
Is a Modular Architecture Enough?

Sarthak Mittal^{1,2} Yoshua Bengio^{1,2} Guillaume Lajoie^{1,2}

Abstract

Inspired from human cognition, machine learning systems are now revealing advantages of sparser and more modular architectures. Recent works demonstrate that not only do some modular architectures generalize well, but they also lead to better out-of-distribution generalization, scaling properties, learning speed, and interpretability. A key intuition behind the success of such systems is that the data generating system for most real-world settings is considered to consist of sparsely interacting parts, promoting the use of similar inductive biases in the models. However, the field has been lacking in a rigorous quantitative assessment of such systems because these real-world data distributions are complex and unknown. Hence, we provide a thorough assessment of common modular architectures, through the lens of simple and known modular data distributions. We highlight the benefits of modularity and sparsity and reveal insights on the challenges faced while optimizing modular systems. We also propose evaluation metrics that highlight the regimes in which these benefits of modularity are substantial, as well as the sub-optimality of current end-to-end learned modular systems as opposed to their claimed potential.

1. Introduction

Deep learning research has an established history of drawing inspiration from neuroscience and cognitive science (Baars, 1997; Dehaene et al., 2017; Bengio, 2017; Goyal & Bengio, 2020). It is thus no surprise that modularity and attention have been leveraged, often together, in artificial networks in recent years (Bahdanau et al., 2015; Andreas et al., 2016; Hu et al., 2017; Vaswani et al., 2017; Kipf et al., 2018; Battaglia et al., 2018; Goyal et al., 2019; 2021), with impressive results. Although a number of recent results hinge on such modular architectures (Graves et al., 2014; Andreas et al.,

2016; Hu et al., 2017; Vaswani et al., 2017; Kipf et al., 2018; Santoro et al., 2018; Battaglia et al., 2018; Goyal et al., 2019; 2021; Locatello et al., 2020; Mittal et al., 2020; Madan et al., 2021; Ke et al., 2021), the abundance of tricks and proposed architectural modifications makes it challenging to parse real, usable architectural principles. It is also unclear whether the merits obtained by such Mixture-of-Experts (MoE) based systems are actually due to good specialization, as is often claimed, or due to other confounding factors.

In this work, we extend the analysis from (Rosenbaum et al., 2019; Maziarz et al., 2019; Cui & Jaech, 2020; Csordás et al., 2020) and propose a principled approach to evaluate, quantify, and analyse common ingredients of modular architectures. To do so, we develop a series of benchmarks and metrics aimed at probing the efficacy of a wide range of modular networks, where computation is factorized. This reveals valuable insights and helps identify not only where current approaches succeed but also *when and how they fail*. Given the recent increased interest in sparse modular systems (Rahaman et al., 2021; Fedus et al., 2021; Du et al., 2021; Mittal et al., 2021), we believe that this work will provide a test-bed for investigating the workings of such models and allow for research into inductive biases that can push such models to achieve good specialization. Through detailed experiments and evaluation metrics, we make the following observations and contributions:

- We develop benchmark tasks and metrics based on probabilistically selected rules to quantify two important phenomena in modular systems, the extent of *collapse* and *specialization*.
- We distill commonly used modularity inductive biases and systematically evaluate them through a series of models aimed at extracting commonly used architectural attributes (*Monolithic*, *Modular*, *Modular-op* and *GT-Modular* models).
- We find that specialization in modular systems leads to significant boosts in performance when there are many underlying rules within a task, but not so much with only few rules.
- We find standard modular systems to be often sub-optimal in both their capacity on focusing on the right information as well as in their ability to specialize, suggesting the need for additional inductive biases.

¹Mila ²Universite de Montreal. Correspondence to: Sarthak Mittal <sarthmit@gmail.com>.

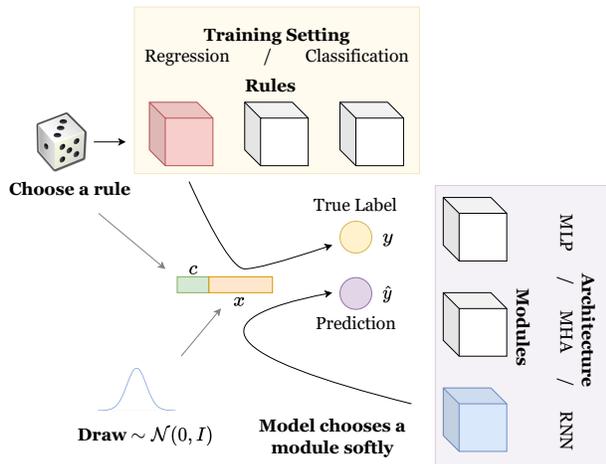


Figure 1. **Illustration of modularity evaluation framework.** Task configurations define the rules in the data-generating process while model parameters define the kind of model to be trained.

2. Notation / Terminology

In this paper, we study how a family of modular systems performs on a common set of tasks, prescribed by a synthetic data generating process which we call *rule-based data*. Below, we introduce the notation for key ingredients: (1) *rules* and how they form *tasks*, (2) *modules* and how they can take different *model architectures*, (3) *specialization* and how we evaluate models. We refer the reader to Figure 1 for an illustration of our setup.

Rules. To properly understand modular systems and analyze their benefits and shortcomings, we consider synthetic settings that allow fine-grained control over different aspects of task requirements. In particular, operations must be learned on the data-generating distribution illustrated in Equations 1-2, which we also refer to as *rules*.

$$\mathbf{c} \sim \text{Categorical}(\cdot) \quad \mathbf{x} \sim p_x(\cdot) \quad (1)$$

$$\mathbf{y} \mid \mathbf{x}, \mathbf{c} \sim p_y(\cdot \mid \mathbf{x}, \mathbf{c}). \quad (2)$$

Here, we define a rule as an expert of this distribution, i.e., rule r is defined as $p_y(\cdot \mid \mathbf{x}, \mathbf{c} = r)$ where \mathbf{c} is a categorical variable representing context, and \mathbf{x} is an input sequence. For example, consider $\mathbf{x} = (1, 2)$ and \mathbf{c} to select between addition and multiplication. Then, depending on \mathbf{c} , the correct output would be either $\mathbf{y} = 3$ or $\mathbf{y} = 2$. Systems will be trained to infer \mathbf{y} given \mathbf{c} and \mathbf{x} . We amend the above scheme to Multilayer Perceptron (MLP), Multi-Head Attention (MHA) and Recurrent Neural Network (RNN) based settings for concrete experimentation, the details of which are presented in Appendix A.

Tasks. A task is described by a set of rules (data-generating distribution) illustrated in Equations 1-2. Different sets of $\{p_y(\cdot \mid \mathbf{x}, \mathbf{c})\}_{\mathbf{c}}$ imply different tasks. For a given number

Model	Functional Form
<i>Monolithic</i>	$\hat{\mathbf{y}} = f(\mathbf{x}, \mathbf{c})$
<i>Modular</i>	$\hat{\mathbf{y}}_m, p_m = f_m(\mathbf{x}, \mathbf{c})$ $\hat{\mathbf{y}} = \sum_{m=1}^R p_m \hat{\mathbf{y}}_m$
<i>Modular-op</i>	$\hat{\mathbf{y}}_m = f_m(\mathbf{x}, \mathbf{c})$ $\mathbf{p} = g(\mathbf{c})$ $\hat{\mathbf{y}} = \sum_{m=1}^R p_m \hat{\mathbf{y}}_m$
<i>GT-Modular</i>	$\hat{\mathbf{y}}_m = f_m(\mathbf{x}, \mathbf{c})$ $\hat{\mathbf{y}} = \sum_{m=1}^R c_m \hat{\mathbf{y}}_m$

Table 1. **Functional Forms of Different Models.** Exact functional forms of the different models, given the data (\mathbf{x}, \mathbf{c}) . Depending on context, f and f_m are either MLP, MHA or RNN architectures.

of rules, we train models on multiple tasks to remove bias towards any particular task.

Modules. A modular system comprises a set of neural network modules, each of which can contribute to the overall output. One can see this through the functional form $\mathbf{y} = \sum_{m=1}^M p_m \mathbf{y}_m$, where \mathbf{y}_m denotes the output and p_m the activation of the m^{th} module. Details about the different modular systems are outlined in Appendix B.

From this point onwards, we exclusively use *rules* to refer to the specialized components in the data-generating process, and *modules* to refer to the experts that are learned by a modular system. Further, for ease of quantitative assessment, we always set the number of modules equal to the number of rules, except when evaluating monolithic models (with a single module). Modules can be implemented in three different architectures, as described next.

Models. We consider four different class of models, (a) *Monolithic*: one big model, (b) *Modular*: a system comprising of modules which are softly selected based on input, (c) *Modular-op*: similar to modular, but the selection depends only on the code \mathbf{c} which is privileged information, and (d) *GT-Modular*: a perfectly specializing system where \mathbf{c} directly guides selection of modules. The differences are outlined in Table 1 and additional details are discussed in Appendix B

Model Architectures. Model architectures describe the choice of architecture considered for each module of a modular system, or the single module in a monolithic system (MLP, MHA, or RNN). Importantly, the rules (or data generating distributions) are adapted to the model architecture, and we often refer to them as such (e.g. MLP based rules). Details about the data distributions and models considered in this work are provided in Appendices A and B respectively.

Perfect Specialization. When training modular systems on rule-based data, we would like the modules to specialize according to the rules in the data-generating distribution. Thus, there is an important need to quantify what constitutes perfect specialization of the system to the data. To allow for

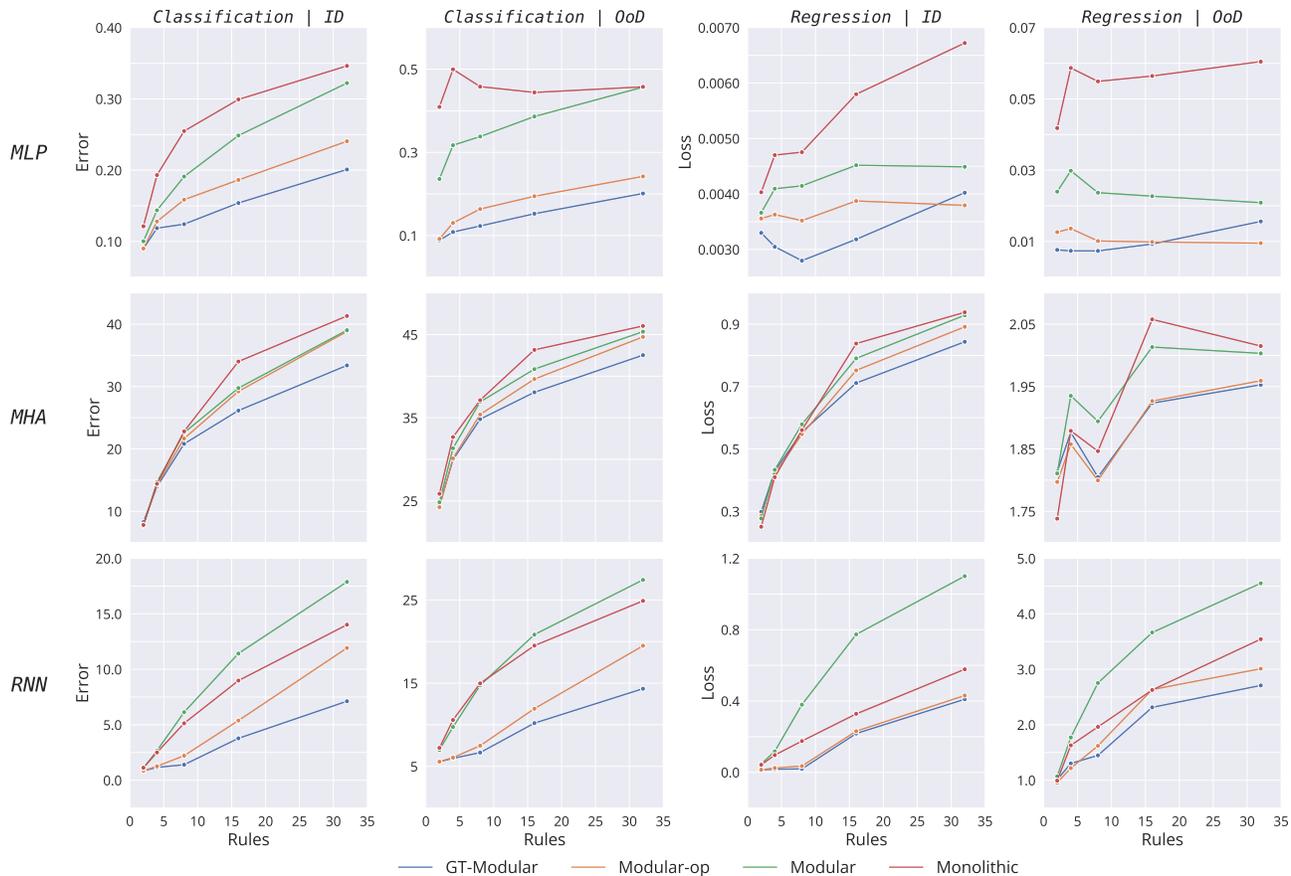


Figure 2. Performance Results. Performance of different models with MLP (top row), MHA (middle row) and RNN (bottom row) architectures against varying number of rules, evaluated both in-distribution and out-of-distribution. Modular systems generally outperform monolithic ones (*lower is better*) but a typical end-to-end trained modular system (green) is neither able to concentrate on the right information (compare with orange) nor able to get optimal specialization (compare with blue).

easier quantification, we always consider an equal number of modules and rules.

3. Metrics

To reliably evaluate modular systems, we propose metrics that not only gauge the performance benefits of such systems but also evaluate them across two important modalities: *collapse* and *specialization*, which we use to analyse the extent of resource allocation (in terms of parameters/modules) and specialization respectively of a modular system.

Performance. The first set of evaluation metrics are based on both in-distribution as well as out-of-distribution (OoD) performance. These metrics capture how well the different models perform on a wide variety of different tasks. For classification settings, we report the classification error while for regression settings, we report the loss. We highlight our performance based metrics in Figure 2, which are averaged over multiple seeds, tasks and model capacities. We refer the readers to Appendices C,D and G-I for details.

Collapse Metrics. We propose a set of metrics *Collapse-Avg* and *Collapse-Worst* that quantify the amount of collapse

suffered by a modular system. Collapse refers to the degree of under-utilization of the modules. An example of this is illustrated in Figure 4 in the Appendix. We consider the setting where all the data rules are equi-probable and the number of modules in the model are set to be the same as the number of data rules, to R .

Collapse-Avg. Let $p(m)$ be the marginal probability distribution of activation of module m . Then, $C_A = \frac{R}{R-1} \sum_{m=1}^R \max(0, \frac{1}{R} - p(m))$ captures the amount of under-utilization of all the modules of the system. A lower number is preferable for this metric, as a lower number demonstrates that all the modules are equally utilized.

Collapse-Worst. Given the same setting as above, $C_W = 1 - R \min_m p(m)$ captures the amount of under-utilization of the least used module of the system. Again, a low number is preferable here.

Specialization Metrics. To complement collapse metrics, we propose, (1) *Alignment*, (2) *Adaptation* and (3) *Inverse Mutual Information* to quantify the amount of specialization obtained by modular systems. These metrics are aimed at capturing how well the modules specialize to the rules, i.e.,

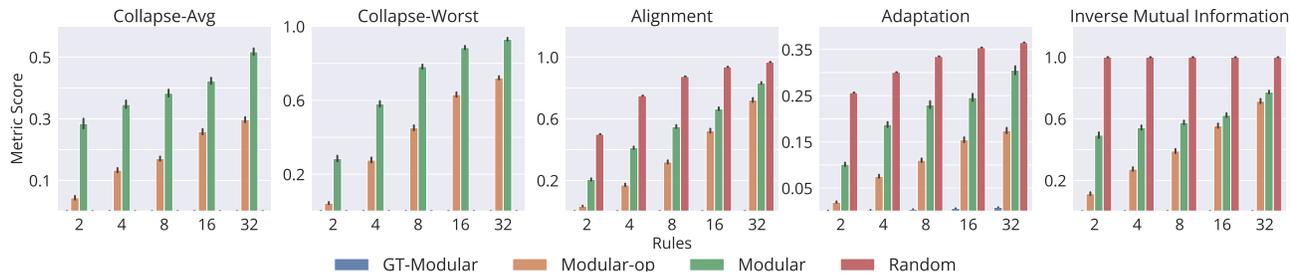


Figure 3. **Metrics against Increasing Number of Rules.** Evaluation of different modular systems through collapse (first two columns) and specialization (next three columns) (*lower is better*) under different number of rules.

whether different modules align to different rules.

Alignment. In a modular system, one can obtain the activation matrix \mathbf{A} , where \mathbf{A}_{rm} denotes $p(\text{module} = m \mid \text{rule} = r)$, i.e., the probability of activation of module m conditioned on rule r . Then, $s_d = \min_{\mathbf{P} \in \mathbf{S}_R} d(\mathbf{A}, \mathbf{P})$ defines the distance of the modular system to perfect specialization, where d denotes the normalized L_1 distance and \mathbf{S}_R denotes the space of permutation matrices over R objects. $s_d \rightarrow 0$ implies that each module specializes to a *unique* rule, signifying perfect specialization. An efficient computation of this metric is done through the Hungarian algorithm (Kuhn, 1955) and a lower score implies better specialization.

Inverse Mutual Information. Let the joint distribution $p(m, r)$ denote the activation probability of module m on rule r . Then, we define $S_{IMI} = 1 - \frac{1}{\log R} \mathbb{E}_{p(m, r)} \left[\log \frac{p(m, r)}{p(m) \times p(r)} \right]$, where a low inverse mutual information metric is preferable as it denotes that the modules are more specialized to the rules as opposed to multiple modules contributing to a single rule.

Adaptation. Let \mathcal{P} be a distribution over the R -dimensional simplex and $p(\cdot)$ the distribution over rules (not equiprobable in this metric) and $q(\cdot)$ the corresponding distribution obtained over modules, which is dependent on $p(\cdot)$. We then define $S_A = \mathbb{E}_{p \sim \mathcal{P}} \left[\sum_{i=1}^R |p(\hat{r}_i) - q(\hat{m}_i)| \right]$, where \hat{r}_i and \hat{m}_i are such that $p(\hat{r}_1) \leq \dots \leq p(\hat{r}_R)$ and $q(\hat{m}_1) \leq \dots \leq q(\hat{m}_R)$ and \mathcal{P} is a dirichlet distribution. This metric can be understood as the amount by which the modules adapt to changes in the rule distributions. A low adaptation score implies that the marginal distribution of the modules adapt well according to the distribution of the rules.

