


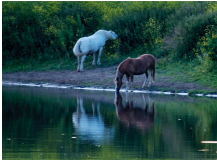
Image		
Question	Is the plate blue and round?	Do you see horses near the bushes?
MMN Results	<pre> graph TD A[Select(plate)] --> B[Image] B --> C[Verify(blue)] B --> D[Verify(round)] C --> E[True] D --> F[False] E --> G[And()] F --> G G --> H[False X] </pre>	<pre> graph TD A[Select(bush)] --> B[Image] B --> C[RelateInverse(near, horse)] C --> D[Image] D --> E[Exist()] E --> F[False X] </pre>
ProTo Results	<pre> graph TD A[Select(plate)] --> B[Image] B --> C[Verify(blue)] B --> D[Verify(round)] C --> E[True] D --> F[True] E --> G[And()] F --> G G --> H[True ✓] </pre>	<pre> graph TD A[Select(bush)] --> B[Image] B --> C[RelateInverse(near, horse)] C --> D[Image] D --> E[Exist()] E --> F[True ✓] </pre>

Figure 1: Visualization of validation results on the GQA dataset (Part 1). We use the green tick and the red cross to denote right and wrong reasoning result respectively. The MMN approach fails both examples while ProTo successfully addresses them.



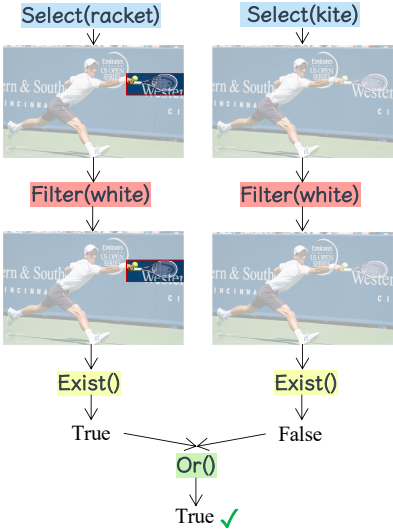
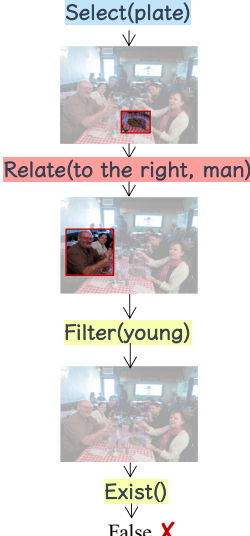
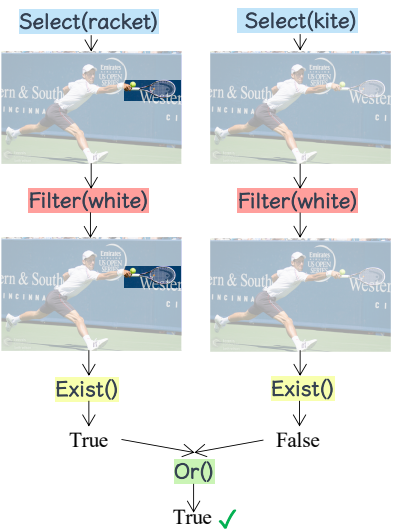
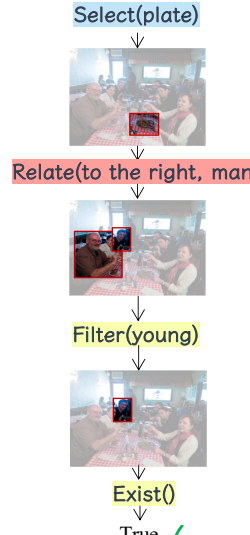
Image		
Question	Is there a kite or a racket that is white?	Is there a young man to the left of a plate?
MMN Results		
ProTo Results		

Figure 2: Visualization of validation results on the GQA dataset (Part 2). We use the green tick and the red cross to denote right and wrong reasoning results, respectively. Both models predict the correct answer in the first example, and the reasoning traces are the same. For the second example, our model gets the correct answer, but the MMN baseline fails.

