving tensors are the (color, channel) and (example, height, width, channel) tensors, which are interpreted in Figure 18.

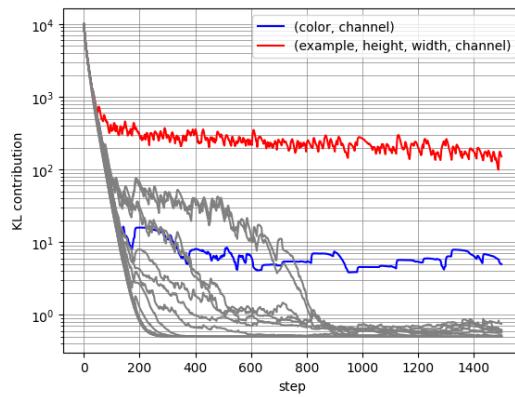


Figure 17: Breaking down the KL loss during training into contributions from each individual shaped tensor in the multitensor  $z$ .

1664  
1665  
1666  
1667  
1668  
1669  
1670  
1671  
1672  
1673

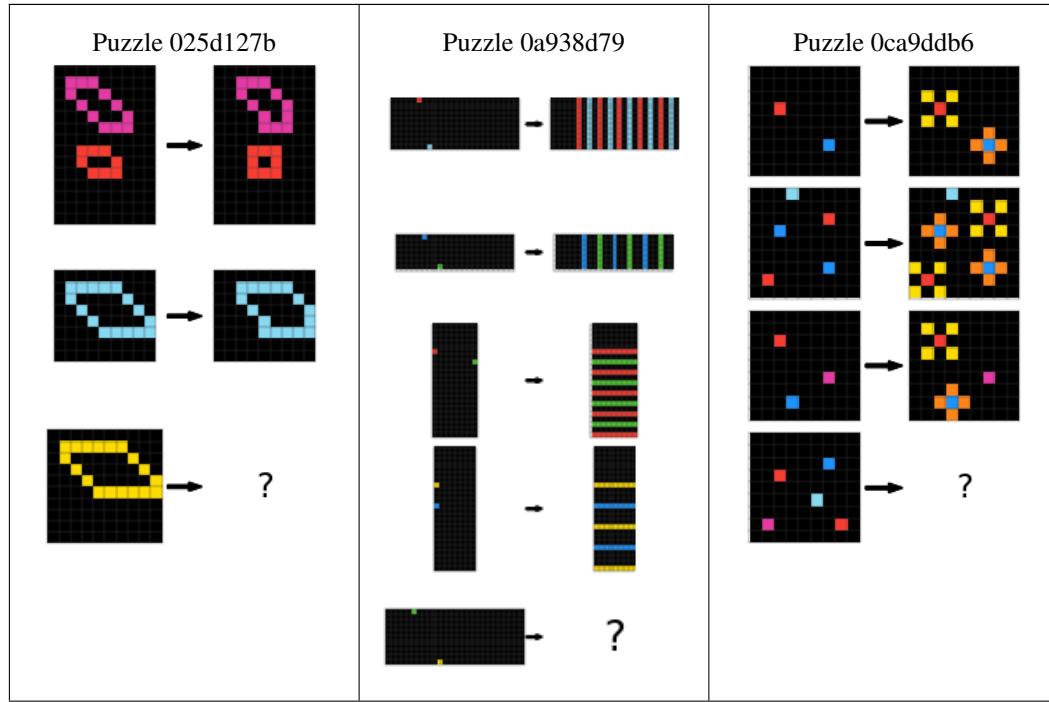
1674  
 1675  
 1676  
 1677  
 1678  
 1679  
 1680  
 1681  
 1682  
 1683  
 1684  
 1685  
 1686  
 1687  
 1688  
 1689  
 1690  
 1691  
 1692  
 1693  
 1694  
 1695  
 1696  
 1697  
 1698  
 1699  
 1700  
 1701  
 1702  
 1703  
 1704  
 1705  
 1706  
 1707  
 1708  
 1709  
 1710  
 1711  
 1712  
 1713  
 1714  
 1715  
 1716  
 1717  
 1718  
 1719  
 1720  
 1721  
 1722  
 1723  
 1724  
 1725  
 1726  
 1727

(a) **(example, height, width, channel) tensor.** The top principal component is 2496 times stronger than the second principal component. **This tensor codes for the centers of the bubbles.** In the KL contribution plot, we can see that the information content of this tensor is decreasing over time. Likely, CompressARC is in the process of eliminating the plus shaped representation, and replacing it with a pixel instead, which takes fewer bits.