We refer the readers to Figure 3 which shows that the problems of collapse and specialization increase with the number of rules, and is often mitigated in the presence of privileged information (Modular-op). Note that uniform random activation patterns lead to low collapse metrics but high alignment, adaptation and inverse mutual information metrics, implying little collapse but poor specialization, as expected. On the contrary, *GT-Modular* systems lead to low collapse, alignment, adaptation and inverse mutual

information, denoting little collapse and good specialization, which is expected since specialization is given as oracle.

4. Conclusion and Discussion

We provide a benchmark suitable for the analysis of modular systems and provide metrics that not only evaluate them on in-distribution and out-of-distribution performance, but also on collapse and specialization. Through our large-scale analysis, we uncover many intriguing properties of modular systems and highlight potential issues that could lead to poor scaling properties of such systems.

Perfect Specialization. We discover that perfect specialization indeed helps in boosting performance both in-distribution and out-of-distribution, especially in the regime of many rules. On the contrary, monolithic systems often do comparatively when there are only a few rules.

End-to-End Trained Modular systems. While Modular systems outperform Monolithic ones, the margin of improvement is often small since these models do not discover perfect specialization. In fact, the problem of poor specialization and high collapse becomes worse with increasing number of rules, even when privileged contextual information is used explicitly, as in Modular-op.

In summary, through systematic and extensive experiments, this work shows that modularity, when supporting good and distributed specialization (i.e. little collapse), can outperform monolithic models both in and out of distribution testing. However, we also find that although perfectly specialized solutions are attainable by modular networks, end-to-end training does not recover them, often even with explicit information about task context (as in Modular-op). An important conclusion is that additional inductive biases are required to learn adequately specialized solutions. These could include regularization schemes (e.g. load-balancing (Fedus et al., 2021)) or optimization strategies (e.g. learning rate scheduling), among others, to promote specialization. We refer the reader to Appendix E for further discussion on these exciting prospects. We believe the framework proposed in this work is ideal to drive research into such inductive biases and a necessary stepping stone for applications of these designs at scale.

Acknowledgements

SM would like to acknowledge the support of UNIQUE and IVADO towards his research, as well as Calcul Québec and Compute Canada for providing the computing resources. YB and GL acknowledge the support from Canada CIFAR AI Chair Program, as well as Samsung Electronics Co., Ltd. GL acknowledges NSERC Discovery Grant [RGPIN-2018-04821].

References

- Andreas, J., Rohrbach, M., Darrell, T., and Klein, D. Neural module networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 39–48, 2016.
- Baars, B. J. In the theatre of consciousness. global workspace theory, a rigorous scientific theory of consciousness. *Journal of Consciousness Studies*, 4(4):292–309, 1997.
- Bahdanau, D., Cho, K., and Bengio, Y. Neural machine translation by jointly learning to align and translate. *ICLR’2015*, arXiv:1409.0473, 2015.
- Battaglia, P. W., Hamrick, J. B., Bapst, V., Sanchez-Gonzalez, A., Zambaldi, V., Malinowski, M., Tacchetti, A., Raposo, D., Santoro, A., Faulkner, R., et al. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*, 2018.
- Bengio, Y. The consciousness prior. *arXiv preprint arXiv:1709.08568*, 2017.
- Csordás, R., van Steenkiste, S., and Schmidhuber, J. Are neural nets modular? inspecting functional modularity through differentiable weight masks. *arXiv preprint arXiv:2010.02066*, 2020.
- Cui, L. and Jaech, A. Re-examining routing networks for multi-task learning. 2020.
- Dehaene, S., Lau, H., and Kouider, S. What is consciousness, and could machines have it? *Science*, 358(6362): 486–492, 2017.
- Du, N., Huang, Y., Dai, A. M., Tong, S., Lepikhin, D., Xu, Y., Krikun, M., Zhou, Y., Yu, A. W., Firat, O., et al. Glam: Efficient scaling of language models with mixture-of-experts. *arXiv preprint arXiv:2112.06905*, 2021.
- Fedus, W., Zoph, B., and Shazeer, N. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *arXiv preprint arXiv:2101.03961*, 2021.
- Goyal, A. and Bengio, Y. Inductive biases for deep learning of higher-level cognition. *arXiv preprint arXiv:2011.15091*, 2020.
- Goyal, A., Lamb, A., Hoffmann, J., Sodhani, S., Levine, S., Bengio, Y., and Schölkopf, B. Recurrent independent mechanisms. *arXiv preprint arXiv:1909.10893*, 2019.
- Goyal, A., Didolkar, A., Ke, N. R., Blundell, C., Beaudoin, P., Heess, N., Mozer, M., and Bengio, Y. Neural production systems. *arXiv preprint arXiv:2103.01937*, 2021.
- Graves, A., Wayne, G., and Danihelka, I. Neural Turing machines. *arXiv preprint arXiv:1410.5401*, 2014.
- Hu, R., Andreas, J., Rohrbach, M., Darrell, T., and Saenko, K. Learning to reason: End-to-end module networks for visual question answering. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 804–813, 2017.
- Ke, N. R., Didolkar, A. R., Mittal, S., Goyal, A., Lajoie, G., Bauer, S., Rezende, D. J., Mozer, M. C., Bengio, Y., and Pal, C. Systematic evaluation of causal discovery in visual model based reinforcement learning. 2021.
- Kipf, T., Fetaya, E., Wang, K.-C., Welling, M., and Zemel, R. Neural relational inference for interacting systems. *arXiv preprint arXiv:1802.04687*, 2018.
- Kuhn, H. W. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955.
- Locatello, F., Weissenborn, D., Unterthiner, T., Mahendran, A., Heigold, G., Uszkoreit, J., Dosovitskiy, A., and Kipf, T. Object-centric learning with slot attention. *arXiv preprint arXiv:2006.15055*, 2020.
- Madan, K., Ke, R. N., Goyal, A., Schölkopf, B. B., and Bengio, Y. Fast and slow learning of recurrent independent mechanisms. *arXiv preprint arXiv:2105.08710*, 2021.
- Masoudnia, S. and Ebrahimpour, R. Mixture of experts: a literature survey. *Artificial Intelligence Review*, 42(2): 275–293, 2014.
- Maziarz, K., Kokipoulou, E., Gesmundo, A., Sbaiz, L., Bartok, G., and Berent, J. Flexible multi-task networks by learning parameter allocation. *arXiv preprint arXiv:1910.04915*, 2019.
- Mittal, S., Lamb, A., Goyal, A., Voleti, V., Shanahan, M., Lajoie, G., Mozer, M., and Bengio, Y. Learning to combine top-down and bottom-up signals in recurrent neural networks with attention over modules. In *International Conference on Machine Learning*, pp. 6972–6986. PMLR, 2020.
- Mittal, S., Raparthy, S. C., Rish, I., Bengio, Y., and Lajoie, G. Compositional attention: Disentangling search and retrieval. *arXiv preprint arXiv:2110.09419*, 2021.

- Rahaman, N., Gondal, M. W., Joshi, S., Gehler, P., Bengio, Y., Locatello, F., and Schölkopf, B. Dynamic inference with neural interpreters. *Advances in Neural Information Processing Systems*, 34, 2021.
- Rosenbaum, C., Cases, I., Riemer, M., and Klinger, T. Routing networks and the challenges of modular and compositional computation. *arXiv preprint arXiv:1904.12774*, 2019.
- Santoro, A., Faulkner, R., Raposo, D., Rae, J., Chrzanowski, M., Weber, T., Wierstra, D., Vinyals, O., Pascanu, R., and Lillicrap, T. Relational recurrent neural networks. *arXiv preprint arXiv:1806.01822*, 2018.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. In *Advances in neural information processing systems*, pp. 5998–6008, 2017.
- Yuksel, S. E., Wilson, J. N., and Gader, P. D. Twenty years of mixture of experts. *IEEE transactions on neural networks and learning systems*, 23(8):1177–1193, 2012.

A. Data Generating Process

Since we aim to study modular systems through synthetic data, here we flesh out the data-generating processes operating based on the rules scheme described above (see Equations 1-2). We use a simple Mixture-of-Experts (MoE) (Yuksel et al., 2012; Masoudnia & Ebrahimpour, 2014) styled data-generating process, where we expect different modules to specialize to the different experts in the rules. It is important to note that this system is slightly different from the traditional flat MoE since the experts are more plug-and-play and can be composed to solve a particular problem. As an example, if we consider a mixture of recurrent systems, different tokens (time-points) in the input sequence can undergo computations according to different rules (e.g. a switching linear dynamical system), as opposed to the choice of expert being governed by the whole sequence.

We now look at more specific setups of the data-generating systems in consideration, the general template of which was outlined above. To do so, we explain the data-generating processes amenable to our three model architectures: MLP, MHA, and RNN. Additionally, each of the following tasks have two versions: regression, and classification. These are included to explore potential differences these distinct loss types may induce.

$$c \sim \mathcal{U}\{1, R\} \quad (3)$$

$$\mathbf{x}_1, \mathbf{x}_2 \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \mathbf{I}) \quad (4)$$

$$\mathbf{y} = \alpha_c \mathbf{x}_1 + \beta_c \mathbf{x}_2 \quad (5)$$

MLP. Here, we define the data scheme that is amenable for learning of modular MLP-based systems. In this synthetic data-generating scheme, a data sample consists of two independent numbers and a choice of rule being sampled from some distribution. Different rules lead to different linear combinations of the two numbers to give the output. That is, the choice of linear combination is dynamically instantiated based on the rule drawn. This is mathematically formulated in Equations 3-5, where α_c and β_c are the data parameters and \mathbf{y} denotes the label for the regression tasks and $\text{sign}(\mathbf{y})$ for the classification tasks.

Hence, the data comes from a MoE distribution where c denotes which linear combination governs the conditional distribution $p_y(\cdot | \mathbf{x}_1, \mathbf{x}_2, c)$. When training modular architectures on such data, one expects each module in the trained system to specialize according to a unique rule.

$$c_n \stackrel{\text{iid}}{\sim} \mathcal{U}\{1, R\} \quad (6)$$

$$\mathbf{q}_{nr}, \mathbf{q}'_{nr}, \mathbf{v}_{nr}, \mathbf{v}'_{nr} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \mathbf{I}) \quad (7)$$

$$s_n = \min_{i \neq n} d(\mathbf{q}_{nc_n}, \mathbf{q}_{ic_n}) \quad (8)$$

$$s'_n = \min_{i \neq n} d(\mathbf{q}'_{nc_n}, \mathbf{q}'_{ic_n}) \quad (9)$$

$$\mathbf{y}_n = \alpha_{c_n} \mathbf{v}_{s_n c_n} + \beta_{c_n} \mathbf{v}'_{s'_n c_n} \quad (10)$$

MHA. Now, we define the data scheme that is tuned for learning in modular MHA based systems. Essentially, a MHA module can be understood through a set of searches (query-key interactions), a set of corresponding retrievals (values) and then some computation of the retrieved values, as explained by Mittal et al. (2021). Accordingly, we design the data-generating distribution with the following properties: Each rule is composed of a different notion of search, retrieval and the final linear combination of the retrieved information respectively. We mathematically describe the process in Equations 6-10, where $n = 1, \dots, N$ and $r = 1, \dots, R$ with N as the sequence length and R the number of rules. We denote the tuple $(\mathbf{q}_{nr}, \mathbf{q}'_{nr}, \mathbf{v}_{nr}, \mathbf{v}'_{nr})$ as \mathbf{x}_n . Further, \mathbf{y}_n denotes the label for the regression tasks while for classification, we consider the categorical label to be $\text{sign}(\mathbf{y}_n)$.

Thus, we can see that c_n denotes the rule for the n^{th} token. This rule governs which two tokens are closest to the n^{th} token, demonstrated as s_n and s'_n . It also governs what features are retrieved from the searched tokens, which are $\mathbf{v}_{s_n c_n}$ and $\mathbf{v}_{s'_n c_n}$. These retrieved features then undergo a rule-dependent linear combination (on c_n). Here, too, when training a modular MHA architecture, we want each MHA module in the system to be able to specialize to a unique MHA rule in the data system.

$$c_n \stackrel{\text{iid}}{\sim} \mathcal{U}\{1, R\} \quad (11)$$

$$\mathbf{x}_n \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \mathbf{I}) \quad (12)$$

$$s_n = A_{c_n} s_{n-1} + B_{c_n} \mathbf{x}_n \quad (13)$$

$$\mathbf{y}_n = \mathbf{w}^T s_n \quad (14)$$

RNN. For recurrent systems, we define a rule as a kind of linear dynamical system, where one of multiple rules can be triggered at any time-point. Mathematically, this process can be defined through Equations 11-14, where $n = 1, \dots, N$, with N describing the sequence length. Each rule thus describes a different procedure for the update of the state s_t as well as the effect of the input \mathbf{x}_t to the state. Thus, we can see that c_n denotes the rule to be used at the n^{th} time-point. Further, \mathbf{y}_n denotes the label for the regression tasks while for classification, we consider the labels as $\text{sign}(\mathbf{y}_n)$.

Hence, in all settings, the data comes from a MoE distribution where c denotes the rule and governs the conditional $p_y(\cdot | \mathbf{x}, c)$. When training modular architectures on such data, one expects each module in the trained system to specialize according to a unique rule. Our aim is to use these synthetic rule-based data setting to study and analyse modular systems and understand whether end-to-end trained modular systems concentrate on the right information to

specialize based on, i.e. based on \mathbf{c} , whether they do learn perfect specialization and whether perfect specialization actually helps in these settings. To properly understand this, we detail the different kinds of models considered in Section B as well as the different metrics proposed in Section C to analyse trained systems.

For this work, we limit our analysis to infinite-data regime where each training iteration operates on a new data sample. Future work would perform similar analysis in the regime of limited data.

B. Models

Several works claim that end-to-end trained modular systems outperform their monolithic counterparts, especially in out-of-distribution settings. However, there is a lack of step-by-step analysis on the benefits of such systems and whether they actually specialize according to the data generating distribution or not. To perform an in-depth analysis, we consider four different types of models that allow for varying levels of specialization, which are: *Monolithic*, *Modular*, *Modular-op*, and *GT-Modular*. We give the formulations for each of these models below and then discuss the different analysis we can perform through them. We also illustrate these models in Table 1 and depending on the data-generating procedure described in Appendix A, f and f_m can be implemented as either MLP, MHA or RNN cells in this work.

Monolithic. A monolithic system is a big neural network that takes the entire data (\mathbf{x}, \mathbf{c}) as input and makes predictions $\hat{\mathbf{y}}$ based on it. There is no inductive bias about modularity or sparsity explicitly baked in the system and it is completely up to back-propagation to learn whatever functional form is needed to solve the task. An example of such a system is a traditional Multi-Head Attention (MHA) based system, eg. a Transformer.

Modular. A modular system is composed of a number of modules, each of which is a neural network of a given architectural type (MLP, MHA, or RNN). Each module m takes the data (\mathbf{x}, \mathbf{c}) as input and computes an output $\hat{\mathbf{y}}_m$ and a confidence score, normalized across modules into an activation probability p_m . The activation probability reflects the contribution of each module’s output to the final output $\hat{\mathbf{y}}$ of the system. Thus, there is an explicit baked-in inductive bias of modularity but it is still up to system-wide back-propagation to figure out the right specialization. An example of such a system is a mixture of MLPs or reusable RNNs, reusable across different time/positions.

Modular-op. A modular-op (for *operation only*) system is very similar to the modular system with just one small difference. Instead of the activation probability p_m of module m being a function of (\mathbf{x}, \mathbf{c}) , we instead make

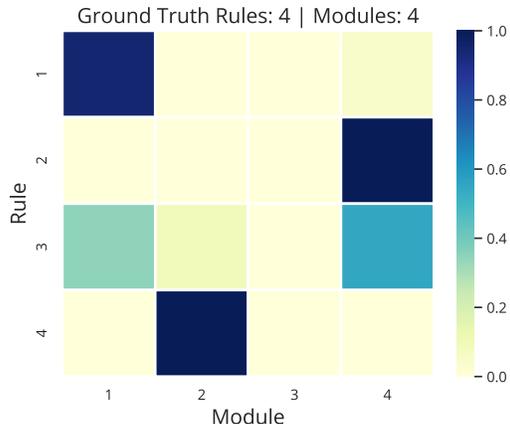


Figure 4. **Example of Collapse.** Entry (i, j) denotes the activation probability of module j on rule i . We see that Module 3 never activates, signifying collapse, while Module 4 covers two rules.

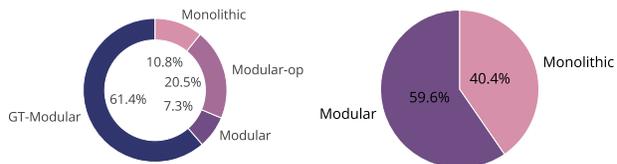


Figure 5. **Ranking Metric.** Spread indicates the number of experiments where the corresponding model did better. Plots include **Left:** All models, **Right:** Models trained without explicit rule-based module selection. Note that (a) explicit specialization (*GT-Modular*) helps, and (b) *Modular* systems outperform *Monolithic* but with small margin.

sure that the activation is decided only by the rule context \mathbf{c} . Hence, unlike modular systems, modular-op cannot be distracted by \mathbf{x} in figuring out specialization of different modules. Even though the operation required is explicitly provided, this model still needs to learn specialization through back-propagation.

GT-Modular. A GT-modular system (for *ground truth*) serves as an oracle benchmark, i.e., a modular system that specializes perfectly. In particular, the activation probability p'_m s of modules are just set according to \mathbf{c} , which is the indicator present in the data (\mathbf{x}, \mathbf{c}) . Thus, this is a perfectly specializing system that chooses different modules sparsely and perfectly according to the different data rules.

Given enough capacity, we can see that there is a hierarchy of models based on the functions they can implement, with $GT\text{-Modular} \subseteq Modular\text{-op} \subseteq Modular \subseteq Monolithic$. Put differently, models from *Monolithic* to *GT-Modular* increasingly incorporate the inductive biases for modularity and sparsity. This is proved in Appendix F by inspecting the function classes implemented by these models.

In what follows, we want to analyse the benefits of having

simple end-to-end trained modular systems as opposed to monolithic ones. This can be understood through a comparison of various performance based metrics between *Monolithic* and *Modular* models, explained in the next section. This will allow us to answer if a modular architecture is always better for various distinct rule-based data generating systems. For instance, a comparison between the *Modular* and *Modular-op* models will show whether the standard modular systems are able to focus on the right information and ignore the distractors in driving specialization. To study this, we will look at performance as well as collapse and specialization metrics between these class of models. A comparison between *GT-Modular* and *Modular-op* will show the benefits of having a sparse activation pattern with proper resource allocation of modules as opposed to an end-to-end learned specialization on the right information (without distractors).

Finally, we note that *GT-Modular* is a modular system which obtains perfect specialization. Through this model, we aim to analyse whether perfect specialization is in-fact important and if so, how far are typical modular systems from obtaining similar performance and specialization through end-to-end training. We now describe the metrics used for these evaluations.