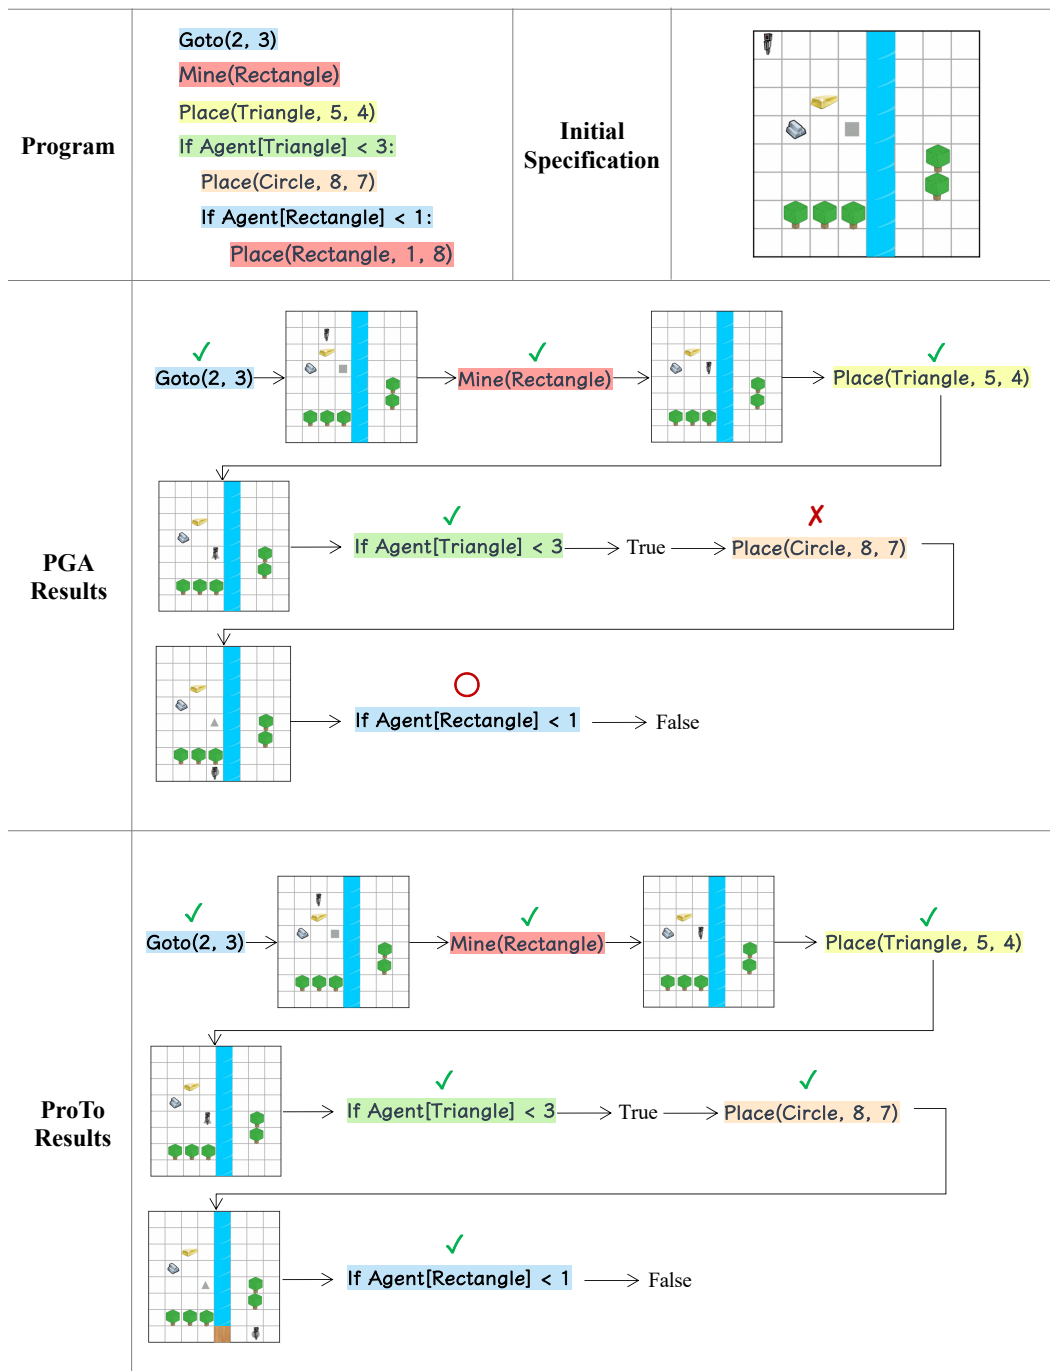


Figure 3: Visualization of test results on the Minecraft environment in comparison to program-guided agent [81] (Part 1). Due to space limitations, we omit the middle steps for executing routines and directly show the execution results of routines. We use the green tick and the red cross to denote right and wrong reasoning results, respectively. Additionally, we use a red circle to indicate that the agent has failed some routine (therefore, the program execution is failed). We observe that the ProTo agent can build a bridge to cross a river to place a circle on the other side of the river, but PGA fails to do so.



Figure 4: Visualization of test results on the Minecraft environment in comparison to program-guided agent [81] (Part 2). Due to space limitations, we omit the middle steps for executing routines and directly show the execution results of routines. We use the green tick and the red cross to denote right and wrong reasoning results, respectively. We observe that the ProTo agent can mine the gold in the corner, but PGA fails to do so within the time constraint.

A Extended Related Work

Transformers with Masked Attention We notice a few transformer-based architectures adopt masked attention to achieve different goals. First, the masked self-attention [87] is adopted to restrict the model from seeing subsequent positions of tokens for machine translation. Second, mask attention networks [28] uses mask matrices to enforce the localness modeling ability of transformers. Third, in the vision domain, attention masks are used to highlight specific classes of visual content (e.g., foreground objects) [78, 91]. One concurrent unpublished work [34] uses masked attention to model data flow relationship for source code summarization.

Neural Symbolic Learning The goal of neural symbolic learning is to pursue a coherent, unified view of symbolic logic-based computation and neural computation [9, 35, 36]. Previous neural symbolic systems involve visual reasoning [99, 64], logic induction [63], and reading comprehension [19]. We propose to leverage transformer architecture to integrate perception, reasoning, and decision. We believe transformer architecture would advance neural-symbolic systems because of its ability of compositional representation learning.

B Program Details

In this section, we provide more details of the programs on two experimented tasks.

B.1 Program-guided Visual Reasoning

All types of routines on the GQA dataset [47] are provided in Table 8. No loop or branching routines exist in the GQA datasets. In the GQA dataset, the entrance routine is the first routine in the program, and the exiting routine is the last routine of the program, which would produce the final answer to the program. Some types of routines may use more than one inputs.

There are three types of program inputs and outputs (execution results). The first type is `Objects`, which is a probabilistic distribution over detected objects [64]. The second type is `Boolean` that is either `True` or `False`. Finally, the `Answer` type is a distribution over all answer candidates. We acquire the ground true execution results by executing ground true programs on the ground true scene graphs. Since the dimensions of results for different types of routines are different, we use MLPs with different output dimensions to decode the result embeddings of different routines. Note that ProTo executes the program in a latent space, so no explicit inputs are directly sent to ProTo. Only latent embeddings are sent into ProTo. The latent result embeddings are grounded to the explicit results via execution losses.

All the routines on the GQA datasets are single-step routines, which are finished via only a single forward step. In other words, we always update the pointer at each execution step.

B.2 Program-guided Policy Learning

We list different types of routines on the 2D Minecraft datasets in Table 5. In this dataset, the entrance routine is the first routine in the program, and the exiting routines are routines that might be executed at the end of the program execution⁴.

All routines in 2D Minecraft datasets take the specification as input and output either actions or Boolean results. The actions can either be motor actions or interactive actions as described in Sec 5.2 of the main text. Like the GQA experiments, we also use MLPs with different output dimensions to decode the results embeddings.

The list of ending conditions is also provided in Table 5. Note that the PGA baseline [81] uses the same set of end conditions.

We visualize the distribution of program lengths on both GQA and Minecraft in Figure D6.

⁴In branching cases, the last routines of both branches are possibly executed, so they are all exiting routines.

C Algorithm Details

C.1 Derive Parents of Routines

For the GQA dataset, the parents of routines are provided in the ground truths. For example, in the Figure 2 of main text, the parents of the third routine (`Verify_relation(left)`) are the first routine (`Select(bag)`) and the second routine (`Select(wine)`).

For the Minecraft dataset, the parent of one routine is the previously executed routine. Each routine only depends on the previous routine, and no routines rely on more than one routine.

C.2 Parallel Execution

Since ProTo is a transformer-like [87] architecture, it has a promising ability to execute many routines in parallel when they have no result dependency. A typical example is shown in Figure C5. The parallel version of algorithms is shown in Algorithm 3 and Algorithm 4. We adopt a vector of pointers \mathbf{p} to point to different routines executed in parallel.

The semantic guidance remains the same as Eq. 1. We revise the structure guidance \mathbf{P} to support executing many routines in one time as the following equation:

$$\mathbf{P}^{t(\tau)}[i][j] = \begin{cases} 0 & \text{if } (\exists p \in \mathbf{p}, \text{ s.t. } P_j \in \text{Parents}(P_p) \text{ and } P_i = P_p) \text{ or } (i = j), \\ -\infty & \text{else.} \end{cases} \quad (6)$$

Parallel execution is only enabled in the GQA experiments. On the Minecraft datasets, there is only one agent who needs to perform all the required tasks. So the routines in the Minecraft experiments can not be executed in parallel. We expect this parallel execution feature can be tested on a multi-agent environment [57] in the future.

Note that in the main text, we explain our approach sequentially for ease of understanding.