Figure 18: Breaking down the loss components during training tells us where and how CompressARC prefers to store information relevant to solving a puzzle.

## N LIST OF MENTIONED ARC-AGI-1 PUZZLES

See Table 7 below.



1728

1729

1730

1731

1732

1733

1734

1735

1736

1737

1738

1739

1740

1741

1742

1743

1744

1745

1746

1747

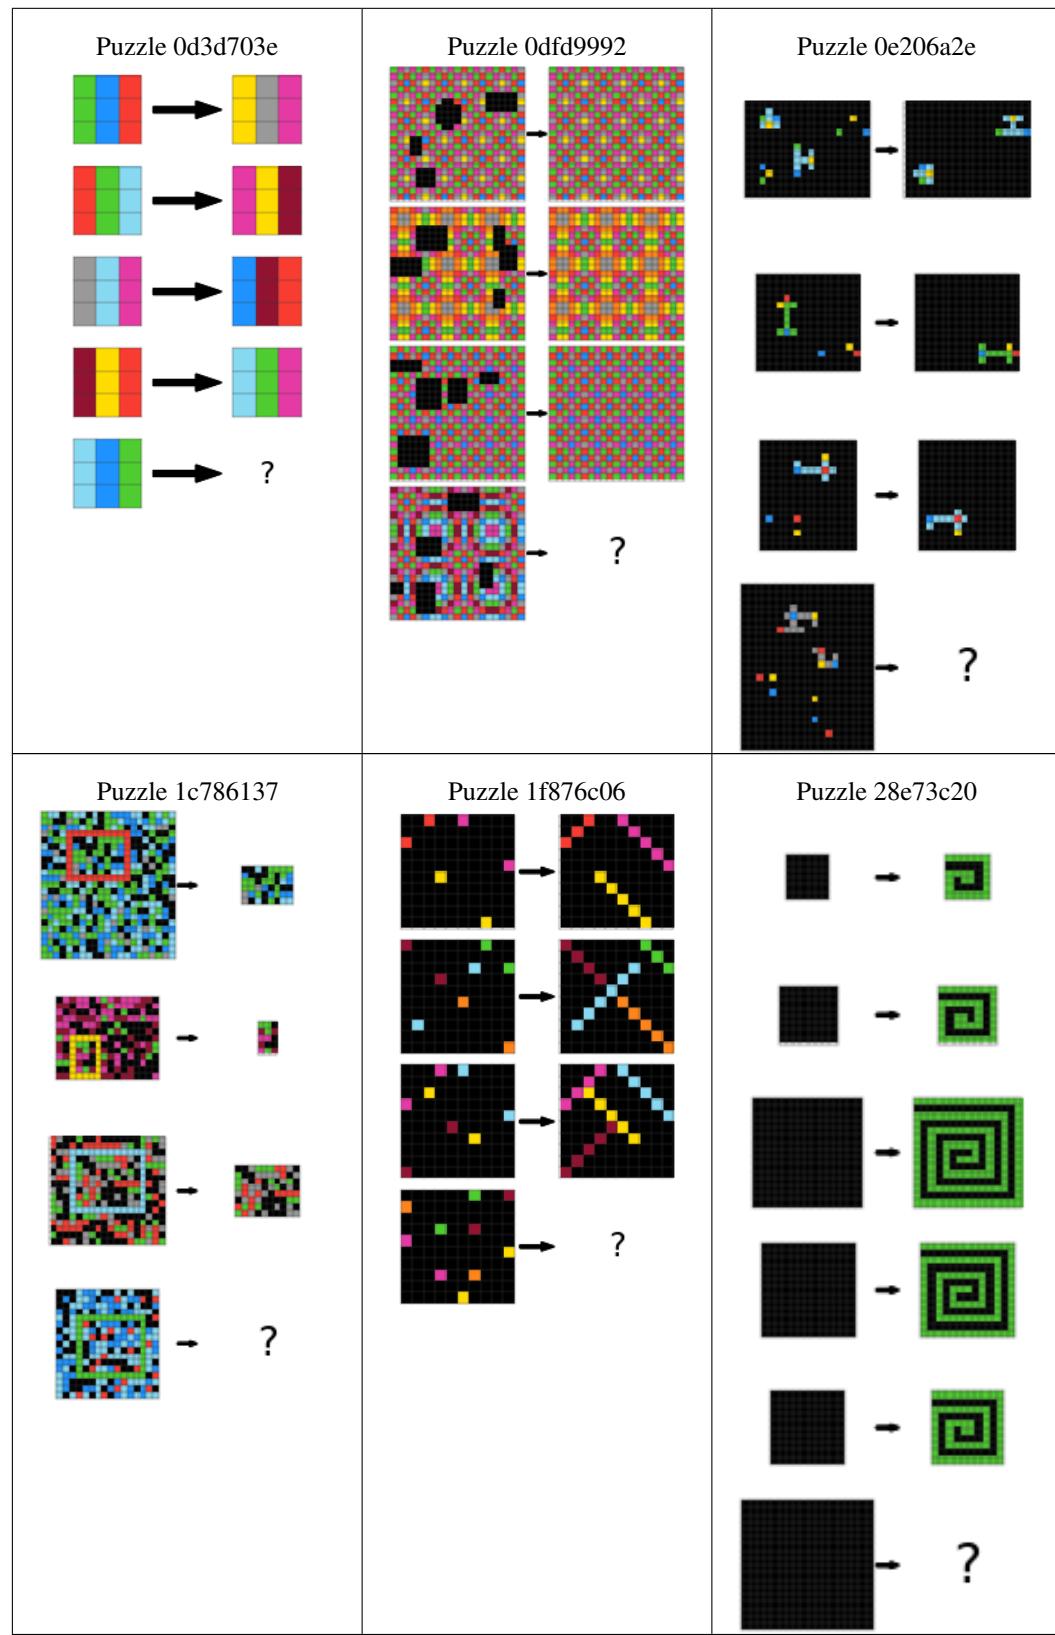