C. Metrics

To reliably evaluate modular systems, we propose a suite of metrics that not only gauge the performance benefits of such systems but also evaluate them across two important modalities: *collapse* and *specialization*, which we use to analyse the extent of resource allocation (in terms of parameters/modules) and specialization respectively of a modular system.

Performance. The first set of evaluation metrics are based on performance of the models in both in-distribution as well as out-of-distribution (OoD) settings. These metrics capture how well the different models perform on a wide variety of different tasks. For classification settings, we report the classification error while for regression settings, we report the loss.

In-Distribution. This refers to the in-distribution performance, evaluated by looking at both the final performance as well as convergence speeds of the different models.

Out-of-Distribution. This refers to the OoD performance of different models. We consider very simple forms of OoD generalization: either (a) change in distribution of \mathbf{x} by increasing variance, or (b) different sequence lengths, wherever the possibility presents (eg. in MHA and RNN).

Collapse Metrics. We propose a set of metrics *Collapse-Avg* and *Collapse-Worst* that quantify the amount of collapse

suffered by a modular system. Collapse refers to the degree of under-utilization of the modules. An example of this is illustrated in Figure 4, where we can see that Module 3 is never used. We consider the setting where all the data rules are equi-probable and the number of modules in the model are set to be the same as the number of data rules, to R . High collapse thus refers to under-utilization of resource (parameters) provided to the model, illustrating that certain modules are never being used and concurrently meaning that certain modules are being utilized for multiple rules.

$$C_A = \frac{R}{R-1} \sum_{m=1}^R \max\left(0, \frac{1}{R} - p(m)\right) \quad (15)$$

Collapse-Avg. Given the data-setting with R equi-probable rules, and hence R modules in the model, we let $p(m)$ be the marginal probability distribution of activation of module m . Then, we define the *Collapse-Avg* metric C_A as in Equation 15, where $\frac{R}{R-1}$ is for normalization. This metric captures the amount of under-utilization of all the modules of the system. A lower number is preferable for this metric, as a lower number demonstrates that all the modules are equally utilized.

$$C_W = 1 - R \min_m p(m) \quad (16)$$

Collapse-Worst. Given the same data and model setting as above, the *Collapse-Worst* metric C_W is defined as in Equation 16. This metric captures the amount of under-utilization of the least used module of the system. Again, a low number is preferable as it signifies that even the least used module is decently utilized by the model.

Specialization Metrics. To complement collapse metrics, we also propose a set of metrics, (1) *Alignment*, (2) *Adaptation* and (3) *Inverse Mutual Information* to quantify the amount of specialization obtained by the modular systems. We again consider the setting of equi-probable rules and the same number of modules and rules R . These metrics are aimed at capturing how well the modules specialize to the rules, that is, whether different modules stick to different rules (good specialization) or whether all modules contribute almost equally to all rules (poor specialization).

$$s_d = \min_{\mathbf{P} \in \mathcal{S}_R} d(\mathbf{A}, \mathbf{P}) \quad (17)$$

Alignment. Given a modular system trained on rule-based data with R rules and modules, one can obtain the activation matrix \mathbf{A} , where $\mathbf{A}_{r,m}$ denotes $p(\text{module} = m \mid \text{rule} = r)$, that is, the probability of activation of module m conditioned on rule r . Further, given a distance metric $d(\cdot, \cdot)$ over

the space of matrices, perfect specialization can be quantified through Equation 17, where \mathbf{S}_R denotes the space of permutation matrices over R objects. We consider $d(\cdot, \cdot)$ as a normalized L_1 distance. The score s_d demonstrates the distance between the activation matrix A and its closest permutation matrix, with distances computed according to the metric $d(\cdot, \cdot)$. Note that $s_d \rightarrow 0$ implies that each module specializes to a *unique* rule, thereby signifying perfect specialization. Since the space of permutation matrices \mathbf{S}_R grows exponentially at the rate of $\Theta(R!)$, computing s_d naively soon becomes intractable. However, we use the Hungarian algorithm (Kuhn, 1955) to compute it in polynomial time. This metric shows how close the learned modular system is to a perfectly specializing one, where a low score implies better specialization.

$$S_{IMI} = 1 - \frac{1}{\log R} \mathbb{E}_{p(m,r)} \left[\log \frac{p(m,r)}{p(m) \times p(r)} \right] \quad (18)$$

Inverse Mutual Information. Given R as the number of rules and modules and let the joint distribution $p(m, r)$ denote the activation probability of module m on rule r , the *Inverse Mutual Information* metric S_{IMI} is defined as in Equation 18. A low inverse mutual information metric is preferable as it denotes that the modules are more specialized to the rules as opposed to multiple modules contributing to a single rule.

$$S_A = \mathbb{E}_{p \sim \mathcal{P}} \left[\sum_{i=1}^R \left| p(\hat{r}_i) - q(\hat{m}_i) \right| \right] \quad (19)$$

Adaptation. Let R be the number of rules and modules and \mathcal{P} a distribution over the R -dimensional simplex. Further, let $p(\cdot)$ be the distribution over rules (not equi-probable in this metric) and $q(\cdot)$ the corresponding distribution obtained over the modules. Note that the distribution $q(\cdot)$ is dependent on $p(\cdot)$. Given these distributions, we define the *Adaptation* metric S_A in Equation 19, where \hat{r}_i and \hat{m}_i are such that $p(\hat{r}_1) \leq p(\hat{r}_2) \leq \dots \leq p(\hat{r}_R)$ and $q(\hat{m}_1) \leq q(\hat{m}_2) \leq \dots \leq q(\hat{m}_R)$ and \mathcal{P} is a dirichlet distribution.

This metric can be understood as the amount by which the modules adapt (signified through the distribution $q(\cdot)$) to changes in the rule distributions (which are $p(\cdot)$ sampled from \mathcal{P}). The matching between the rule and module is obtained through a simple sort as defined above. A low adaptation score implies that the marginal distribution of the modules adapt well according to the distribution of the rules. That is, when a rule is weakly present in the data, there exists a module which weakly contributes in the corresponding output, averaged over multiple different rule distributions.

To understand these metrics, note that uniform random activation patterns for the modules lead to low collapse

metrics but high alignment, adaptation and inverse mutual information metrics, implying little collapse but poor specialization, as expected. On the other hand, *GT-Modular* systems necessarily lead to low collapse metrics as well as low alignment, adaptation and inverse mutual information, denoting little collapse and good specialization, which is expected since specialization is given as oracle.

D. Experiments

We are now ready to report experiments on the models outlined in Section B with associated data generation processes described in Appendix A. For each level of modularity (i.e. Monolithic, Modular, Modular-op, GT-Modular), we analyse models learning over five different number of rules, ranging from few (2) to many (32), five different model capacities (number of parameters) and two different training settings, i.e. regression and binary classification. To remove any biases towards particular task parameters (e.g. α_c, β_c in Equation 5), we randomly select new rules to create five different tasks per setting and, train five seeds per task. In essence, we train $\sim 20,000$ models¹ to properly analyse the benefits of modularity, the level of specialization obtained by end-to-end trained systems, the impact of number of rules and the impact of model capacity.

Performance. We refer the readers to Figure 5 for a compressed overview on the performance of various models. We see that GT-Modular system wins most of the times (*left*), indicating the benefits of perfect specialization. We also see that between standard end-to-end trained Modular and Monolithic systems, the former outperforms but not by a huge gap. Together, these two pie charts indicate that current end-to-end trained modular systems do not achieve good specialization and are thus sub-optimal by a substantial margin.

We then look at the specific architectural choices (MLP, MHA and RNN cells for functions f and f_m in Table 1) and analyse their performance and trends across increasing number of rules. Figure 2 shows that while there are concrete benefits of a perfectly specializing system (GT-Modular) or even models that know what information to drive specialization from (Modular-op), typical end-to-end trained Modular systems are quite sub-optimal and not able to realize these benefits, especially with increasing number of rules which is where we see substantial benefits of good specialization (contrast Modular vs GT-Modular and Modular-op). Moreover, while such end-to-end Modular systems do generally outperform the Monolithic ones, it is often only by a small margin.

We also see the training pattern of different models averaged over all other settings, with the average containing error

¹All models are trained on single V100 GPUs, each taking a few hours.

Is a Modular Architecture Enough?

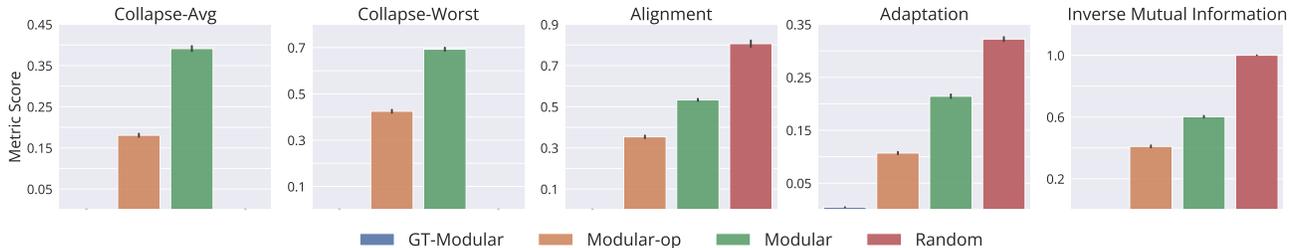


Figure 6. Metrics for Different Models. While end-to-end training of activation decisions leads to reduced collapse (first two columns) and better specialization (next three columns) (*lower is better*) than random activations, it is still far from a perfectly specializing system. This signifies that the models are not able to learn good specialization and actually suffer from increased collapse when learned solely through back-propagation.

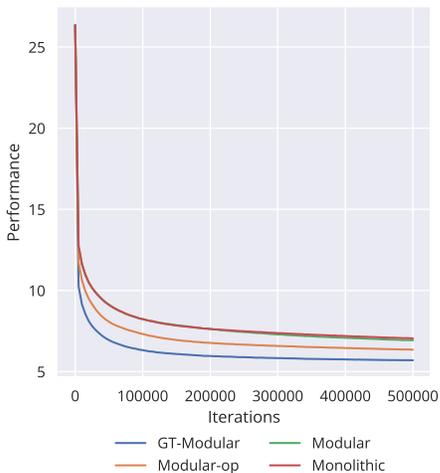


Figure 7. Training Curve. Averaged over different model architectures, training settings, model capacities and number of rules. We see that end-to-end trained modular systems are still far from the benefits of perfect specialization.

for classification and loss for regression, in Figure 7. We can see that good specialization not only leads to better performances but also faster training.

Collapse. We evaluate all the models on the two collapse metrics outlined in Section C. Figure 3 shows the two collapse metrics, *Collapse-Avg* and *Collapse-Worst*, for different models against varying number of rules, averaged over the different model architectures (MLP, MHA and RNN), training settings (Classification and Regression), model capacities, tasks and seeds. First, we notice that a Random activation baseline and the GT-Modular system do not have any collapse, which is expected. Next, we notice that both Modular and Modular-op suffer from the problems of collapse and this problem becomes worse with increasing number of rules. Figure 6 further shows similar information averaged over the number of rules too, highlighting that Modular-op has less collapse than Modular

in general. However, we still see that the problem of collapse is significant whenever back-propagation is tasked with finding the right activation patterns, especially in the regime of large number of rules. This clearly indicates the need for investigation into different forms of regularizations to alleviate some of the collapse problems.

Specialization. Next, we evaluate through the proposed specialization metrics in Section C whether the end-to-end trained modular systems actually specialize according to the data-generating distribution. Figure 3 shows the three specialization metrics, *Alignment*, *Adaptation* and *Inverse Mutual Information*, for different models against varying number of rules, again averaged over different model architectures, training settings, model capacities, tasks and seeds. As expected, we see that the Random activation baseline has poor specialization (high metrics) while the GT-Modular system has very good specialization. We further see that end-to-end trained Modular systems as well as Modular-op suffer from sub-optimal specialization, as indicated by the high metrics. As with collapse, we again see that it becomes harder to reach optimal specialization with increasing number of rules. Figure 6 shows that while Modular-op has marginally better specialization than standard Modular systems, they are indeed quite sub-optimal when compared to a perfectly specializing system, i.e. GT-Modular.

We refer the readers to Appendix G, H and I for training details as well as additional experiments regarding the effect of model sizes for MLP, MHA and RNN architectures respectively.

E. Future Work

We perform large-scale analysis on a variety of modular systems through simple *rule-based* data generating distributions. Through our analysis, we uncover several interesting insights into the regimes where modularity, sparsity and perfect specialization helps and how sub-optimal standard modular systems are in terms of collapse and specialization.

We believe that this is a first step towards better benchmarking and understanding of modular systems. However, there are still a number of important and interesting directions that have not been explored in this work. We discuss some of these important future directions here.

Stochasticity. In Appendix A, we see that we consider deterministic formulations of $p_y(\cdot | \mathbf{x}, \mathbf{c})$ and there is no labeling noise. One direction of exploration is to extend the settings considered here to noisy domains and investigate whether similar analyses still hold.

Hard Attention. In this work, we only considered simple soft-attention based activation decisions in end-to-end trained modular systems described in Section B. An interesting future work involves benchmarking hard-attention based modular systems to explore whether they perform or specialize better, and also if the problem of collapse is exacerbated in such models.

Finite-Data Regime. Since this is the first work that provides such an analysis in this field, we decided to stick to the simplest setting of infinite-data regime to limit the effects of overfitting. We believe that a useful future work would be to consider a low-data regime to see whether the inductive biases of modularity and sparsity lead to even better generalization when there is only little data to learn from.

Complex Rules. While we consider the simplest setting for each rule, one could consider more complex distributions where \mathbf{x} is also conditional on the rule \mathbf{c} and where the labeling function, denoted by \mathbf{y} in Appendix A, is a complex non-linear function instead of a simple linear function. Analysis on trends between modular and monolithic models with increasing complexity of rules would not only lead to better understanding of such systems but also bridge it closer to real-world settings.

Harder OoD Settings. Modular systems are often shown to lead to better OoD generalization. In this work, we considered the simplest possible OoD settings that were often heavily correlated with in-distribution performances. Future work should investigate more complex OoD settings where either the support of \mathbf{x} is dis-joint between in-distribution and OoD, or there are combinatorial computations required to obtain labels \mathbf{y} , such that certain rule permutations are with-held in training and used for evaluations.

Better Inductive Biases. A very important next step is to discover and investigate various inductive biases and regularization procedures that bridge the gap between the current modular systems and the perfectly specializing systems. Our benchmark provides the perfect opportunity that allows for analysis into the levels of specialization of different models.

The above points highlight only some of the immediate future works that would paint a richer and more intricate

picture of what current modular systems are capable of, and what are the benefits we *can* obtain through perfect specialization.

Apart from extensions to the rule and module settings, we believe that it would also be important to investigate further into the quantitative evaluation metrics that we consider, in particular -

Generalizing to non equi-probable rules. Our current metrics on *collapse* and *specialization* rely on the need for equi-probable rules in the data generating distribution. It would be important to extend this to a more general setting where certain rules could be present more than the others.

Different number of rules and modules. For ease of quantitative evaluation, we only consider a well-specified system where the number of rules and the number of modules are kept the same. An important next step is to formulate the *collapse* and *specialization* metrics to work in the settings where the number of modules could be more (or less) than the number of rules present in the data.

While investigating all the above possibilities is surely exciting, we believe that our work and setting provides the test-bed to allow for extensions and analysis into all the laid-out possibilities. This would not only allow for a more thorough understanding of modular systems but also lead to investigations into inductive biases that benefit such systems on various real-world settings.

F. Proof of Model Relations

To prove: Given enough representational capacity, we need to show that $GT\text{-Modular} \subseteq Modular\text{-op} \subseteq Modular \subseteq Monolithic$.

We prove this step-by-step by using the functional forms of different models as described in Table 1.

Claim: $GT\text{-Modular} \subseteq Modular\text{-op}$

Proof. Given the formulation of a *GT-Modular* system as

$$\hat{\mathbf{y}}_m = f_m(\mathbf{x}, \mathbf{c}) \quad (20)$$

$$\hat{\mathbf{y}} = \sum_{m=1}^R c_m \hat{\mathbf{y}}_m \quad (21)$$

and of the *Modular-op* system as

$$\hat{\mathbf{y}}_m = f'_m(\mathbf{x}, \mathbf{c}) \quad (22)$$

$$\mathbf{p} = g(\mathbf{c}) \quad (23)$$

$$\hat{\mathbf{y}} = \sum_{m=1}^R p_m \hat{\mathbf{y}}_m \quad (24)$$

We see that we obtain *GT-Modular* from *Modular-op* simply by setting $g(\cdot)$ as the identity function and $f_m = f'_m$

Claim: $Modular\text{-}op \subseteq Modular$

Proof. The formulation of a *Modular-op* system is given as

$$\hat{y}_m = f_m(\mathbf{x}, \mathbf{c}) \quad (25)$$

$$\mathbf{p} = g(\mathbf{c}) \quad (26)$$

$$\hat{y} = \sum_{m=1}^R p_m \hat{y}_m \quad (27)$$

and that of the *Modular* system as

$$\hat{y}_m, p_m = f'_m(\mathbf{x}, \mathbf{c}) \quad (28)$$

$$\hat{y} = \sum_{m=1}^R p_m \hat{y}_m \quad (29)$$

We can describe $f'_m(\mathbf{x}, \mathbf{c}) = (f'_{m_1}(\mathbf{x}, \mathbf{c}), f'_{m_2}(\mathbf{x}, \mathbf{c}))$ for the *Modular* system. Setting $f'_{m_1}(\mathbf{x}, \mathbf{c}) = f_m(\mathbf{x}, \mathbf{c})$ and $f'_{m_2}(\mathbf{x}, \mathbf{c}) = g(\mathbf{c})$ for all \mathbf{x}, \mathbf{c} , we recover *Modular-op* from the *Modular* system.

Claim: $Modular \subseteq Monolithic$

Proof. Let the formulation of a *Modular* system be

$$\hat{y}_m, p_m = f_m(\mathbf{x}, \mathbf{c}) \quad (30)$$

$$\hat{y} = \sum_{m=1}^R p_m \hat{y}_m \quad (31)$$

where $f_m(\mathbf{x}, \mathbf{c}) = (f_{m_1}(\mathbf{x}, \mathbf{c}), f_{m_2}(\mathbf{x}, \mathbf{c}))$. The formulation of the *Monolithic* system is

$$\hat{y} = f(\mathbf{x}, \mathbf{c}) \quad (32)$$

We can recover the *Modular* system from *Monolithic* simply by setting $f(\mathbf{x}, \mathbf{c}) = \sum_{m=1}^R f_{m_2}(\mathbf{x}, \mathbf{c}) \cdot f_{m_1}(\mathbf{x}, \mathbf{c})$.

Given enough capacity of all the systems, all the functional assignment are possible. This completes the proof, and shows that each choice provides an additional inductive bias by potentially restricting the functional class.