Algorithm 3: ProTo Execution in Parallel

Result: Execution results $\{o^{(\tau)}\}_{\tau=1}^T$

- 1 Initialize $\tau = 0$ and a pointer;
- 2 Build \mathbf{P}^s according to Eq. 1;
- 3 **while** not reach an exiting routine **do**
- 4 Observe $s^{(\tau)}$ and set $\tau = \tau + 1$;
- 5 **if** there are routines that can be executed in parallel **then**
- 6 Spawn a vector of pointers \mathbf{p} to point to those routines;
- 7 Build $\mathbf{P}^{t(\tau)}$ according to Eq. 6;
- 8 $o^{(\tau)} =$
- ProToInfer($s^{(\tau)}, \mathbf{P}^s, \mathbf{P}^{t(\tau)}, \mathbf{p}$);
- 9 Output $o^{(\tau)}$;
- 10 ParaUpdatePointer($\mathbf{p}, o^{(\tau)}$);
- 11 **end**

Algorithm 4: ParaUpdatePointer($\mathbf{p}, o^{(\tau)}$)

- 1 **if** parallel routines finishes execution **then**
- 2 Merge pointers \mathbf{p} .
- 3 **for** one pointer p in \mathbf{p} **do**
- 4 **if** P_p does not finish execution **then**
- 5 **return**
- 6 **if** P_p is an *If*-routine **then**
- 7 Point p to the first routine in the *T/F* branch if the result of P_p is *T/F*;
- 8 **else if** P_p is a *While*-routine **then**
- 9 Point p to the first routine inside/outside the loop if the result of P_p is *T/F*;
- 10 **else if** P_p ends a loop **then**
- 11 Point p to the loop condition routine.
- 12 **else**
- 13 Point p to the subsequent routine;
- 14 **end**
- 15 **end**

D Experimental Details and Further Results

In this section, we provide more details and further results on two experimented tasks.

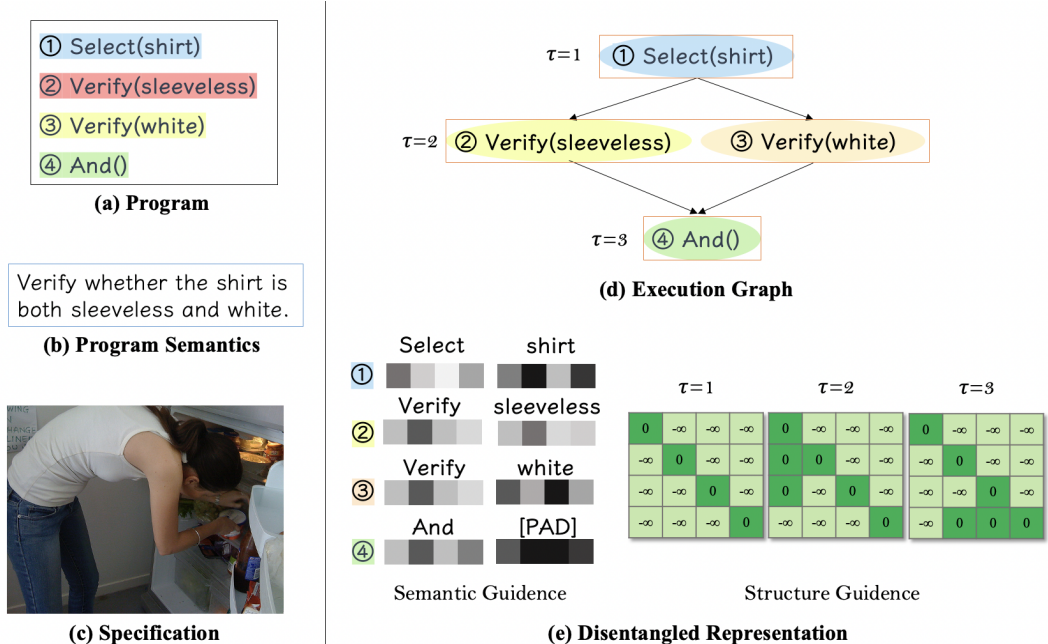


Figure C5: An example of ProTo parallel execution. We show the program and its semantics in (a) and (b). The specification is shown in (c). We show the execution flow in (d) where we can see we execute two routines at $\tau = 2$. In (e), we show the semantic guidance and the structure guidance of ProTo execution process.

D.1 Program-guided Visual Reasoning

D.1.1 Program Synthesis Model

We adopt a simple transformer-based seq2seq model [87] to translate a natural language question into a program. Both the encoder and the decoder of the seq2seq model are composed of six identical self-attention layers with hidden feature dimensions $d_h = 512$. The head number is eight.

The input question is encoded in a token-by-token manner via a learnable dictionary ϕ_q . We turn the ground true program into a sequence by traversing the program tree via pre-order traverse. Segment tokens [SEG] are added between two routines. The predicted sequential program can be recovered via a reverse way. We used beam search with a beam size of 4 and length penalty $\alpha = 0.6$ [87]. The validation accuracy of this model is 98.1%.

D.1.2 Program Representation

In the semantic part of the program, we set $d_m = 256$, $L = 8$ and $d = 2048$. In the implementation of the structure part, we use -1×10^9 as the negative infinity, which is the same as the standard implementation of masked attention in transformers [87].

D.1.3 Visualizations

We provide more visualization of ProTo on the validation datasets in comparison to the concurrent work meta module networks (MMN) [15]. The implementation and hyperparameters of meta module networks follow their official code. Specifically, for the Objects types of results, we visualize the predicted object with a probability $p > 0.5$. For the results with type Boolean or Answer, we present the choice with maximum probability. The visualizations are shown in Figure 1 and Figure 2.

D.1.4 Details of the GAT Encoding

During the ablation study, we compare our model to graph attention networks (GAT) [88]. The node features are semantic embeddings of routines, and the edges represent message passing relationships.

Specifically, the node features are constructed via a concatenation of word embeddings, which is the same as Eq. 1. The embedding dimension is the same as ProTo. One edge exists between a routine and its parents as in Eq. 2. Note that since the GQA programs do not have conditional routines such as `While` and `If`, the edges are determined before execution. The routine embeddings are fed to two cross-attention modules to fuse information of result embeddings and specifications following Eq 3 and Eq 5. We also use an MLP to decode the latent results to get explicit routine results.

The GAT model consists of three layers, where each layer has eight attention heads with 256 features, following by an ELU nonlinearity.

D.1.5 Details of Generalization Experiments

Purpose of Collecting Additional Human-written Programs We have the following reasons for collecting the human-written programs. First, we are curious whether humans can communicate with machines via programs, which has not been done by previous work before. Second, the GQA questions and programs are synthetic, and many of the programs are awkward (e.g., with many unnecessary modifiers such as "the baked good that is on the top of the plate that is on the left of the tray"). Third, the GQA test split programs are not publicly available, and the translated programs from the questions may be inaccurate. Since the validation split has been used for parameter tuning, we wish to benchmark program-guided visual reasoning on the collected independent data points. Forth, this small-scale dataset lays the ground for the construction of our novel dataset for program-guided tasks.