1748

1749

1750

1751

1752



1780

1781

1782

1783

1784

1785

1786

1787

1788

1789

1790

1791

1792

1793

1794

1795

1796

1797

1798

1799

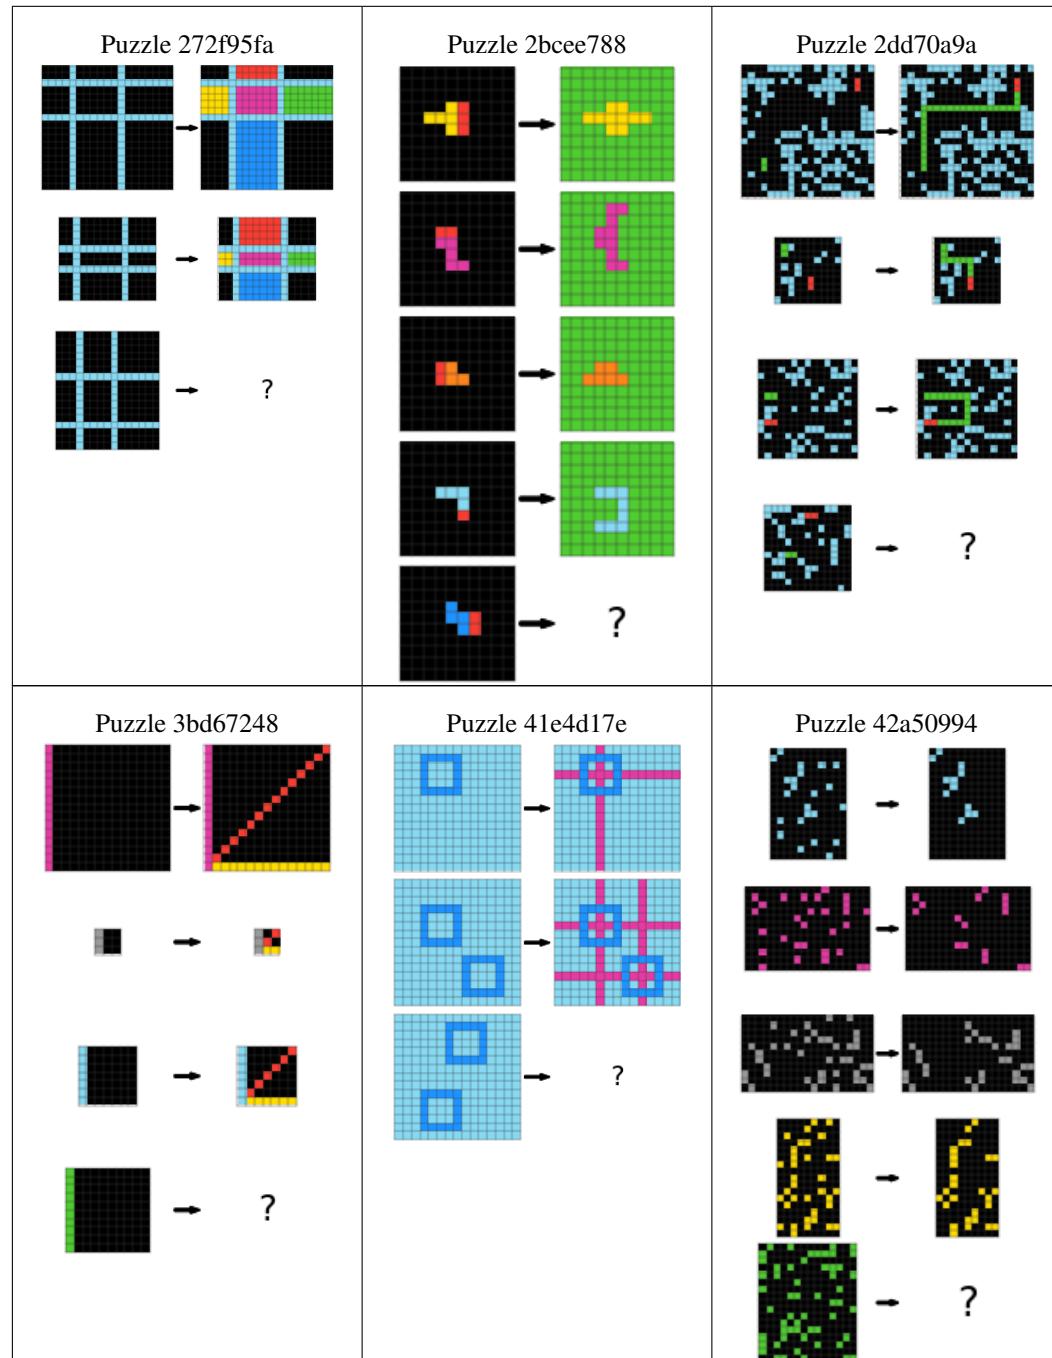
1800

1801

1802

1803

1804



1826

1827

1828

1829

1830

1831

1832

1833

1834

1835

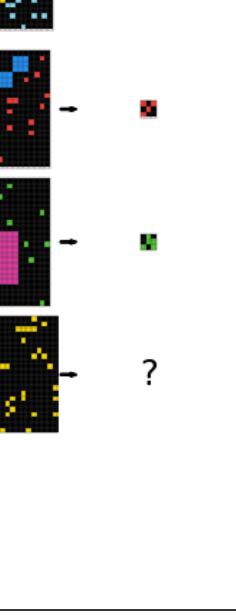
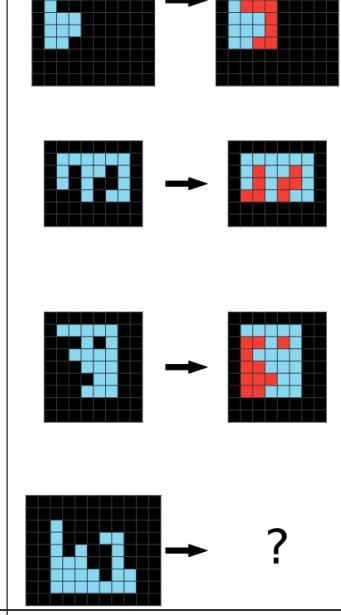