G. MLP

We provide detailed results of our MLP based experiments highlighting the effects of the training setting (Regression or Classification), the number of rules (ranging from 2 to 32) and the different model capacities. In these set of experiments, we use the MLP version of the data generating process (as highlighted in Appendix A) and consider the models (highlighted in Section B) with f and f_m modeled using MLP architectures.

Task and Model Setups. We follow the task setup as described in Appendix A. We consider 1-dimensional data

samples for \mathbf{x}_1 and \mathbf{x}_2 and for the task parameters, we sample $\alpha_c, \beta_c \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$. Further, for the OoD generalization setup, we instead sample input from a different distribution, i.e., $\mathbf{x}_1, \mathbf{x}_2 \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 2\mathbf{I})$.

For the models, we consider a shared non-linear digit encoder that separately encodes \mathbf{x}_1 and \mathbf{x}_2 and an operational encoder that encodes \mathbf{c} . The encoded inputs are then fed together to either a monolithic MLP system or to each MLP module of the modular system. We control for the number of parameters such that all the systems roughly share the same number of parameters. Having obtained an output from the system, we then use a shared decoder to make the prediction.

We train all the models for 100,000 iterations with a batch-size of 256 and the Adam optimizer with learning rate of 0.0001. For the classification tasks, we consider binary cross entropy loss, while for regression we consider the l_1 loss.

Classification. We first look at the results on the binary classification based MLP tasks. For reporting performance metrics, we consider all model capacities as well as number of rules while for collapse and specialization metrics, we consider the smallest, mid-size and largest models and report over the different number of rules.

Performance. For ease of readability, we first provide a snapshot of the results through rankings in Figure 8. The rankings are based on the votes obtained by the different models. Given a task, averaged over the five training seeds, a vote is given to the model that performs the best. This provides a quick view of the number of times each model outperformed the rest.

Next, we refer the readers to Figure 9 for the in-distribution and Figure 10 for the out-of-distribution performance of the various models across both different model capacities as well as different number of rules.

Collapse-Avg. For each rule setting, we report the *Collapse-Avg* metric score of the different models and the three different model capacities in Figure 11.

Collapse-Worst. For each rule setting, we report the *Collapse-Worst* metric score of the different models and the three different model capacities in Figure 12.

Inverse Mutual Information. For each rule setting, we report the *Inverse Mutual Information* metric score of the different models and the three different model capacities in Figure 13.

Adaptation. For each rule setting, we report the *Adaptation* metric score of the different models and the three different model capacities in Figure 14.

Is a Modular Architecture Enough?

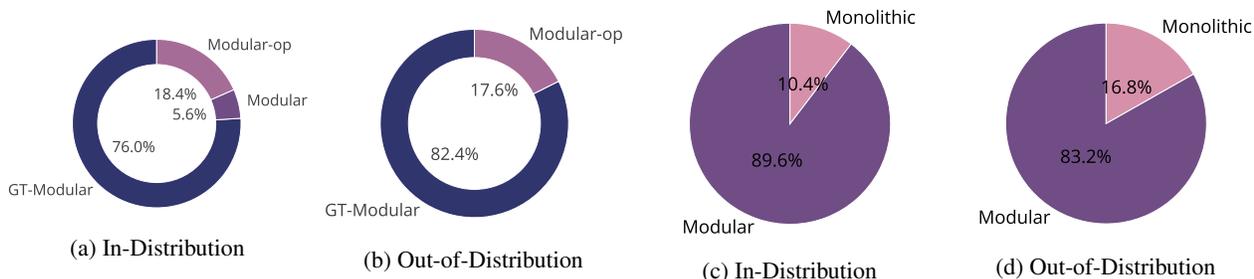


Figure 8. Ranking Metric for MLP-Classification. Each pie chart shows the number of times a model wins the competition (*higher is better*), which means outperforms the other models on a single task. Ranking is based on all models with in-distribution ranking in (a) and out-of-distribution ranking in (b). On the contrary, rankings are based only on completely end-to-end trained Modular and Monolithic systems with in-distribution ranking in (c) and out-of-distribution ranking in (d).

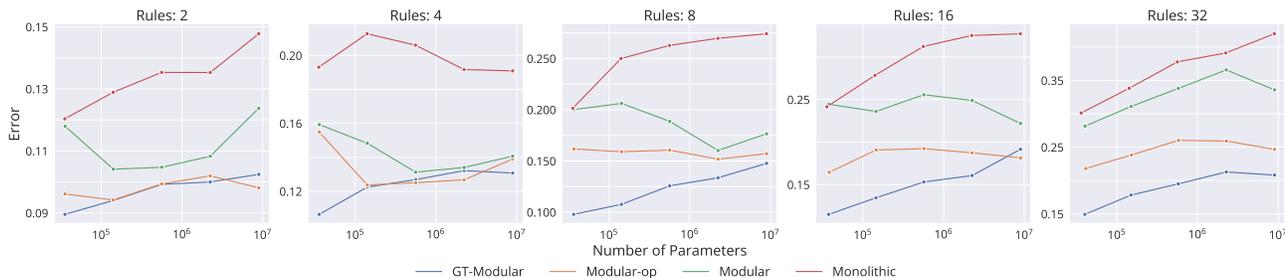


Figure 9. In-Distribution Performance of MLP-Classification Models. Performance (*lower is better*) of different models of varying capacities trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

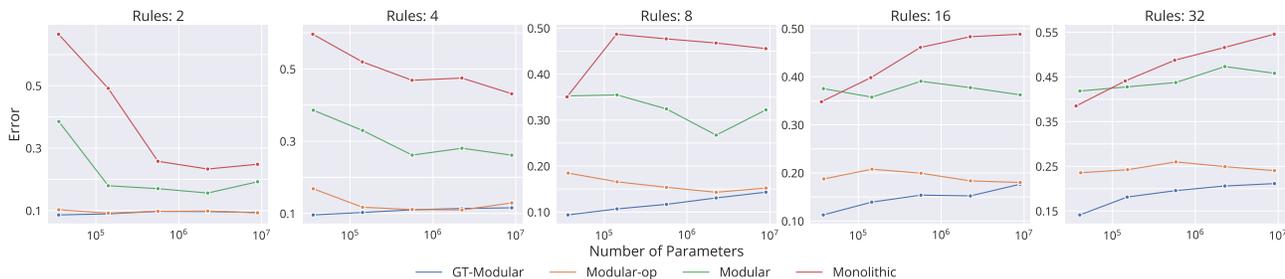


Figure 10. Out-of-Distribution Performance of MLP-Classification Models. Performance (*lower is better*) of different models of varying capacities trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

Alignment. For each rule setting, we report the *Alignment* metric score of the different models and the three different model capacities in Figure 15.

Regression. Next, we look at the results on the regression based MLP tasks. For reporting performance metrics, we consider all model capacities as well as number of rules while for collapse and specialization metrics, we consider the smallest, mid-size and largest models and report over the different number of rules.

Performance. For ease of readability, we first provide a snapshot of the results through rankings in Figure 16. The

rankings are based on the votes obtained by the different models. Given a task, averaged over the five training seeds, a vote is given to the model that performs the best. This provides a quick view of the number of times each model outperformed the rest.

Next, we refer the readers to Figure 17 for the in-distribution and Figure 18 for the out-of-distribution performance of the various models across both different model capacities as well as different number of rules.

Collapse-Avg. For each rule setting, we report the *Collapse-Avg* metric score of the different models and the three differ-

Is a Modular Architecture Enough?

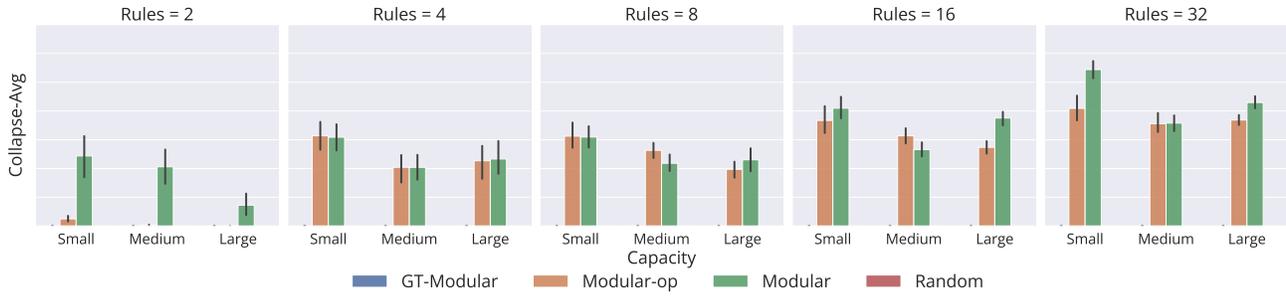


Figure 11. **Collapse-Avg Metric for MLP-Classification Models.** Highlights the amount of collapse (*lower is better*) suffered by different models of varying capacities trained across different number of rules. Each bar on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

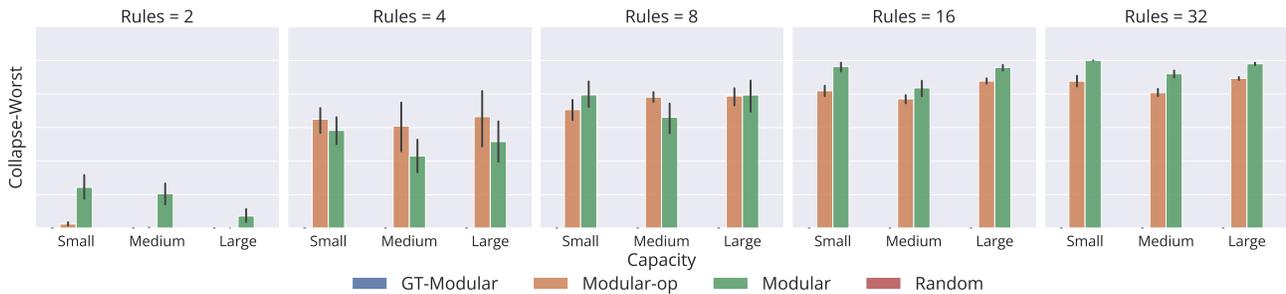


Figure 12. **Collapse-Worst Metric for MLP-Classification Models.** Highlights the amount of collapse (*lower is better*) suffered by different models of varying capacities trained across different number of rules. Each bar on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

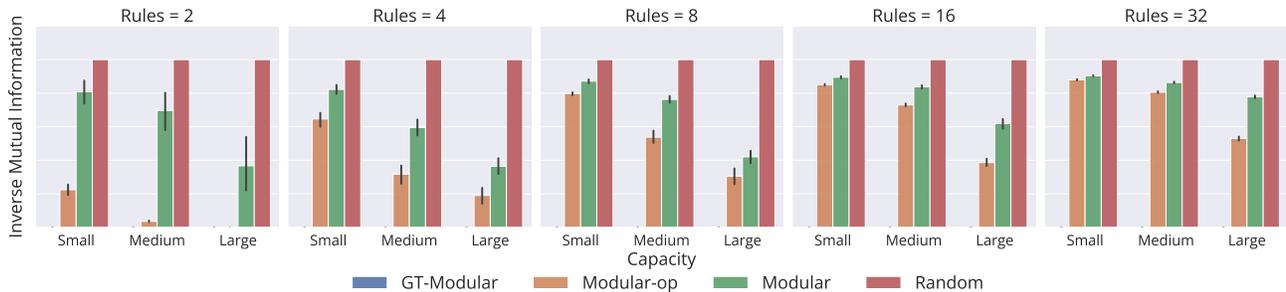


Figure 13. **Inverse Mutual Information Metric for MLP-Classification Models.** Highlights how poor the specialization (*lower is better*) is of different models of varying capacities trained across different number of rules. Each bar on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

ent model capacities in Figure 19.

Collapse-Worst. For each rule setting, we report the *Collapse-Worst* metric score of the different models and the three different model capacities in Figure 20.

Inverse Mutual Information. For each rule setting, we report the *Inverse Mutual Information* metric score of the different models and the three different model capacities in Figure 21.

Adaptation. For each rule setting, we report the *Adaptation*

metric score of the different models and the three different model capacities in Figure 22.

Alignment. For each rule setting, we report the *Alignment* metric score of the different models and the three different model capacities in Figure 23.

H. MHA

We provide detailed results of our MHA based experiments highlighting the effects of the training setting (Regression

Is a Modular Architecture Enough?

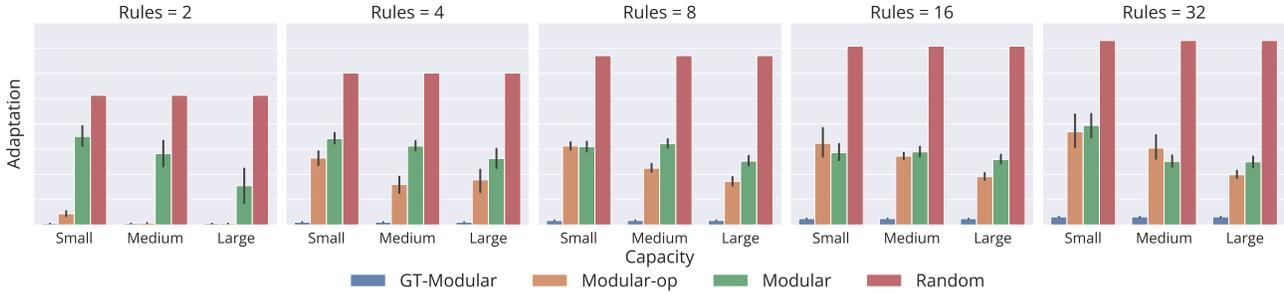


Figure 14. **Adaptation Metric for MLP-Classification Models.** Highlights how poor the specialization (*lower is better*) is of different models of varying capacities trained across different number of rules. Each bar on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

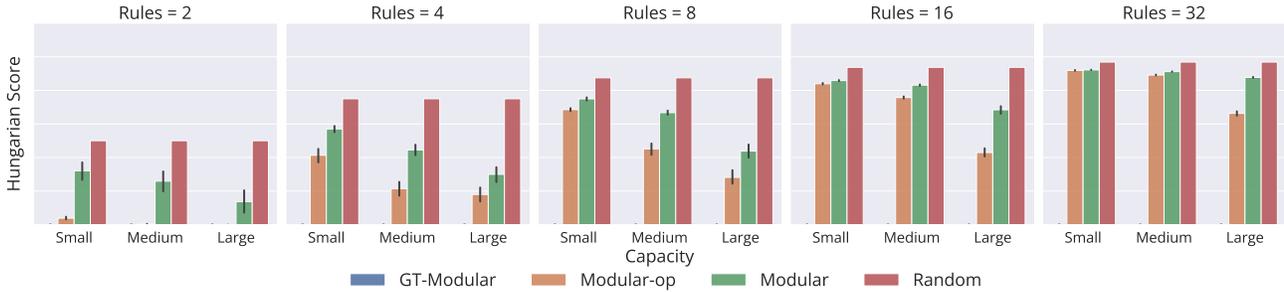


Figure 15. **Alignment Metric for MLP-Classification Models.** Highlights how poor the specialization (*lower is better*) is of different models of varying capacities trained across different number of rules. Each bar on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

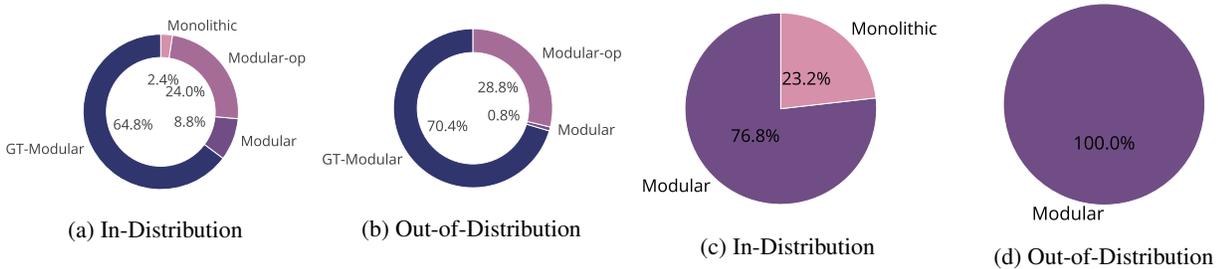


Figure 16. **Ranking Metric for MLP-Regression.** Each pie chart shows the number of times a model wins the competition (*higher is better*), which means outperforms the other models on a single task. Ranking is based on all models with in-distribution ranking in (a) and out-of-distribution ranking in (b). On the contrary, rankings are based only on completely end-to-end trained Modular and Monolithic systems with in-distribution ranking in (c) and out-of-distribution ranking in (d).

or Classification), the number of rules (ranging from 2 to 32) and the different model capacities. In these set of experiments, we use the MHA version of the data generating process (as highlighted in Appendix A) and consider the models (highlighted in Section B) with f and f_m modeled using MHA architectures.

Task and Model Setups. We follow the task setup as described in Appendix A. We consider the sequence length as 10 for training and the input \mathbf{v}_{nr} and \mathbf{v}'_{nr} to be 1-dimensional per token per rule. We consider two different

notions of search $d(\cdot, \cdot)$, outlined as

- Search-Version 1: Here \mathbf{q}_{nr} and \mathbf{q}'_{nr} are 1-dimensional per token per rule. The search rule is defined as $d(a, b) = |a - b|$.
- Search-Version 2: Here \mathbf{q}_{nr} and \mathbf{q}'_{nr} are 2-dimensional per token per rule and are instead sampled independently from a 2-dimensional hyper-sphere. The search rule is defined as $d(a, b) = a^T b$.

Is a Modular Architecture Enough?