GQA-Humam-Program Dataset Collection Process For the Human Program Guidance experiments, we create the GQA-Humam-Program dataset to diagnose whether humans can guide the reasoning process via programs. We employ five volunteers to write 500 programs and answers on the GQA validation dataset. The estimated hourly wage is ten dollars, and the total amount spent on volunteers is two thousand dollars. A parser checks the written programs to ensure that they follow the domain specification language of GQA. We encourage the volunteers to write longer and more complex programs. Two volunteers cross-check the correctness of programs and answers. For fairness of comparison, we retrain the meta neural module networks [15] on the training split of GQA while preventing it from seeing the natural language questions. The screenshot of the data collection tool is provided in Figure D7.

Rationale Behind Unseen Programs Experiments In the experiments of Unseen Programs, the models are required to learn combinatorial word-level semantics to execute unseen programs. The training set contains `verify_size` and `filter_shape`, the models may generalize compositionally to the unseen program `verify_shape`.

Restricted Data Regime Experiments We repeat for three random seeds and found the standard deviation of the results is smaller than 2%.

D.1.6 Computational Resources

We train our model and the baselines on a 48 core Ubuntu 16.04 Linux server with eight Nvidia Titan-X GPU. The CPU is Intel Silver 4116 CPU @ 2.10GHz. The total training time is around 48 hours.

D.1.7 License and Permissions

The GQA dataset is built upon Visual Genome [53], which is under Creative Commons Attribution 4.0 International License. The GQA dataset is publicly available so that we can use it for research purposes. The GQA dataset is used in many published literature [47, 45], and we do not found offensive content in this dataset.




Figure D6: The program length distribution on the GQA and the Minecraft dataset.

D.2 Program-guided Policy Learning

D.2.1 Environment Details

The major environmental resources that the agent can interact with are gold, wood, or iron. The environment may contain a river, and the agent cannot go across unless a bridge is built. The size of the grid world ranges from five to eight. There could be two to four merchants in the environment. The environment is randomly initialized, and the agent is also randomly initialized while ensuring the program can be finished. If the agent fails to finish the program within 300 timesteps, the execution would be terminated (timeout).

D.2.2 A2C and Hyper-parameters

We use the same implementation of the A2C algorithm as [81] (provided by its authors). The A2C algorithm uses a learning rate of 1×10^{-3} , 64 environments running in parallel with 64 number of workers. The number of roll-out steps for each update is five. The agent is trained for 10^7 timesteps. The balance of the entropy regularization term is $\beta = 0.1$.

D.2.3 More Visualizations

We provide more visualizations of the results on the test splits in Figure 3 and Figure 4.

D.2.4 Architecture Details and Computational Costs

We list the computational costs and details of ProTo and the baselines in Table 6.

The server that we used is the same as the GQA experiments, but we only use a single GPU in the Minecraft experiments following [81]. The total training time is around 80 hours.

D.2.5 Details about the Planner Used in Dense Supervision

We create a planner to generate ground true execution traces to train the ProTo baselines with full supervision. The planner uses a hard-coded interpreter to parse the programs. For all the actions that need navigation, the planner uses an A* search algorithm [26] to find the shortest path. For the Bridge, Mine and Sell action, we would find the nearest river, the nearest item or the nearest merchant. For the If and While routines, the planner uses the symbolic information in the environment (e.g. `env[gold]=3`) to decide whether the conditions are satisfied.

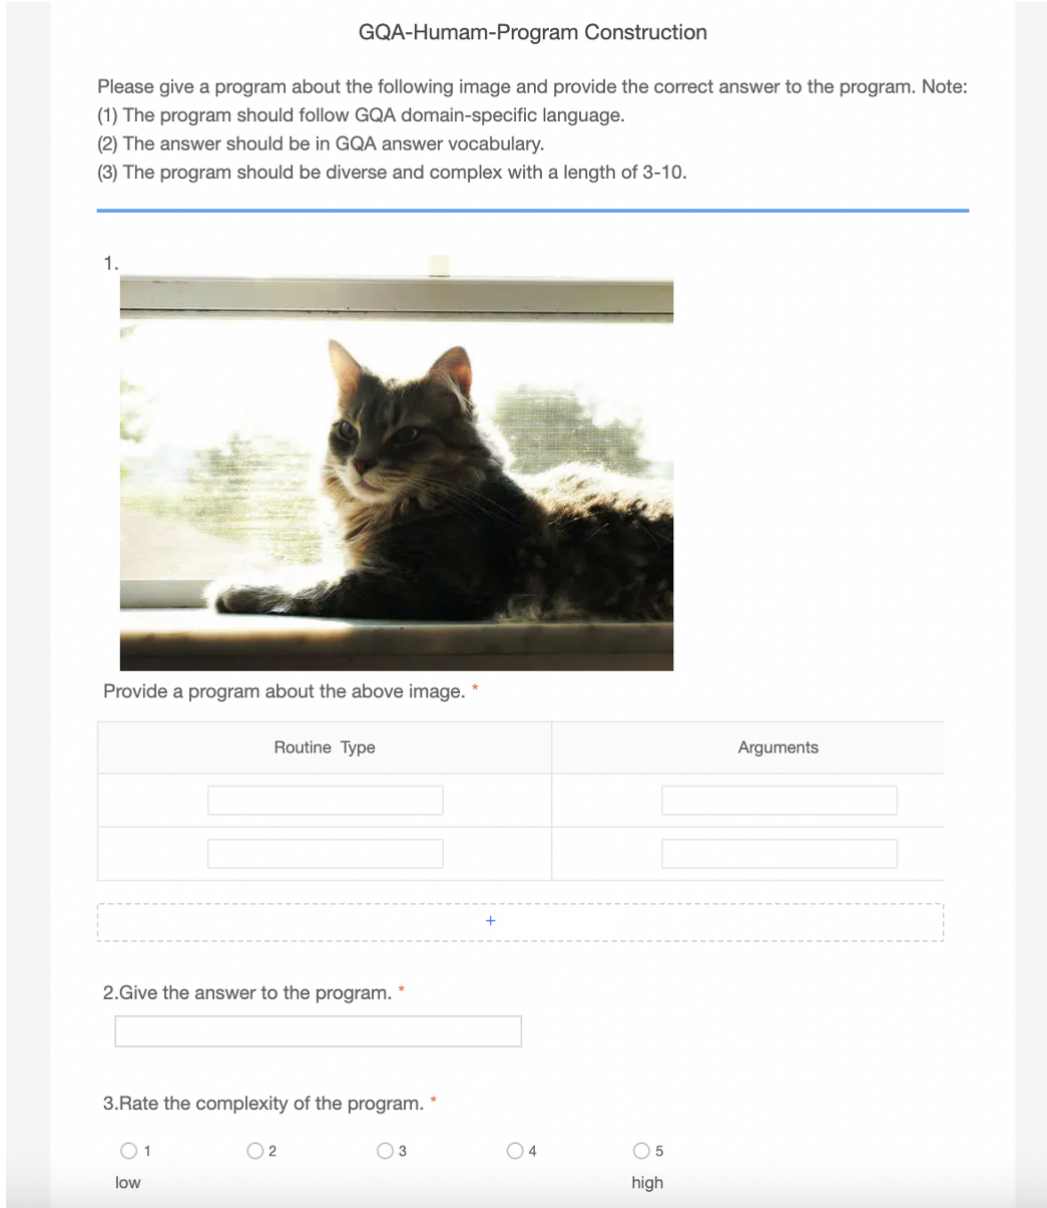


Figure D7: Screenshot of the UI for the construction of the GQA-Humam-Program dataset. The submitted programs are checked by the parser and the collected answers are cross-checked by two more annotators.