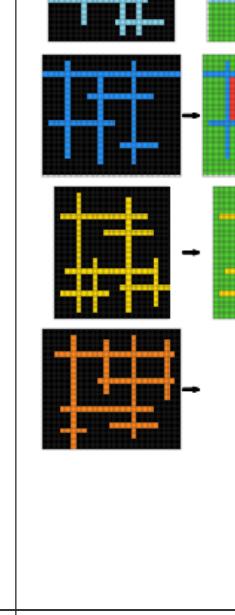
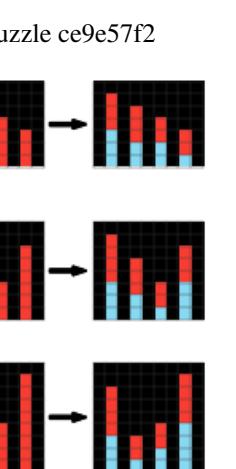
|  |   |  |
|--|---|--|
| <p>Puzzle 5ad4f10b</p>   | <p>Puzzle 6d75e8bb</p>  | <p>Puzzle 7b6016b9</p>  |
| <p>Puzzle ce9e57f2</p>  |   |  |

Table 7: List of Mentioned ARC-AGI=1 Puzzles. All these puzzles are part of the training split.

Q CODE

Code for this project is provided in the supplemental materials.