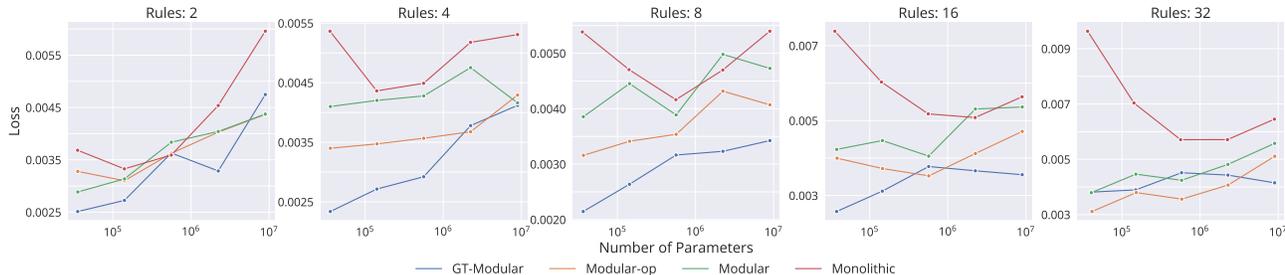


Figure 17. **In-Distribution Performance of MLP-Regression Models.** Performance (*lower is better*) of different models of varying capacities trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

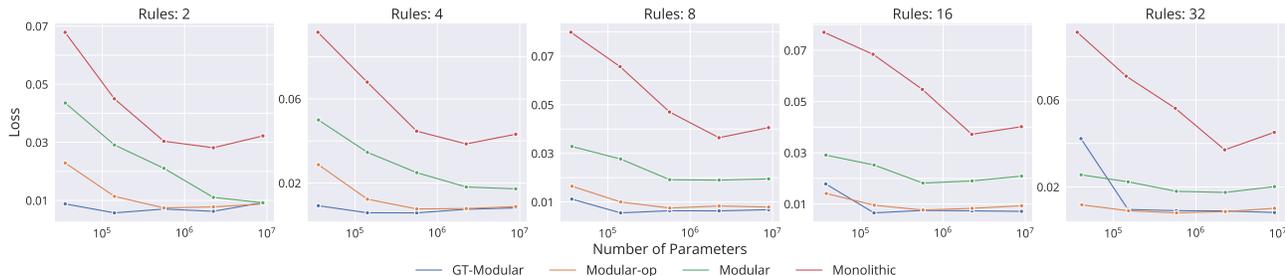


Figure 18. **Out-of-Distribution Performance of MLP-Regression Models.** Performance (*lower is better*) of different models of varying capacities trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

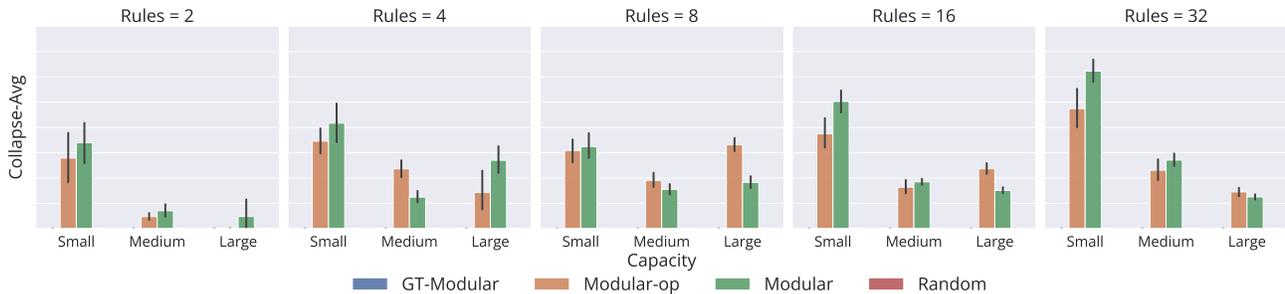


Figure 19. **Collapse-Avg Metric for MLP-Regression Models.** Highlights the amount of collapse (*lower is better*) suffered by different models of varying capacities trained across different number of rules. Each bar on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

For the task parameters, we sample $\alpha_c, \beta_c \stackrel{iid}{\sim} \mathcal{N}(0, 1)$. Further, for the OoD generalization setup, we instead sample input from a different distribution, i.e., we use $\mathcal{N}(0, 2I)$ instead of standard normal and in case of sampling from hypersphere, we instead consider one with double the radius. We also test with different sequence lengths, specifically 3, 5, 10, 20 and 30.

For the models, we consider a shared non-linear encoder that encodes each $(\mathbf{x}_n, \mathbf{c}_n)$ independently. The encoded inputs are then fed to either a monolithic MHA system or to each MHA module of the modular system. We control for the

number of parameters and the number of heads such that all the systems roughly share the same number of parameters. For modular systems, we consider the number of heads per module to be 2 which is the number required to solving the task. On the contrary, we give similar capacity to monolithic systems by having the number of heads as $2 \times (\text{Number of Rules})$. Having obtained an output from the system, we then use a shared decoder to make the prediction.

We train all the models for 500,000 iterations with a batch-size of 256 and the Adam optimizer with learning rate of 0.0001. For the classification tasks, we consider binary

Is a Modular Architecture Enough?

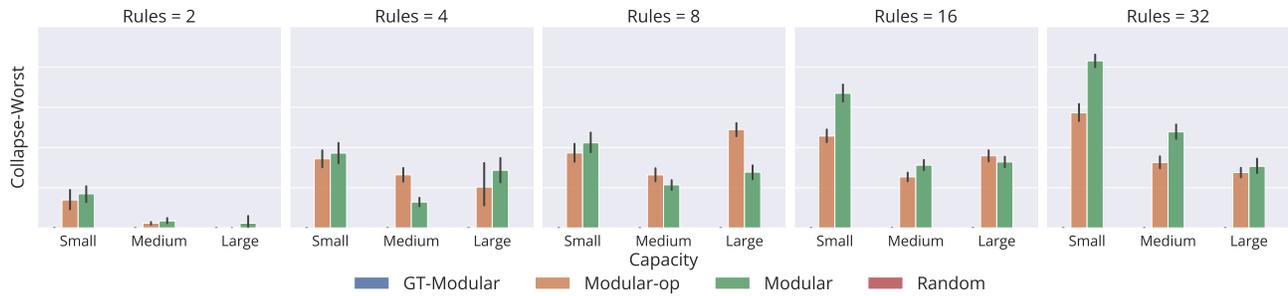


Figure 20. Collapse-Worst Metric for MLP-Regression Models. Highlights the amount of collapse (*lower is better*) suffered by different models of varying capacities trained across different number of rules. Each bar on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

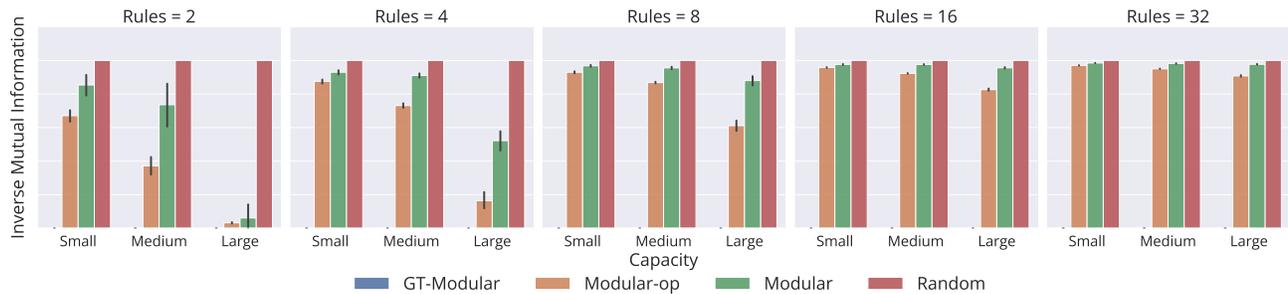


Figure 21. Inverse Mutual Information Metric for MLP-Regression Models. Highlights how poor the specialization (*lower is better*) is of different models of varying capacities trained across different number of rules. Each bar on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

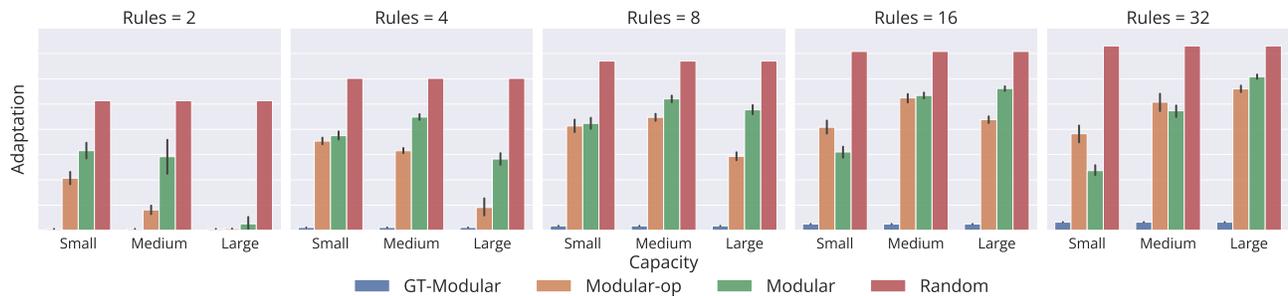


Figure 22. Adaptation Metric for MLP-Regression Models. Highlights how poor the specialization (*lower is better*) is of different models of varying capacities trained across different number of rules. Each bar on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

cross entropy loss, while for regression we consider the l_1 loss.

Classification. We first look at the results on the binary classification based MHA tasks. For reporting performance metrics, we consider all model capacities, all number of rules as well as the two search versions while for collapse and specialization metrics, we consider the smallest, mid-size and largest models and report over the different number of rules as well the search versions.

Performance. For ease of readability, we first provide a snapshot of the results through rankings in Figure 24. The rankings are based on the votes obtained by the different models. Given a task, averaged over the five training seeds, a vote is given to the model that performs the best. This provides a quick view of the number of times each model outperformed the rest.

We refer the readers to Figure 25 for the in-distribution performance and Figure 26 for the most extreme out-of-distribution performance (where we use the largest sequence

Is a Modular Architecture Enough?

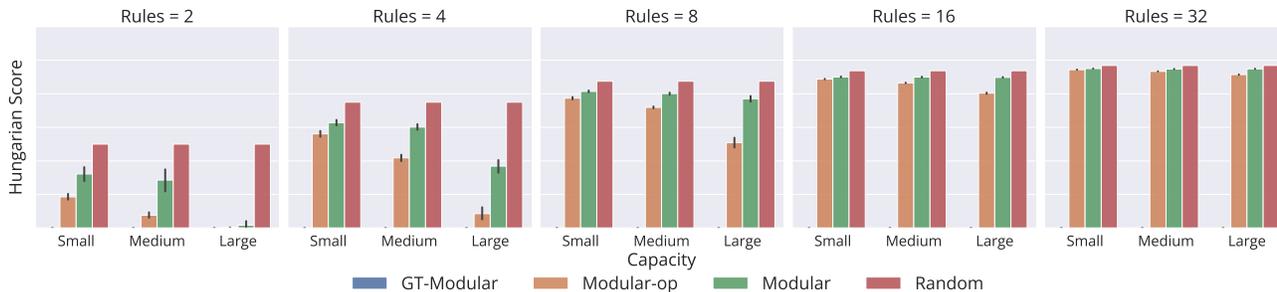


Figure 23. **Alignment Metric for MLP-Regression Models.** Highlights how poor the specialization (*lower is better*) is of different models of varying capacities trained across different number of rules. Each bar on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

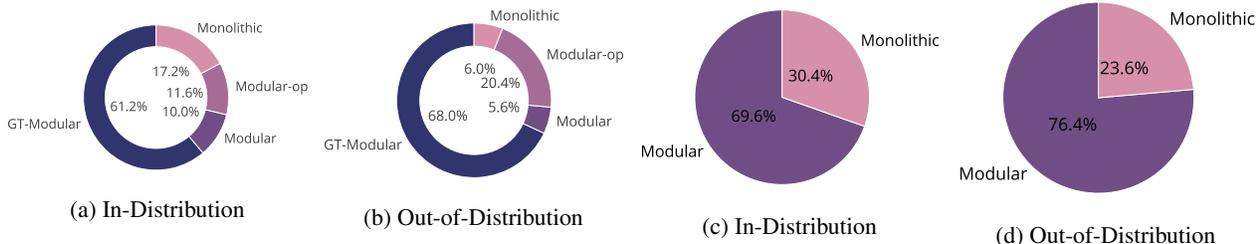


Figure 24. **Ranking Metric for MHA-Classification.** Each pie chart shows the number of times a model wins the competition (*higher is better*), which means outperforms the other models on a single task. Ranking is based on all models with in-distribution ranking in (a) and out-of-distribution ranking in (b). On the contrary, rankings are based only on completely end-to-end trained Modular and Monolithic systems with in-distribution ranking in (c) and out-of-distribution ranking in (d). For OoD, we only consider the extreme setting with largest sequence length and change in distribution of individual tokens.

length (30) and also change the distribution from which individual tokens are sampled) of the various models across both different model capacities, search functions as well as different number of rules. We also refer the reader to Figures 56 - 59 for the out-of-distribution case where only the sequence length is varied in the data and to Figures 60 - 64 where both the sequence length is potentially altered and also the distribution from which individual data points are sampled.

Collapse-Avg. For each rule setting and search version, we report the *Collapse-Avg* metric score of the different models and the three different model capacities in Figure 27.

Collapse-Worst. For each rule setting and search version, we report the *Collapse-Worst* metric score of the different models and the three different model capacities in Figure 28.

Inverse Mutual Information. For each rule setting and search version, we report the *Inverse Mutual Information* metric score of the different models and the three different model capacities in Figure 29.

Adaptation. For each rule setting and search version, we report the *Adaptation* metric score of the different models

and the three different model capacities in Figure 30.

Alignment. For each rule setting and search version, we report the *Alignment* metric score of the different models and the three different model capacities in Figure 31.

Regression. Next, we look at the results on the regression based MHA tasks. For reporting performance metrics, we consider all model capacities, all number of rules as well as the two search versions while for collapse and specialization metrics, we consider the smallest, mid-size and largest models and report over the different number of rules as well the search versions.

Performance. For ease of readability, we first provide a snapshot of the results through rankings in Figure 32. The rankings are based on the votes obtained by the different models. Given a task, averaged over the five training seeds, a vote is given to the model that performs the best. This provides a quick view of the number of times each model outperformed the rest.

We refer the readers to Figure 33 for the in-distribution performance and Figure 34 for the most extreme out-of-distribution performance (where we use the largest sequence length (30) and also change the distribution from which

Is a Modular Architecture Enough?

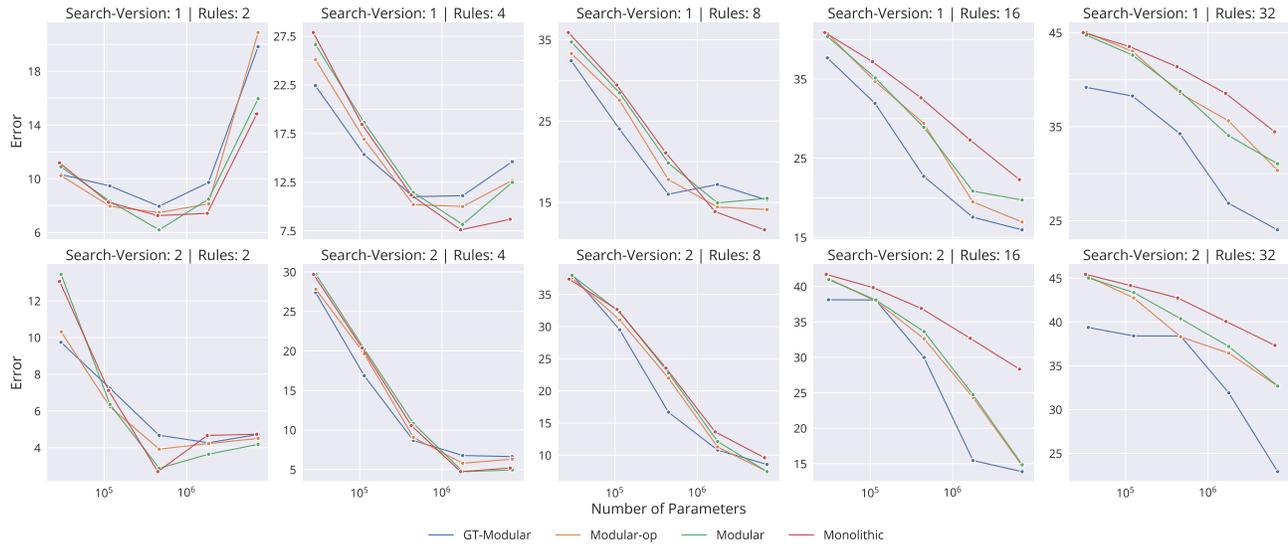


Figure 25. In-Distribution Performance of MHA-Classification Models. Performance (*lower is better*) of different models of varying capacities, different search versions in data and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

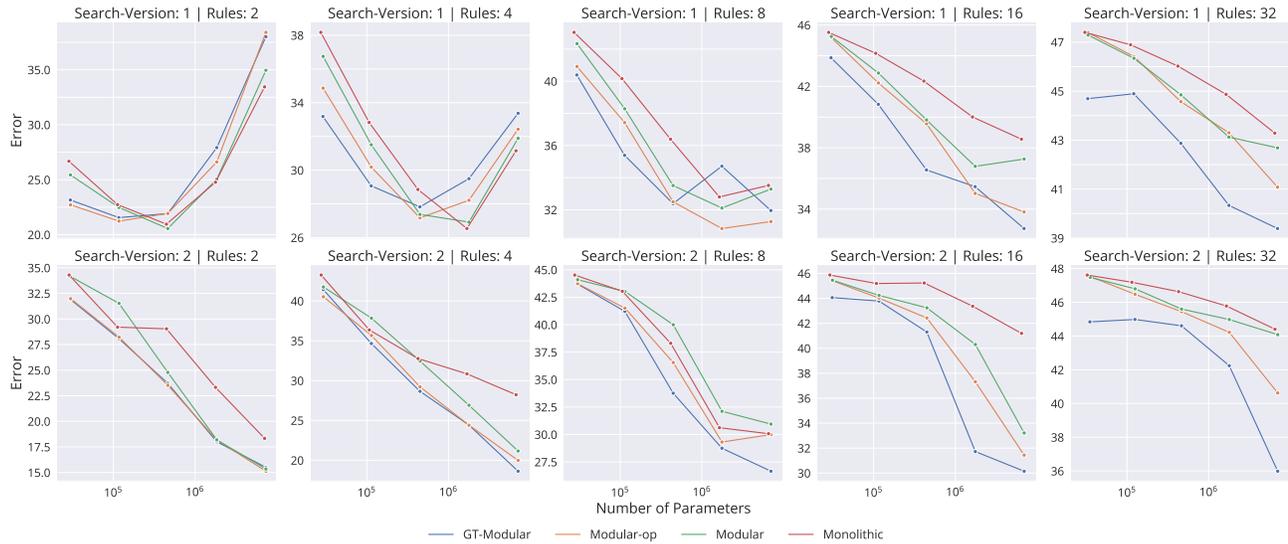


Figure 26. Out-of-Distribution (Sequence Length: 30 - Individual Token Sampling: Altered) Performance on MHA-Classification Models. Performance (*lower is better*) of different models of varying capacities, different search versions in data and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

Is a Modular Architecture Enough?