D.2.6 License and Permissions

The Minecraft dataset is under Creative Commons Attribution 4.0 International License. We acquire this dataset and its license from the authors of [81]. Since it's a synthetic dataset, we don't think it has offensive content.

E Additional Discussions about the Neural-Symbolic Baselines

The comparison between different neural symbolic approaches is a significant aspect of our work. On the first domain of visual reasoning, our paper shows that a mixed parametric neural-symbolic model would outperform pure symbolic non-parametric executors (Table 2). On the second domain of

Table 5: All types of routines in the domain-specific language of the 2D Minecraft dataset [81]. We use Spec. to denote specification. Some routines require many arguments, which are numbered by integers with brackets (e.g., (1), (2), etc.).

Type	Arguments	Input	Output	Semantics and Ending Conditions
Mine	Triangle, circle, rectangle, gold, wood, or iron	Spec.	Action	Go to the item and pick it up. Ended when the action "Mine" is predicted and the agent successfully mined a gold.
BuildBridge	-	Spec.	Action	Go to the river and build a bridge. Ended when the action "Bridge" is predicted and the agent successfully builds a bridge.
Goto	Coordinates	Spec.	Action	Go to the coordinates. Ended when the agent reaches the target.
Place	(1) Triangle, circle or rectangle (2) Coordinates	Spec.	Action	Place the object on the specified coordinates. Ended when the action "Place" is predicted.
Sell	Triangle, circle, rectangle, gold, wood, or iron	Spec.	Action	Go to the merchant and sell mined items. Ended when the action "Sell" is predicted.
If / If-Else	(1) Agent, env, or is_there (2) Gold, wood, iron, bridge, river, merchant, wall, or flat (3) Operators (>, ≥, =, <, ≤) (4) Number	Spec.	Boolean	Perform branching based on the condition in the if-clause. Ended in a single execution step.
While	(1) Agent, env, or is_there (2) Gold, wood, iron, bridge, river, merchant, wall, or flat (3) Operators (>, ≥, =, <, ≤) (4) Number	Spec.	Boolean	Perform looping based on the condition in the while-clause. Ended in a single execution step.

Minecraft, we have compared neural baselines (Table 4). We also have a pure symbolic (non-learning) method: the planner. The advantages of the learned executor beyond the symbolic planner are as follows.

First, the hard-coded planner requires a lot of ad-hoc engineering work to handle complex cases. But our method can learn from a sparse reward signal. For example, our method can learn to build a bridge to cross the river to fetch gold on the other side of the river without explicit supervision (just given a reward signal). However, a planner needs to handle this case with special treatment.

Second, we experiment on Minecraft to validate the ability of ProTo to generalize across unseen routines. Specifically, we remove a routine from the training split (e.g., Mine(Gold)) and test on programs that contain the removed routine. This experiment validates the compositional generalization ability (e.g., generalize to Mine(Gold) after seeing Is_there(gold) and Mine(Silver)). Only a

Table 6: Architecture details of ProTo and the baselines on the 2D MineCraft dataset.

Model	Parameters	Architecture Details
ProTo	1.30M	Eight attention heads with an intermediate size of 64; shared computation between routines; the state map is encoded via a two-layer MLP with hidden size 256 and output size $d = 2048$; the inventory is encoded via another two-layer MLP with hidden size 128 and output size $d = 2048$; the inventory feature is added to the state map to produce the final specification feature; the output MLP is also a two-layer MLP with hidden size 256.
PGA	1.21M	The state map is encoded via a batch of CNNs with channel sizes of 32, 64, 96, and 128. Each convolutional layer has kernel size three and stride 2, which is followed by ReLU nonlinearity. The inventory is encoded via a two-layer MLP with a channel size of 256. The goal is encoded via a two-layer MLP with a channel size of 64. The features are fused via a modulation mechanism proposed by PGA [81].
Vanilla Transformer	2.63M	Two attentional layers stacked, with eight attention heads, a hidden size of 128, and an intermediate size of 256.
Tree-RNN	0.51M	Program embeddings are of dimension 128. Attention LSTM size of 128. Tree-RNN uses a composition module to aggregate all the children representation of a node, which is of size $[128 \times 128]$, and output projection weights of size $[128 \times 128]$, with a bias of size 128. The program embeddings are average pooled across one routine so that each routine will be mapped to a fixed dimension. The composition layer is applied when combining pooled embedding from all the children of a node.

reward signal is used for training. Other experimental details are the same as described in Sec 5.2 (The planner is set to choose a random legal action when meeting an unseen routine). The results on the validation dataset are presented in Tab 7.

Third, the planner cannot scale up to a large number of states. A planner cannot work well on a large-scale scenario such as the game GO [76].

Fourth, the planner can never work on raw image observations. Note our transformer-based architecture can work on raw image input after dividing the observation into patches [25] or detecting objects in the raw image. But the symbolic planner has no way to work on raw image inputs.

F Limitations

Despite our contributions to task formulation and ProTo models, our work has several limitations. First, since the program-based approaches need to leverage the program guidance and dense supervision, it cannot easily leverage large-scale datasets for self-supervised pretraining [61]. We would work on building large-scale datasets with program annotations to alleviate this limitation. Second,

Table 7: Comparison between the generalization ability of the planner and ProTo on the Minecraft dataset.

Removed Routine	Planner Acc	ProTo Acc
Mine(Gold)	11.4	59.3
Is_there(River)	43.8	65.6
Agent [Silver]	46.2	77.1

we hypothesize that one can still improve proTo’s architecture design since the community has not exploited the power of transformers. We would incorporate recent advances in transformers [59, 100] to improve ProTo. Furthermore, we could conduct more ablation studies to reveal the importance of different components in ProTo.

G Broader Impact

Our findings provide a simple yet effective approach to address program-guided tasks. Potentially, people may leverage ProTo models to instruct robots via programs [69]. However, the learned executors may be attacked by adversarial training [62]. And instructing machines via programs might require humans to have more advanced knowledge (e.g., knowing the basic concepts of programs and how to follow program syntax). Therefore, those of a low educational level may not be able to leverage the benefits of our approach, which might aggravate social inequalities.

Table 8: All types of routines in the domain-specific language of the GQA dataset [47].