Figure 27. Collapse-Avg Metric for MHA-Classification Models. Highlights the amount of collapse (*lower is better*) suffered by different models of varying capacities trained across different number of rules. Each bar on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

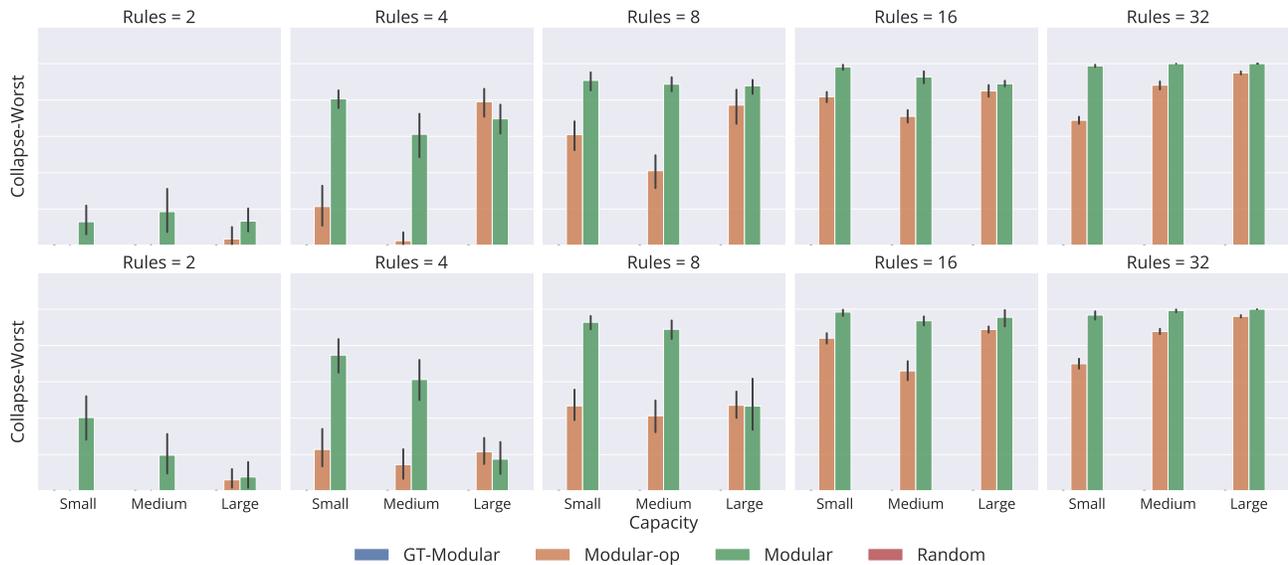


Figure 28. Collapse-Worst Metric for MHA-Classification Models. Highlights the amount of collapse (*lower is better*) suffered by different models of varying capacities trained across different number of rules. Each bar on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

individual tokens are sampled) of the various models across both different model capacities, search functions as well as different number of rules. We also refer the reader to Figures 65 - 68 for the out-of-distribution case where only the sequence length is varied in the data and to Figures 69 - 73 where both the sequence length is potentially altered and also the distribution from which individual data points are sampled.

Collapse-Avg. For each rule setting and search version, we report the *Collapse-Avg* metric score of the different models and the three different model capacities in Figure 35.

Collapse-Worst. For each rule setting and search version, we report the *Collapse-Worst* metric score of the different models and the three different model capacities in Figure 36.

Is a Modular Architecture Enough?

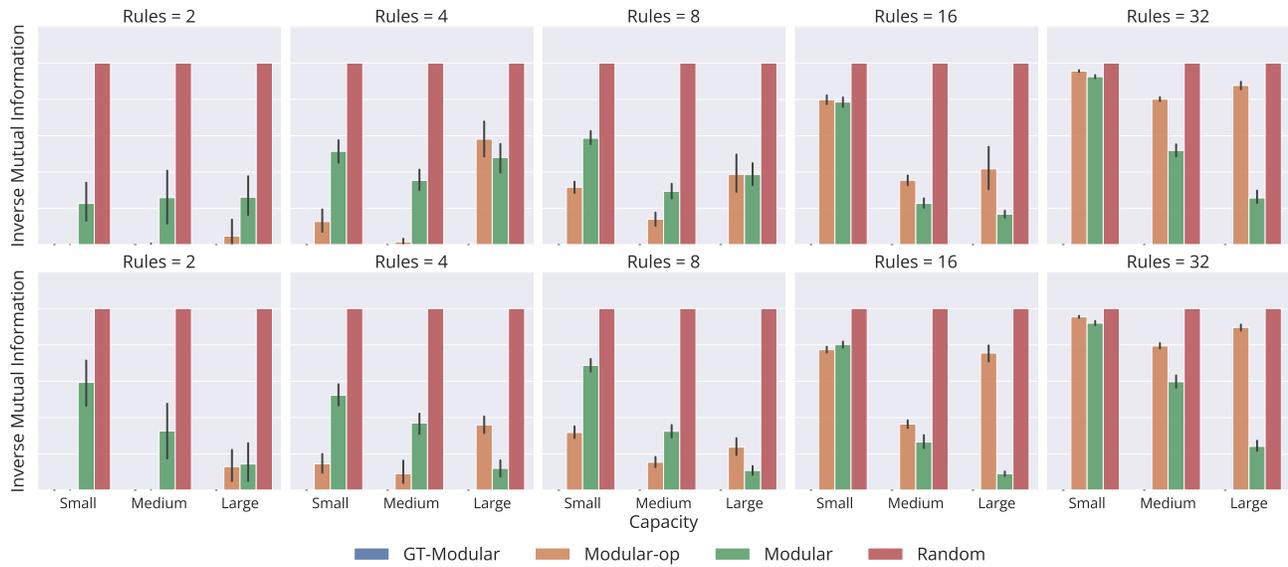


Figure 29. **Inverse Mutual Information Metric for MHA-Classification Models.** Highlights how poor the specialization (*lower is better*) is of different models of varying capacities trained across different number of rules. Each bar on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.



Figure 30. **Adaptation Metric for MHA-Classification Models.** Highlights how poor the specialization (*lower is better*) is of different models of varying capacities trained across different number of rules. Each bar on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

Inverse Mutual Information. For each rule setting and search version, we report the *Inverse Mutual Information* metric score of the different models and the three different model capacities in Figure 37.

Adaptation. For each rule setting and search version, we report the *Adaptation* metric score of the different models and the three different model capacities in Figure 38.

Alignment. For each rule setting and search version, we report the *Alignment* metric score of the different models and the three different model capacities in Figure 39.

I. RNN

We provide detailed results of our RNN based experiments highlighting the effects of the training setting (Regression

Is a Modular Architecture Enough?



Figure 31. **Alignment Metric for MHA-Classification Models.** Highlights how poor the specialization (*lower is better*) is of different models of varying capacities trained across different number of rules. Each bar on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

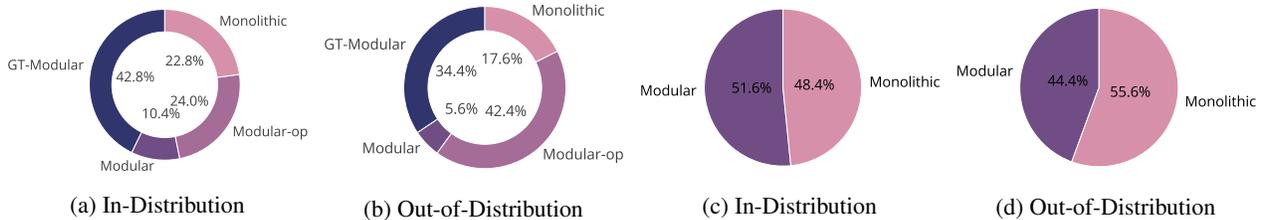


Figure 32. **Ranking Metric for MHA-Regression.** Each pie chart shows the number of times a model wins the competition (*higher is better*), which means outperforms the other models on a single task. Ranking is based on all models with in-distribution ranking in (a) and out-of-distribution ranking in (b). On the contrary, rankings are based only on completely end-to-end trained Modular and Monolithic systems with in-distribution ranking in (c) and out-of-distribution ranking in (d). For OoD, we only consider the extreme setting with largest sequence length and change in distribution of individual tokens.

or Classification), the number of rules (ranging from 2 to 32) and the different model capacities. In these set of experiments, we use the RNN version of the data generating process (as highlighted in Appendix A) and consider the models (highlighted in Section B) with f and f_m modeled using simple RNN architectures.

Task and Model Setups. We follow the task setup as described in Appendix A. We consider the sequence length as 10 for training and the input \mathbf{x}_n to be 32-dimensional.

For the task parameters, we sample $A_c, B_c \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \frac{1}{\sqrt{32}}\mathbf{I})$ and $w \sim \mathcal{N}(0, \mathbf{I})$. Further, for the OoD generalization setup, we instead sample input from a different distribution, i.e., we use $\mathcal{N}(0, 2\mathbf{I})$ instead of standard normal. We also test with different sequence lengths, specifically 3, 5, 10, 20 and 30.

For the models, we consider a shared non-linear encoder that encodes each $(\mathbf{x}_n, \mathbf{c}_n)$ independently. The encoded inputs are then fed to either a monolithic RNN system or to each RNN module of the modular system. We control for the number of parameters such that all the systems roughly share the same number of parameters. Having obtained an updated state from the system, we then use a shared decoder to make the prediction.

We train all the models for 500,000 iterations with a batch-size of 256 and the Adam optimizer with learning rate of 0.0001, with gradient clipping at 1.0. For the classification tasks, we consider binary cross entropy loss, while for regression we consider the l_1 loss.

Classification. We first look at the results on the binary classification based RNN tasks. For reporting performance metrics, we consider all model capacities and all number

Is a Modular Architecture Enough?

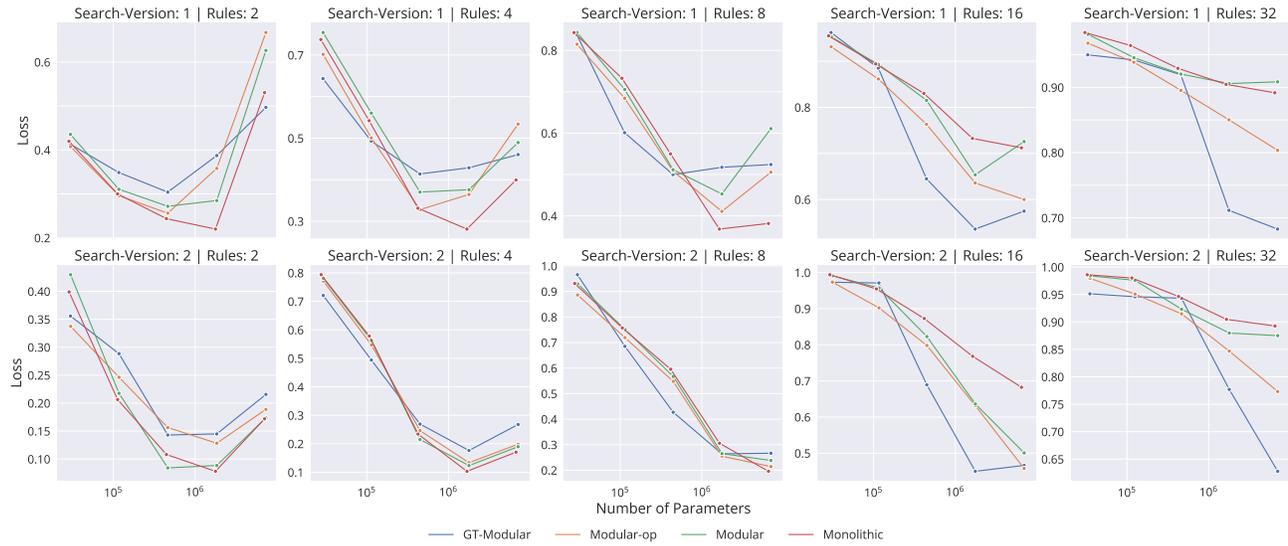


Figure 33. In-Distribution Performance of MHA-Regression Models. Performance (*lower is better*) of different models of varying capacities, different search versions in data and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

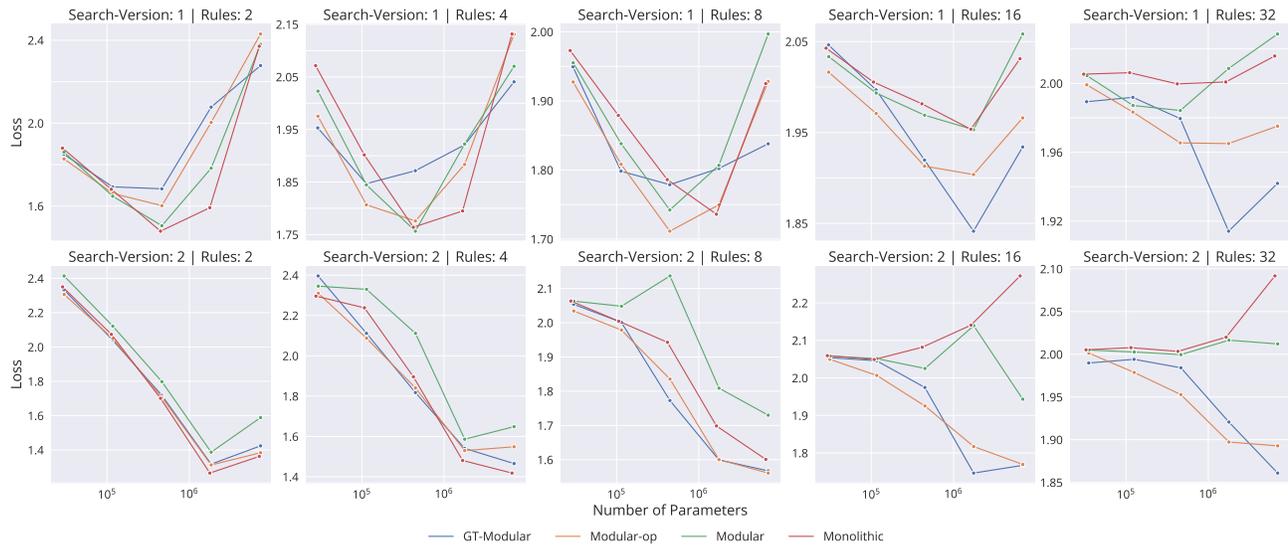


Figure 34. Out-of-Distribution (Sequence Length: 30 - Individual Token Sampling: Altered) Performance on MHA-Regression Models. Performance (*lower is better*) of different models of varying capacities, different search versions in data and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

Is a Modular Architecture Enough?



Figure 35. Collapse-Avg Metric for MHA-Regression Models. Highlights the amount of collapse (*lower is better*) suffered by different models of varying capacities trained across different number of rules. Each bar on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

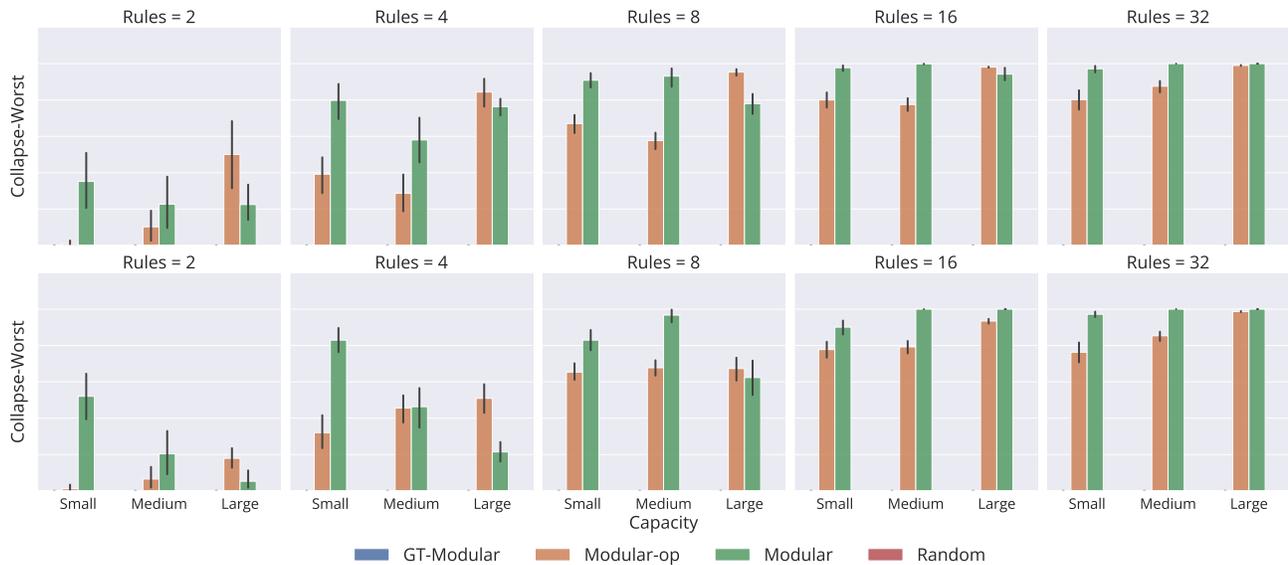


Figure 36. Collapse-Worst Metric for MHA-Regression Models. Highlights the amount of collapse (*lower is better*) suffered by different models of varying capacities trained across different number of rules. Each bar on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

of rules while for collapse and specialization metrics, we consider the smallest, mid-size and largest models and report over the different number of rules.

Performance. For ease of readability, we first provide a snapshot of the results through rankings in Figure 40. The rankings are based on the votes obtained by the different models. Given a task, averaged over the five training seeds,

a vote is given to the model that performs the best. This provides a quick view of the number of times each model outperformed the rest.

We refer the readers to Figure 41 for the in-distribution performance and Figure 42 for the most extreme out-of-distribution performance (where we use the largest sequence length (30) and also change the distribution from which in-

Is a Modular Architecture Enough?

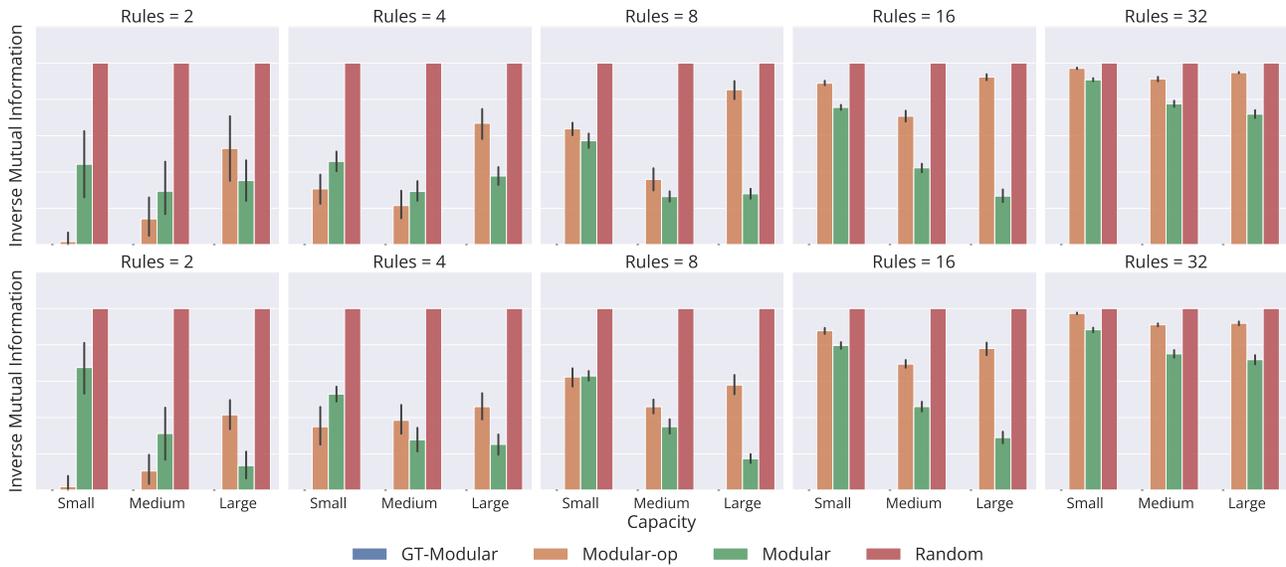


Figure 37. Inverse Mutual Information Metric for MHA-Regression Models. Highlights how poor the specialization (*lower is better*) is of different models of varying capacities trained across different number of rules. Each bar on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.



Figure 38. Adaptation Metric for MHA-Regression Models. Highlights how poor the specialization (*lower is better*) is of different models of varying capacities trained across different number of rules. Each bar on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

dividual tokens are sampled) of the various models across both different model capacities as well as different number of rules. We also refer the reader to Figures 74 - 77 for the out-of-distribution case where only the sequence length is varied in the data and to Figures 78 - 82 where both the sequence length is potentially altered and also the distribution from which individual data points are sampled.

Collapse-Avg. For each rule setting, we report the *Collapse-Avg* metric score of the different models and the three different model capacities in Figure 43.

Collapse-Worst. For each rule setting, we report the *Collapse-Worst* metric score of the different models and the three different model capacities in Figure 44.

Inverse Mutual Information. For each rule setting, we report

Is a Modular Architecture Enough?

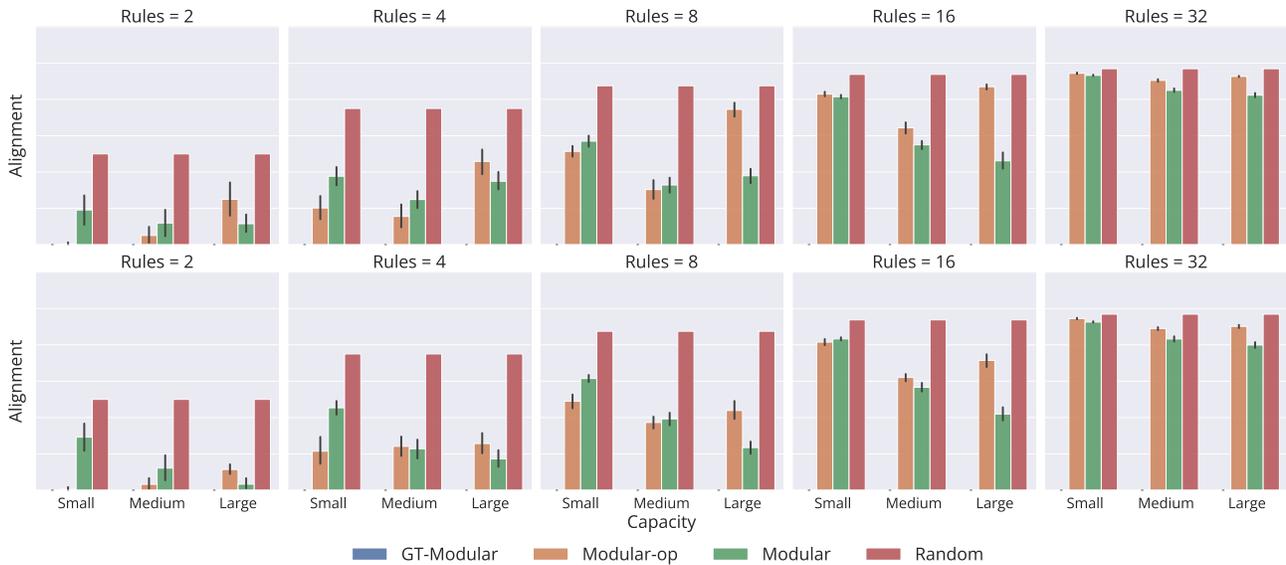


Figure 39. **Alignment Metric for MHA-Classification Models.** Highlights how poor the specialization (*lower is better*) is of different models of varying capacities trained across different number of rules. Each bar on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

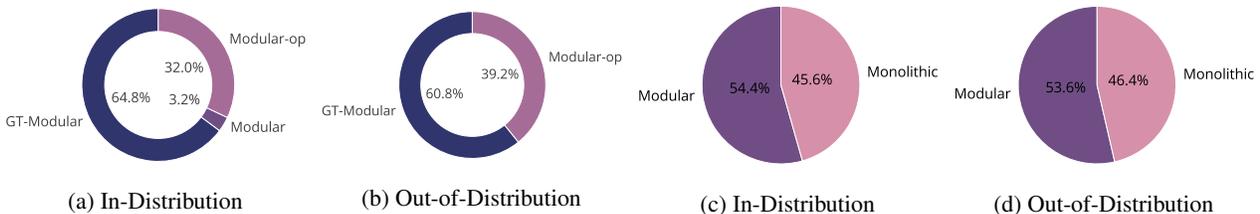


Figure 40. **Ranking Metric for RNN-Classification.** Each pie chart shows the number of times a model wins the competition (*higher is better*), which means outperforms the other models on a single task. Ranking is based on all models with in-distribution ranking in (a) and out-of-distribution ranking in (b). On the contrary, rankings are based only on completely end-to-end trained Modular and Monolithic systems with in-distribution ranking in (c) and out-of-distribution ranking in (d). For OoD, we only consider the extreme setting with largest sequence length and change in distribution of individual tokens.

the *Inverse Mutual Information* metric score of the different models and the three different model capacities in Figure 45.

Adaptation. For each rule setting, we report the *Adaptation* metric score of the different models and the three different model capacities in Figure 46.

Alignment. For each rule setting, we report the *Alignment* metric score of the different models and the three different model capacities in Figure 47.

Regression. Next, we look at the results on the regression based RNN tasks. For reporting performance metrics, we consider all model capacities and all number of rules while for collapse and specialization metrics, we consider the smallest, mid-size and largest models and report over the different number of rules.

Performance. For ease of readability, we first provide a snapshot of the results through rankings in Figure 48. The rankings are based on the votes obtained by the different models. Given a task, averaged over the five training seeds, a vote is given to the model that performs the best. This provides a quick view of the number of times each model outperformed the rest.

We refer the readers to Figure 49 for the in-distribution performance and Figure 50 for the most extreme out-of-distribution performance (where we use the largest sequence length (30) and also change the distribution from which individual tokens are sampled) of the various models across both different model capacities, search functions as well as different number of rules. We also refer the reader to Figures 83 - 86 for the out-of-distribution case where only the sequence length is varied in the data and to Figures 87 - 91 where both the sequence length is potentially altered and

Is a Modular Architecture Enough?

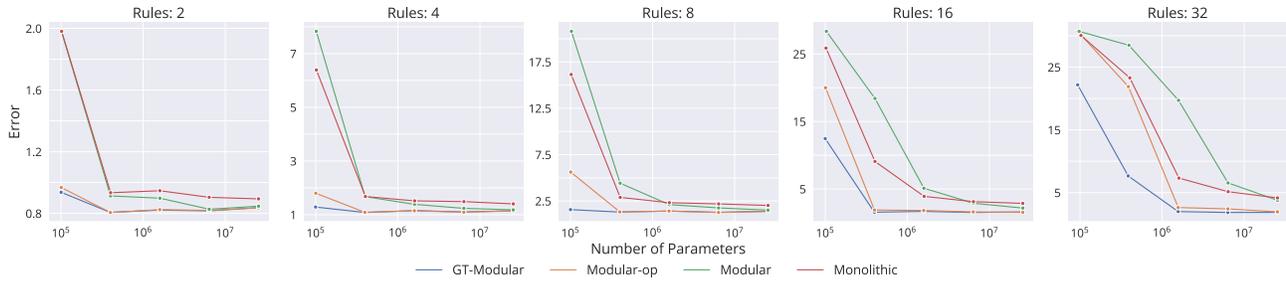


Figure 41. In-Distribution Performance of RNN-Classification Models. Performance (*lower is better*) of different models of varying capacities and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

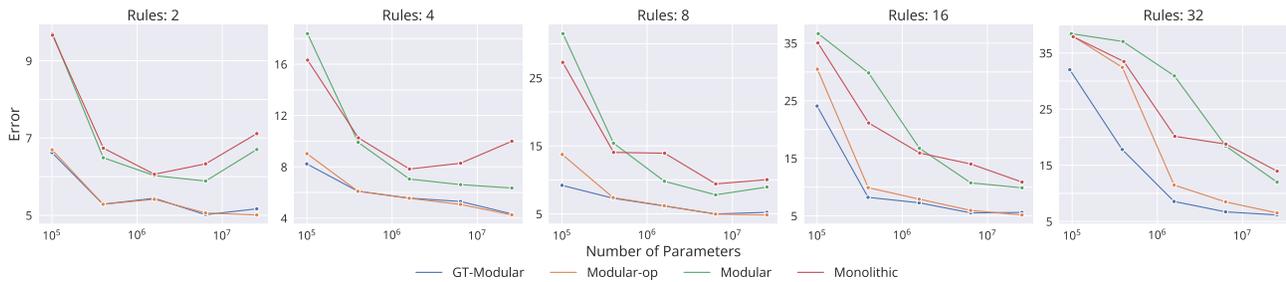


Figure 42. Out-of-Distribution (Sequence Length: 30 - Individual Token Sampling: Altered) Performance on RNN-Classification Models. Performance (*lower is better*) of different models of varying capacities and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

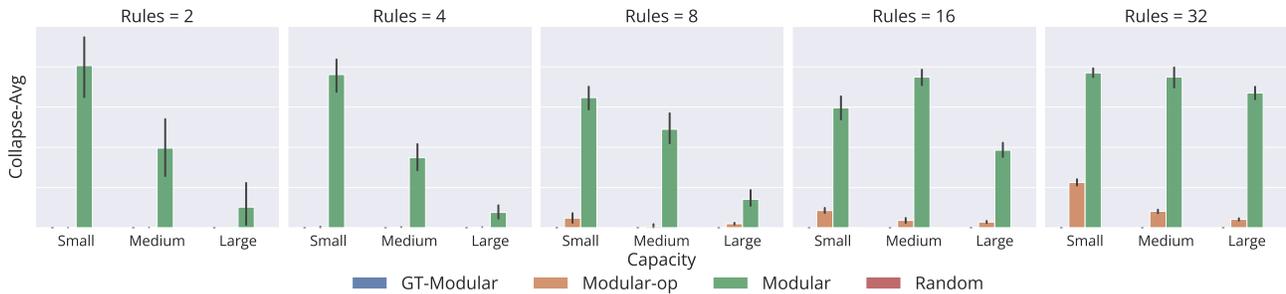


Figure 43. Collapse-Avg Metric for RNN-Classification Models. Highlights the amount of collapse (*lower is better*) suffered by different models of varying capacities trained across different number of rules. Each bar on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

also the distribution from which individual data points are sampled.

Collapse-Avg. For each rule setting, we report the *Collapse-Avg* metric score of the different models and the three different model capacities in Figure 51.

Collapse-Worst. For each rule setting, we report the *Collapse-Worst* metric score of the different models and the three different model capacities in Figure 52.

Inverse Mutual Information. For each rule setting, we report the *Inverse Mutual Information* metric score of the different

models and the three different model capacities in Figure 53.

Adaptation. For each rule setting, we report the *Adaptation* metric score of the different models and the three different model capacities in Figure 54.

Alignment. For each rule setting, we report the *Alignment* metric score of the different models and the three different model capacities in Figure 55.

Is a Modular Architecture Enough?

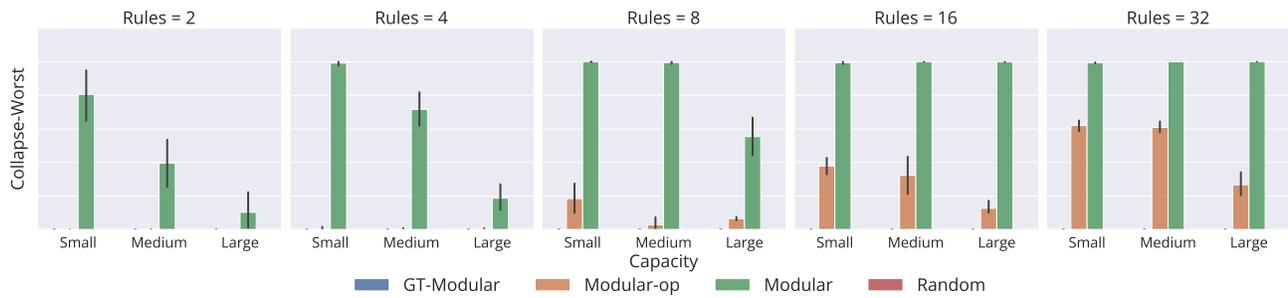


Figure 44. Collapse-Worst Metric for RNN-Classification Models. Highlights the amount of collapse (*lower is better*) suffered by different models of varying capacities trained across different number of rules. Each bar on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

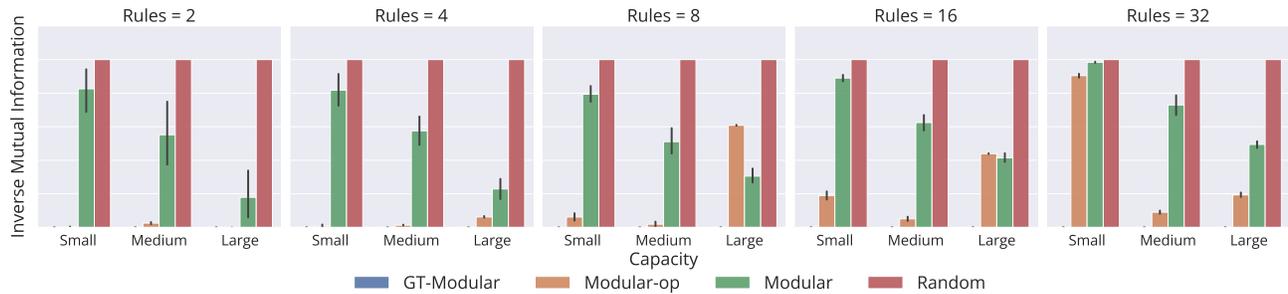


Figure 45. Inverse Mutual Information Metric for RNN-Classification Models. Highlights how poor the specialization (*lower is better*) is of different models of varying capacities trained across different number of rules. Each bar on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

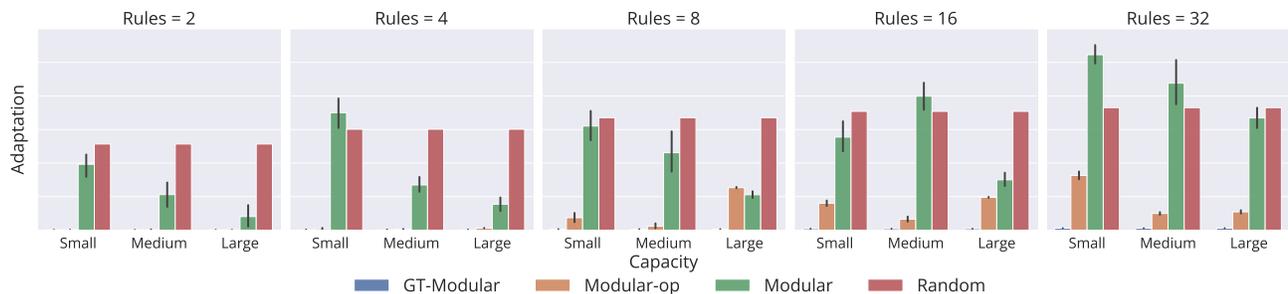


Figure 46. Adaptation Metric for RNN-Classification Models. Highlights how poor the specialization (*lower is better*) is of different models of varying capacities trained across different number of rules. Each bar on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

Is a Modular Architecture Enough?

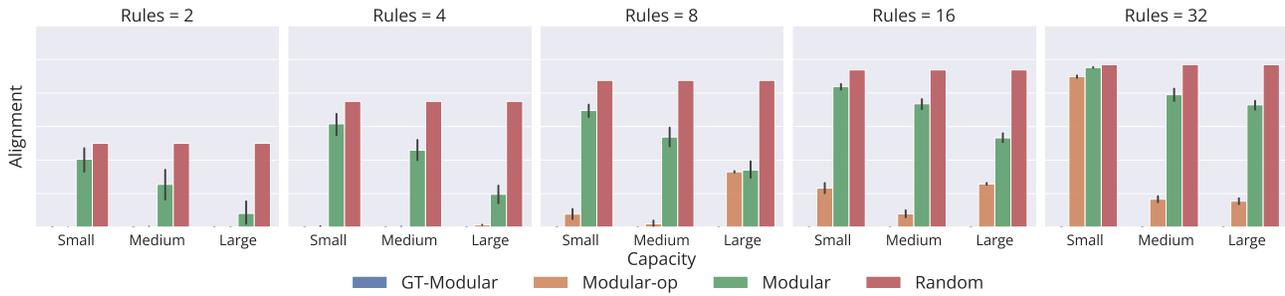


Figure 47. **Alignment Metric for RNN-Classification Models.** Highlights how poor the specialization (*lower is better*) is of different models of varying capacities trained across different number of rules. Each bar on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

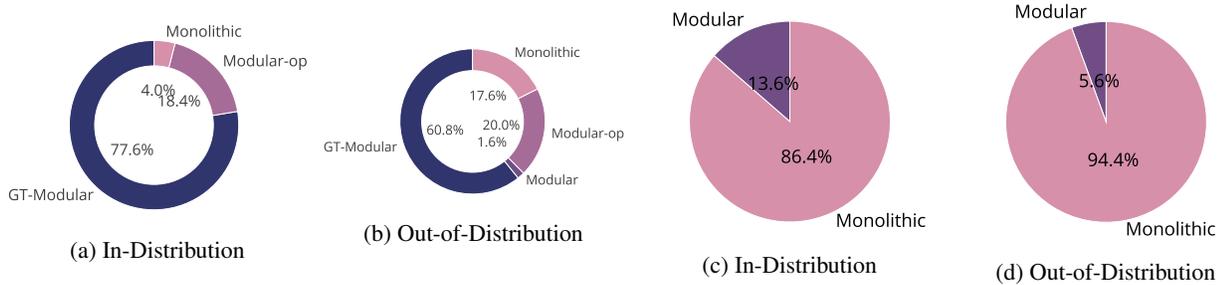


Figure 48. **Ranking Metric for RNN-Regression.** Each pie chart shows the number of times a model wins the competition (*higher is better*), which means outperforms the other models on a single task. Ranking is based on all models with in-distribution ranking in (a) and out-of-distribution ranking in (b). On the contrary, rankings are based only on completely end-to-end trained Modular and Monolithic systems with in-distribution ranking in (c) and out-of-distribution ranking in (d). For OoD, we only consider the extreme setting with largest sequence length and change in distribution of individual tokens.