Type	Arguments	Input	Output	Semantics
Select/Filter	Position, color, material, shape, or activity	Objects	Objects	Filter out a set of objects by the positions, colors, etc from the input objects.
Choose	Name, scene, color, shape, position or attributes	Objects	Answer	Choose one answer (e.g., name, scene, color) from given answer candidates.
Verify	Color, shape, scene, or relation	Objects	Boolean	Verify whether the given concepts (e.g., color, shape) holds true for the input objects.
Relate InverseRelate	Name, or attributes	Objects	Objects	Filter out a set of objects that have the relation concept (e.g., names, attributes) with the input objects.
Query	Name, color, shape, scene or position	Objects	Answer	Query the concept (e.g., name, color) of the input objects.
Common	Color, or material	Two objects	Answer	Query the common concepts (e.g., color, material) of the input objects.
Different	Name, color, or material	Two objects	Boolean	Return whether the concepts of the objects (e.g., name, color) are different.
Same	Name or color	Two objects	Boolean	Return whether the concepts of the objects (e.g., name, color) are same.
And	-	Two booleans	Boolean	Return whether the two input booleans are both True.
Or	-	Two booleans	Boolean	Return whether one of the two input conditions is True.
Exist	-	Objects	Boolean	Return whether the input object set is not empty.

References

- [1] Frances E Allen. Control flow analysis. *ACM Sigplan Notices*, 5(7):1–19, 1970.
- [2] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 6077–6086, 2018.
- [3] Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, Ian Reid, Stephen Gould, and Anton Van Den Hengel. Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3674–3683, 2018.
- [4] David Andre and Stuart J Russell. *Programmable reinforcement learning agents*. PhD thesis, University of California, Berkeley, 2003.
- [5] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pp. 2425–2433, 2015.
- [6] Brenna D Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. A survey of robot learning from demonstration. *Robotics and autonomous systems*, 57(5):469–483, 2009.
- [7] Andrew G Barto and Sridhar Mahadevan. Recent advances in hierarchical reinforcement learning. *Discrete event dynamic systems*, 13(1):41–77, 2003.
- [8] Osbert Bastani, Yewen Pu, and Armando Solar-Lezama. Verifiable reinforcement learning via policy extraction. *arXiv preprint arXiv:1805.08328*, 2018.
- [9] Tarek R. Besold, Artur S. d’Avila Garcez, Sebastian Bader, Howard Bowman, Pedro M. Domingos, Pascal Hitzler, Kai-Uwe Kühnberger, Luís C. Lamb, Daniel Lowd, Priscila Machado Vieira Lima, Leo de Penning, Gadi Pinkas, Hoifung Poon, and Gerson Zaverucha. Neural-symbolic learning and reasoning: A survey and interpretation. *CoRR*, abs/1711.03902, 2017. URL <http://arxiv.org/abs/1711.03902>
- [10] David Bieber, Charles Sutton, Hugo Larochelle, and Daniel Tarlow. Learning to execute programs with instruction pointer attention graph neural networks. *CoRR*, abs/2010.12621, 2020. URL <https://arxiv.org/abs/2010.12621>
- [11] Paul Booth. *An introduction to human-computer interaction (psychology revivals)*. Psychology Press, 2014.
- [12] Rudy Bunel, Matthew Hausknecht, Jacob Devlin, Rishabh Singh, and Pushmeet Kohli. Leveraging grammar and reinforcement learning for neural program synthesis. *arXiv preprint arXiv:1805.04276*, 2018.
- [13] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European Conference on Computer Vision*, pp. 213–229. Springer, 2020.
- [14] David Chen and Raymond Mooney. Learning to interpret natural language navigation instructions from observations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 25, 2011.
- [15] Wenhui Chen, Zhe Gan, Linjie Li, Yu Cheng, William Wang, and Jingjing Liu. Meta module network for compositional visual reasoning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 655–664, 2021.
- [16] Xinyun Chen, Chang Liu, Richard Shin, Dawn Song, and Mingcheng Chen. Latent attention for if-then program synthesis. *arXiv preprint arXiv:1611.01867*, 2016.
- [17] Xinyun Chen, Chang Liu, and Dawn Song. Towards synthesizing complex programs from input-output examples. *arXiv preprint arXiv:1706.01284*, 2017.
- [18] Xinyun Chen, Chang Liu, and Dawn Song. Execution-guided neural program synthesis. In *International Conference on Learning Representations*, 2018.
- [19] Xinyun Chen, Chen Liang, Adams Wei Yu, Denny Zhou, Dawn Song, and Quoc V Le. Neural symbolic reader: Scalable integration of distributed and symbolic representations for reading comprehension. In *International Conference on Learning Representations*, 2019.

- [20] Xinyun Chen, Chen Liang, Adams Wei Yu, Dawn Song, and Denny Zhou. Compositional generalization via neural-symbolic stack machines. *arXiv preprint arXiv:2008.06662*, 2020.
- [21] Abhishek Das, Samyak Datta, Georgia Gkioxari, Stefan Lee, Devi Parikh, and Dhruv Batra. Embodiedquestionanswering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1–10, 2018.
- [22] Jacob Devlin, Rudy Bunel, Rishabh Singh, Matthew Hausknecht, and Pushmeet Kohli. Neural program meta-induction. *arXiv preprint arXiv:1710.04157*, 2017.
- [23] Jacob Devlin, Jonathan Uesato, Surya Bhupatiraju, Rishabh Singh, Abdel-rahman Mohamed, and Pushmeet Kohli. Robustfill: Neural program learning under noisy i/o. In *International conference on machine learning*, pp. 990–998. PMLR, 2017.
- [24] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [25] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [26] František Duchoň, Andrej Babinec, Martin Kajan, Peter Beňo, Martin Florek, Tomáš Fico, and Ladislav Jurišica. Path planning with modified a star algorithm for a mobile robot. *Procedia Engineering*, 96: 59–69, 2014.
- [27] Kevin Ellis, Daniel Ritchie, Armando Solar-Lezama, and Joshua B Tenenbaum. Learning to infer graphics programs from hand-drawn images. *arXiv preprint arXiv:1707.09627*, 2017.
- [28] Zhihao Fan, Yeyun Gong, Dayiheng Liu, Zhongyu Wei, Siyuan Wang, Jian Jiao, Nan Duan, Ruofei Zhang, and Xuanjing Huang. Mask attention networks: Rethinking and strengthen transformer. *arXiv preprint arXiv:2103.13597*, 2021.
- [29] John K Feser, Marc Brockschmidt, Alexander L Gaunt, and Daniel Tarlow. Differentiable functional program interpreters. *arXiv preprint arXiv:1611.01988*, 2016.
- [30] Richard E Fikes and Nils J Nilsson. Strips: A new approach to the application of theorem proving to problem solving. *Artificial intelligence*, 2(3-4):189–208, 1971.
- [31] Roy Fox, Richard Shin, Sanjay Krishnan, Ken Goldberg, Dawn Song, and Ion Stoica. Parametrized hierarchical procedures for neural programming. *ICLR 2018*, 2018.
- [32] Daniel Fried, Ronghang Hu, Volkan Cirik, Anna Rohrbach, Jacob Andreas, Louis-Philippe Morency, Taylor Berg-Kirkpatrick, Kate Saenko, Dan Klein, and Trevor Darrell. Speaker-follower models for vision-and-language navigation. *arXiv preprint arXiv:1806.02724*, 2018.
- [33] Peng Gao, Jiasen Lu, Hongsheng Li, Roozbeh Mottaghi, and Aniruddha Kembhavi. Container: Context aggregation network. *CoRR*, abs/2106.01401, 2021. URL <https://arxiv.org/abs/2106.01401>.
- [34] Shuzheng Gao, Cuiyun Gao, Yulan He, Jichuan Zeng, Lun Yiu Nie, and Xin Xia. Code structure guided transformer for source code summarization. *CoRR*, abs/2104.09340, 2021. URL <https://arxiv.org/abs/2104.09340>.
- [35] Artur d’Avila Garcez, Marco Gori, Luis C Lamb, Luciano Serafini, Michael Spranger, and Son N Tran. Neural-symbolic computing: An effective methodology for principled integration of machine learning and reasoning. *arXiv preprint arXiv:1905.06088*, 2019.
- [36] Artur SD’Avila Garcez, Luis C Lamb, and Dov M Gabbay. *Neural-symbolic cognitive reasoning*. Springer Science & Business Media, 2008.
- [37] Felix A Gers, Jürgen Schmidhuber, and Fred Cummins. Learning to forget: Continual prediction with lstm. *Neural Computation*, 1999.
- [38] Alex Graves, Greg Wayne, and Ivo Danihelka. Neural turing machines. *arXiv preprint arXiv:1410.5401*, 2014.
- [39] Sumit Gulwani. Program synthesis. *Software Systems Safety*, pp. 43–75, 2014.
- [40] Stevan Harnad. Grounding symbols in the analog world with neural nets. *Think*, 2(1):12–78, 1993.

- [41] Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. *arXiv preprint arXiv:1606.03476*, 2016.
- [42] Ronghang Hu, Jacob Andreas, Trevor Darrell, and Kate Saenko. Explainable neural computation via stack neural module networks. In *Proceedings of the European conference on computer vision (ECCV)*, pp. 53–69, 2018.
- [43] Haoyang Huang, Yaobo Liang, Nan Duan, Ming Gong, Linjun Shou, Daxin Jiang, and Ming Zhou. Unicoder: A universal language encoder by pre-training with multiple cross-lingual tasks. *arXiv preprint arXiv:1909.00964*, 2019.
- [44] Wenlong Huang, Igor Mordatch, and Deepak Pathak. One policy to control them all: Shared modular policies for agent-agnostic control. In *International Conference on Machine Learning*, pp. 4455–4464. PMLR, 2020.
- [45] Drew Hudson and Christopher D Manning. Learning by abstraction: The neural state machine. In *Advances in Neural Information Processing Systems*, pp. 5903–5916, 2019.
- [46] Drew A Hudson and Christopher D Manning. Compositional attention networks for machine reasoning. In *International Conference on Learning Representations*, 2018.
- [47] Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6700–6709, 2019.
- [48] Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2901–2910, 2017.
- [49] Justin Johnson, Bharath Hariharan, Laurens Van Der Maaten, Judy Hoffman, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. Inferring and executing programs for visual reasoning. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2989–2998, 2017.
- [50] Łukasz Kaiser and Ilya Sutskever. Neural gpus learn algorithms. *arXiv preprint arXiv:1511.08228*, 2015.
- [51] Vijay R Konda and John N Tsitsiklis. Actor-critic algorithms. In *Advances in neural information processing systems*, pp. 1008–1014. Citeseer, 2000.
- [52] George Konidaris, Scott Kuindersma, Roderic Grupen, and Andrew Barto. Robot learning from demonstration by constructing skill trees. *The International Journal of Robotics Research*, 31(3):360–375, 2012.
- [53] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123(1):32–73, 2017.
- [54] Brenden M Lake, Ruslan Salakhutdinov, and Joshua B Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015.
- [55] Brenden M Lake, Tomer D Ullman, Joshua B Tenenbaum, and Samuel J Gershman. Building machines that learn and think like people. *Behavioral and brain sciences*, 40, 2017.
- [56] Guohao Li, Xin Wang, and Wenwu Zhu. Perceptual visual reasoning with knowledge propagation. In *Proceedings of the 27th ACM International Conference on Multimedia*, MM ’19, pp. 530–538, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450368896. doi: 10.1145/3343031.3350922. URL <https://doi.org/10.1145/3343031.3350922>
- [57] Michael L Littman. Markov games as a framework for multi-agent reinforcement learning. In *Machine learning proceedings 1994*, pp. 157–163. Elsevier, 1994.
- [58] Yunchao Liu and Zheng Wu. Learning to describe scenes with programs. In *International Conference on Learning Representations*, 2019.
- [59] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. *arXiv preprint arXiv:2103.14030*, 2021.

- [60] Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. Hierarchical question-image co-attention for visual question answering. *arXiv preprint arXiv:1606.00061*, 2016.
- [61] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *arXiv preprint arXiv:1908.02265*, 2019.
- [62] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.
- [63] Robin Manhaeve, Sebastijan Dumancic, Angelika Kimmig, Thomas Demeester, and Luc De Raedt. Deepprolog: Neural probabilistic logic programming. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL <https://proceedings.neurips.cc/paper/2018/file/dc5d637ed5e62c36ecb73b654b05ba2a-Paper.pdf>
- [64] Jiayuan Mao, Chuang Gan, Pushmeet Kohli, Joshua B. Tenenbaum, and Jiajun Wu. The Neuro-Symbolic Concept Learner: Interpreting Scenes, Words, and Sentences From Natural Supervision. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=rJgMlhRctm>
- [65] Jiayuan Mao, Xiuming Zhang, Yikai Li, William T. Freeman, Joshua B. Tenenbaum, and Jiajun Wu. Program-guided image manipulators. *CoRR*, abs/1909.02116, 2019. URL <http://arxiv.org/abs/1909.02116>
- [66] Mingyuan Mao, Renrui Zhang, Honghui Zheng, Peng Gao, Teli Ma, Yan Peng, Errui Ding, and Shumin Han. Dual-stream network for visual recognition. *arXiv preprint arXiv:2105.14734*, 2021.
- [67] Saif M Mohammad. Sentiment analysis: Detecting valence, emotions, and other affectual states from text. In *Emotion measurement*, pp. 201–237. Elsevier, 2016.
- [68] Ronald Parr and Stuart Russell. Reinforcement learning with hierarchies of machines. *Advances in neural information processing systems*, pp. 1043–1049, 1998.
- [69] Dean A Pomerleau, Jay Gowdy, and Charles E Thorpe. Combining artificial neural networks and symbolic processing for autonomous robot guidance. *Engineering Applications of Artificial Intelligence*, 4(4): 279–285, 1991.
- [70] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- [71] Scott Reed and Nando De Freitas. Neural programmer-interpreters. *arXiv preprint arXiv:1511.06279*, 2015.
- [72] Karan Samel, Zelin Zhao, Binghong Chen, Kuan Wang, Robin Luo, and Le Song. How to design sample and computationally efficient vqa models. *arXiv preprint arXiv:2103.11537*, 2021.
- [73] Stefan Schaal et al. Learning from demonstration. *Advances in neural information processing systems*, pp. 1040–1046, 1997.
- [74] Kevin J Shih, Saurabh Singh, and Derek Hoiem. Where to look: Focus regions for visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4613–4621, 2016.
- [75] Richard Shin, Illia Polosukhin, and Dawn Song. Towards specification-directed program repair, 2018. URL <https://openreview.net/forum?id=B1iZRFkwz>.
- [76] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.
- [77] Sungryull Sohn, Junhyuk Oh, and Honglak Lee. Hierarchical reinforcement learning for zero-shot generalization with subtask dependencies. *arXiv preprint arXiv:1807.07665*, 2018.
- [78] Chunfeng Song, Yan Huang, Wanli Ouyang, and Liang Wang. Mask-guided contrastive attention model for person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [79] Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. Videobert: A joint model for video and language representation learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 7464–7473, 2019.

- [80] Shao-Hua Sun, Hyeonwoo Noh, Sriram Somasundaram, and Joseph Lim. Neural program synthesis from diverse demonstration videos. In *International Conference on Machine Learning*, pp. 4790–4799. PMLR, 2018.
- [81] Shao-Hua Sun, Te-Lin Wu, and Joseph J Lim. Program guided agent. In *International Conference on Learning Representations*, 2019.
- [82] Kai Sheng Tai, Richard Socher, and Christopher D Manning. Improved semantic representations from tree-structured long short-term memory networks. *arXiv preprint arXiv:1503.00075*, 2015.
- [83] Hao Tan and Mohit Bansal. Lxmert: Learning cross-modality encoder representations from transformers. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 5103–5114, 2019.
- [84] Yee Whye Teh, Victor Bapst, Wojciech Marian Czarnecki, John Quan, James Kirkpatrick, Raia Hadsell, Nicolas Heess, and Razvan Pascanu. Distral: Robust multitask reinforcement learning. *arXiv preprint arXiv:1707.04175*, 2017.
- [85] Yonglong Tian, Andrew Luo, Xingyuan Sun, Kevin Ellis, William T Freeman, Joshua B Tenenbaum, and Jiajun Wu. Learning to infer and execute 3d shape programs. *arXiv preprint arXiv:1901.02875*, 2019.
- [86] Dweep Trivedi, Jesse Zhang, Shao-Hua Sun, and Joseph J. Lim. Learning to synthesize programs as interpretable and generalizable policies, 2021.
- [87] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *arXiv preprint arXiv:1706.03762*, 2017.
- [88] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.
- [89] Abhinav Verma, Vijayaraghavan Murali, Rishabh Singh, Pushmeet Kohli, and Swarat Chaudhuri. Programmatically interpretable reinforcement learning. In *International Conference on Machine Learning*, pp. 5045–5054. PMLR, 2018.
- [90] Ricardo Vilalta and Youssef Drissi. A perspective view and survey of meta-learning. *Artificial intelligence review*, 18(2):77–95, 2002.
- [91] Huiyu Wang, Yukun Zhu, Hartwig Adam, Alan L. Yuille, and Liang-Chieh Chen. Max-deeplab: End-to-end panoptic segmentation with mask transformers. *CoRR*, abs/2012.00759, 2020. URL <https://arxiv.org/abs/2012.00759>
- [92] Aaron Wilson, Alan Fern, Soumya Ray, and Prasad Tadepalli. Multi-task reinforcement learning: a hierarchical bayesian approach. In *Proceedings of the 24th international conference on Machine learning*, pp. 1015–1022, 2007.
- [93] Jiajun Wu, Joshua B Tenenbaum, and Pushmeet Kohli. Neural scene de-rendering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 699–707, 2017.
- [94] Da Xiao, Jo-Yu Liao, and Xingyuan Yuan. Improving the universality and learnability of neural programmer-interpreters with combinator abstraction. *arXiv preprint arXiv:1802.02696*, 2018.
- [95] Wenbin Xu, Xudong Li, Wenda Xu, Liang Gong, Yixiang Huang, Zelin Zhao, Lujie Zhao, Binhao Chen, Haozhe Yang, Li Cao, et al. Human-robot interaction oriented human-in-the-loop real-time motion imitation on a humanoid tri-co robot. In *2018 3rd International Conference on Advanced Robotics and Mechatronics (ICARM)*, pp. 781–786. IEEE, 2018.
- [96] Wenbin Xu, Xudong Li, Liang Gong, Yixiang Huang, Zeyuan Zheng, Zelin Zhao, Lujie Zhao, Binhao Chen, Haozhe Yang, Li Cao, et al. Natural teaching for humanoid robot via human-in-the-loop scene-motion cross-modal perception. *Industrial Robot: the international journal of robotics research and application*, 2019.
- [97] Jianwei Yang, Jiayuan Mao, Jiajun Wu, Devi Parikh, David D Cox, Joshua B Tenenbaum, and Chuang Gan. Object-centric diagnosis of visual reasoning. *arXiv preprint arXiv:2012.11587*, 2020.
- [98] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Ruslan Salakhutdinov, and Quoc V Le. Xlnet: Generalized autoregressive pretraining for language understanding. *arXiv preprint arXiv:1906.08237*, 2019.

- [99] Kexin Yi, Jiajun Wu, Chuang Gan, Antonio Torralba, Pushmeet Kohli, and Joshua B Tenenbaum. Neural-symbolic vqa: Disentangling reasoning from vision and language understanding. *arXiv preprint arXiv:1810.02338*, 2018.
- [100] Yang You, Jing Li, Sashank Reddi, Jonathan Hseu, Sanjiv Kumar, Srinadh Bhojanapalli, Xiaodan Song, James Demmel, Kurt Keutzer, and Cho-Jui Hsieh. Large batch optimization for deep learning: Training bert in 76 minutes. *arXiv preprint arXiv:1904.00962*, 2019.
- [101] QingLong Zhang and Yubin Yang. Rest: An efficient transformer for visual recognition. *CoRR*, abs/2105.13677, 2021. URL <https://arxiv.org/abs/2105.13677>
- [102] Zelin Zhao, Gao Peng, Haoyu Wang, Hao-Shu Fang, Chengkun Li, and Cewu Lu. Estimating 6d pose from localizing designated surface keypoints. *arXiv preprint arXiv:1812.01387*, 2018.
- [103] Zelin Zhao, Chuang Gan, Jiajun Wu, Xiaoxiao Guo, and Joshua Tenenbaum. Augmenting policy learning with routines discovered from a demonstration. *arXiv preprint arXiv:2012.12469*, 2020.
- [104] Bolei Zhou, Yuandong Tian, Sainbayar Sukhbaatar, Arthur Szlam, and Rob Fergus. Simple baseline for visual question answering. *arXiv preprint arXiv:1512.02167*, 2015.