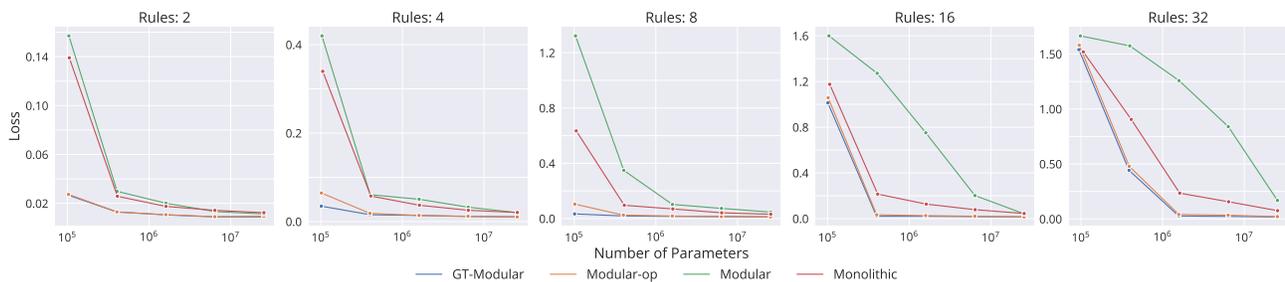


Figure 49. **In-Distribution Performance of RNN-Regression Models.** Performance (*lower is better*) of different models of varying capacities and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

Is a Modular Architecture Enough?

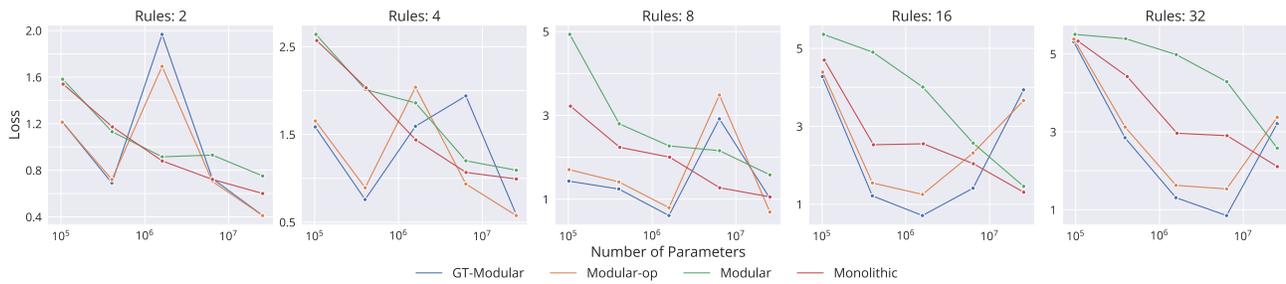


Figure 50. Out-of-Distribution (Sequence Length: 30 - Individual Token Sampling: Altered) Performance on RNN-Regression Models. Performance (*lower is better*) of different models of varying capacities and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

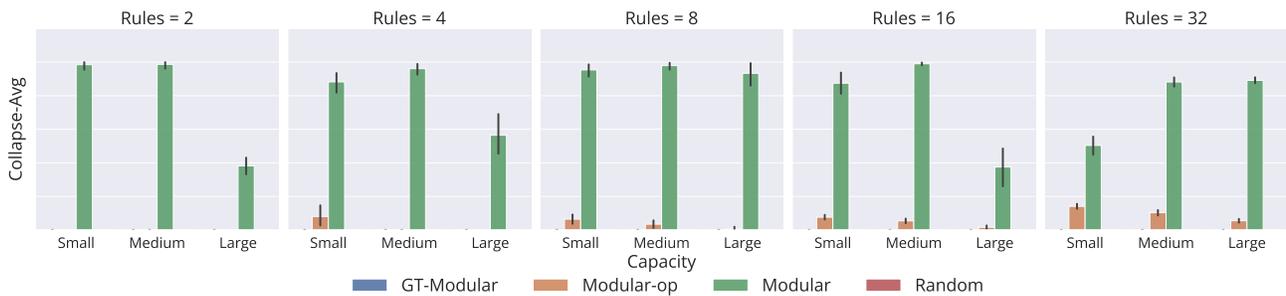


Figure 51. Collapse-Avg Metric for RNN-Regression Models. Highlights the amount of collapse (*lower is better*) suffered by different models of varying capacities trained across different number of rules. Each bar on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

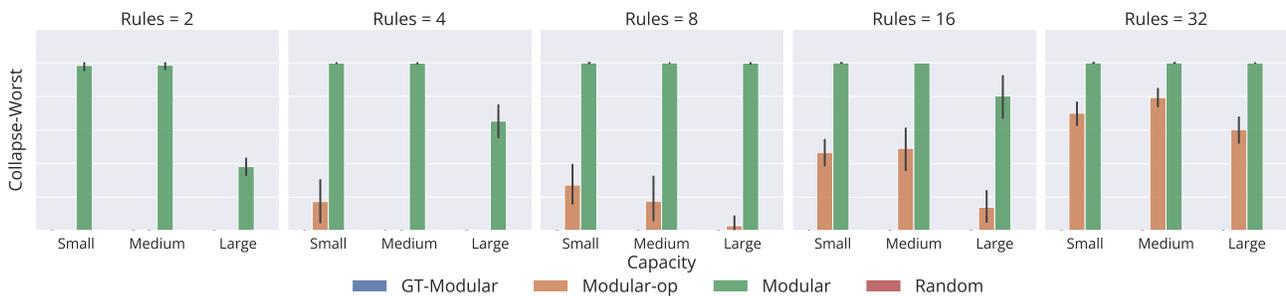


Figure 52. Collapse-Worst Metric for RNN-Regression Models. Highlights the amount of collapse (*lower is better*) suffered by different models of varying capacities trained across different number of rules. Each bar on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

Is a Modular Architecture Enough?



Figure 53. Inverse Mutual Information Metric for RNN-Regression Models. Highlights how poor the specialization (*lower is better*) is of different models of varying capacities trained across different number of rules. Each bar on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

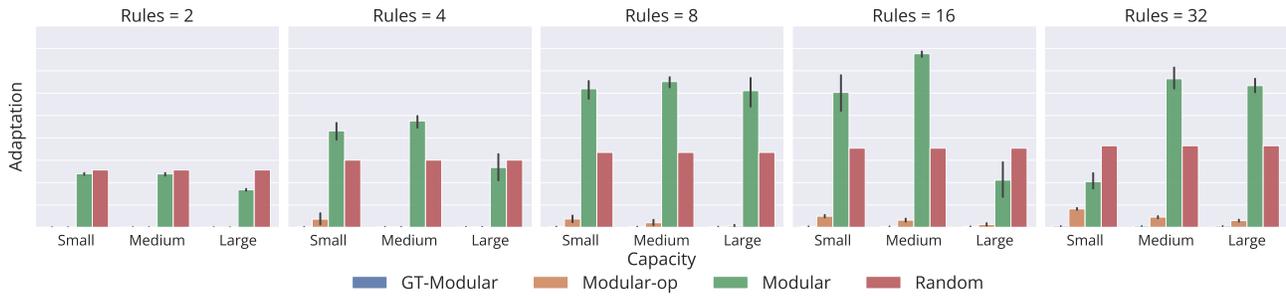


Figure 54. Adaptation Metric for RNN-Regression Models. Highlights how poor the specialization (*lower is better*) is of different models of varying capacities trained across different number of rules. Each bar on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

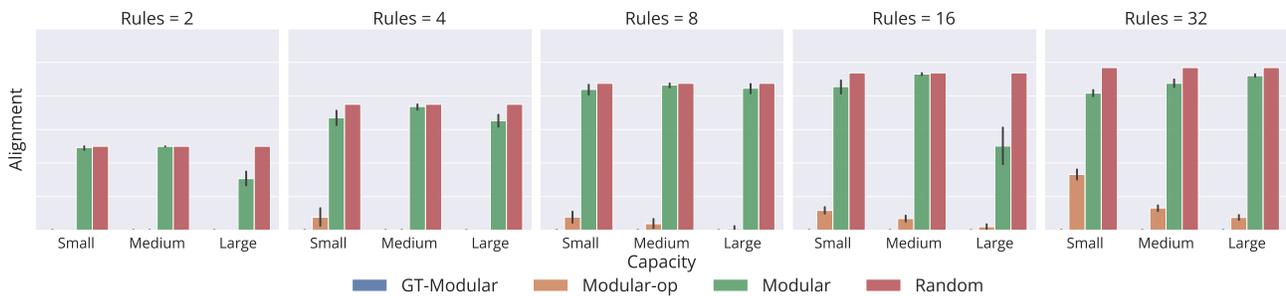


Figure 55. Alignment Metric for RNN-Classification Models. Highlights how poor the specialization (*lower is better*) is of different models of varying capacities trained across different number of rules. Each bar on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

Is a Modular Architecture Enough?

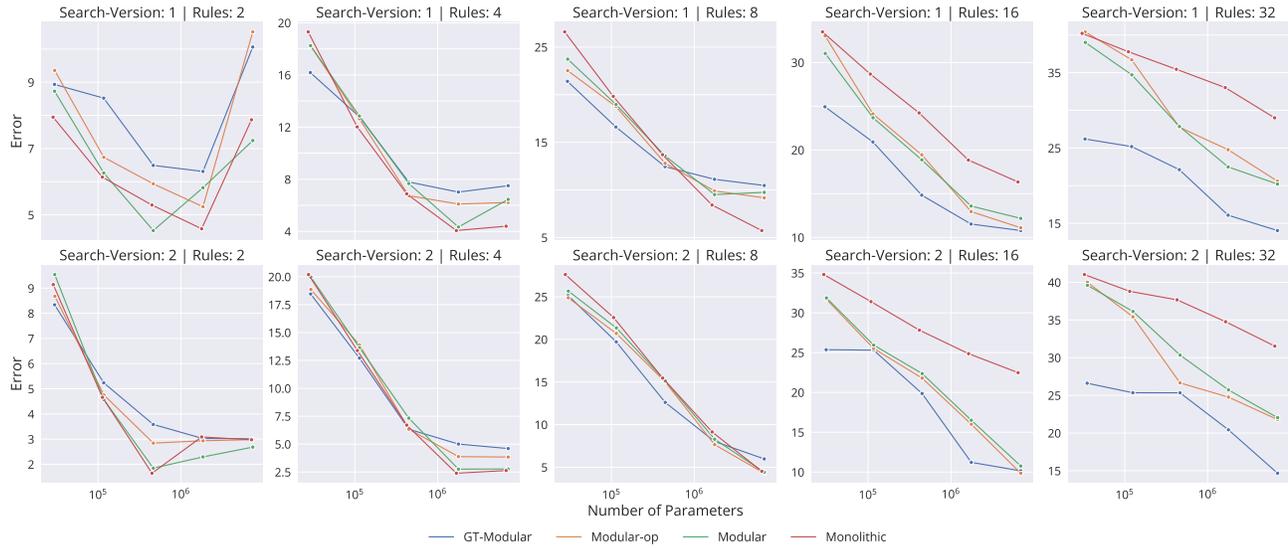


Figure 56. Out-of-Distribution (Sequence Length: 3 - Individual Token Sampling: Same) Performance on MHA-Classification Models. Performance (*lower is better*) of different models of varying capacities, different search versions in data and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

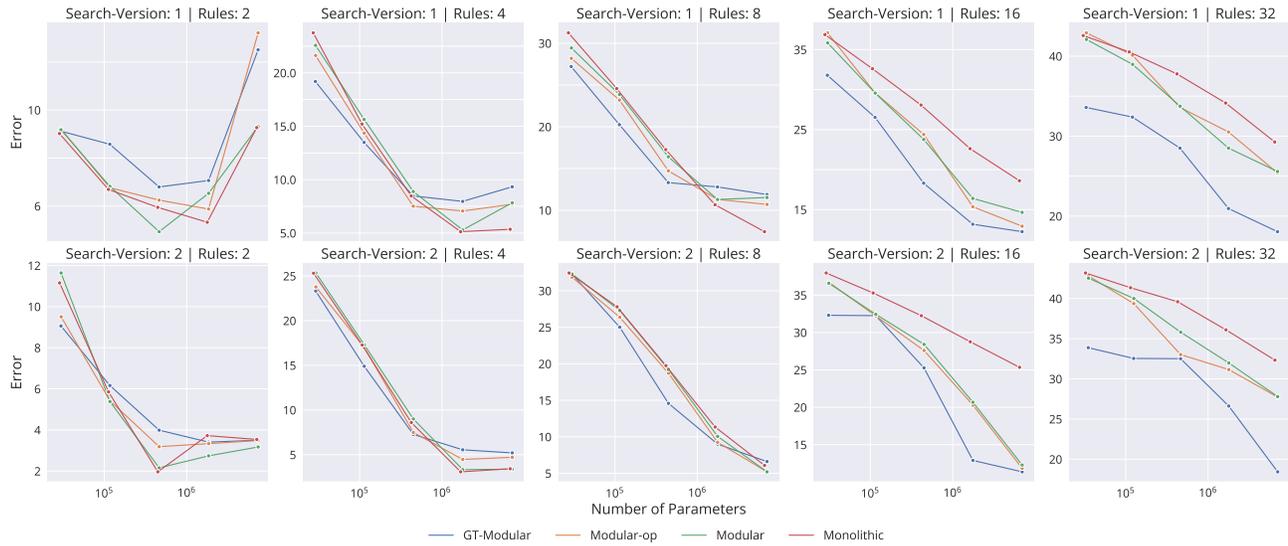


Figure 57. Out-of-Distribution (Sequence Length: 5 - Individual Token Sampling: Same) Performance on MHA-Classification Models. Performance (*lower is better*) of different models of varying capacities, different search versions in data and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

Is a Modular Architecture Enough?

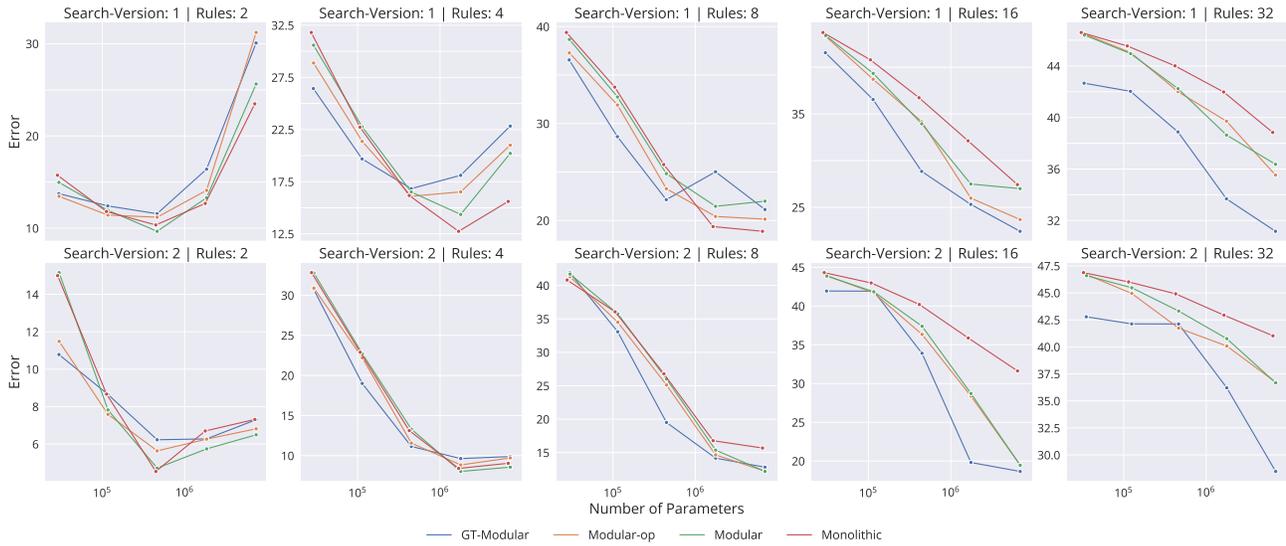


Figure 58. Out-of-Distribution (Sequence Length: 20 - Individual Token Sampling: Same) Performance on MHA-Classification Models. Performance (*lower is better*) of different models of varying capacities, different search versions in data and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

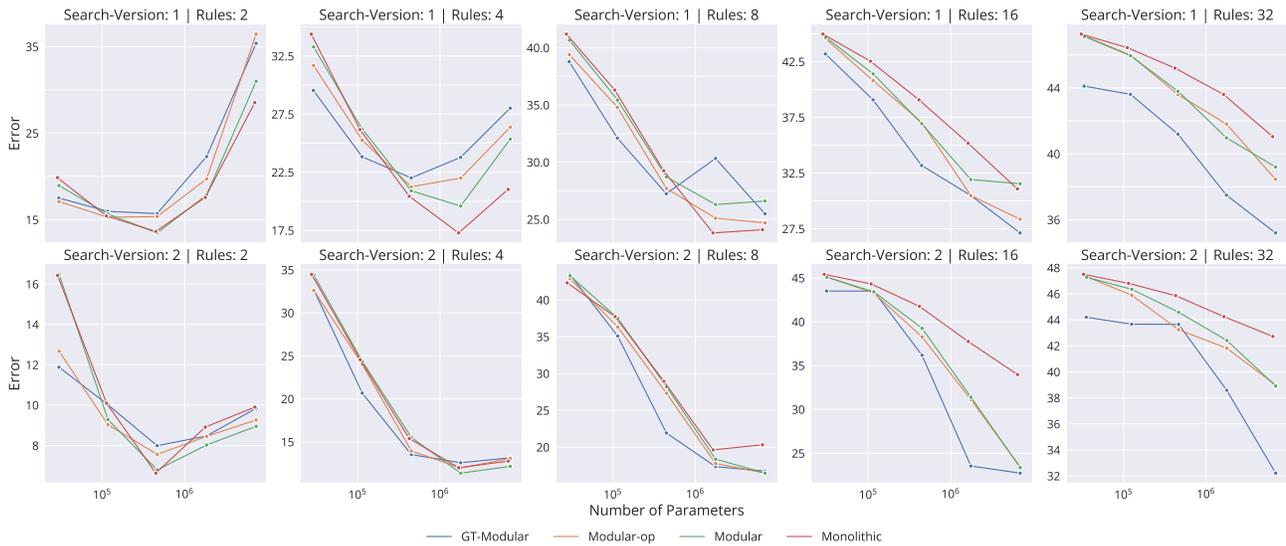


Figure 59. Out-of-Distribution (Sequence Length: 30 - Individual Token Sampling: Same) Performance on MHA-Classification Models. Performance (*lower is better*) of different models of varying capacities, different search versions in data and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

Is a Modular Architecture Enough?

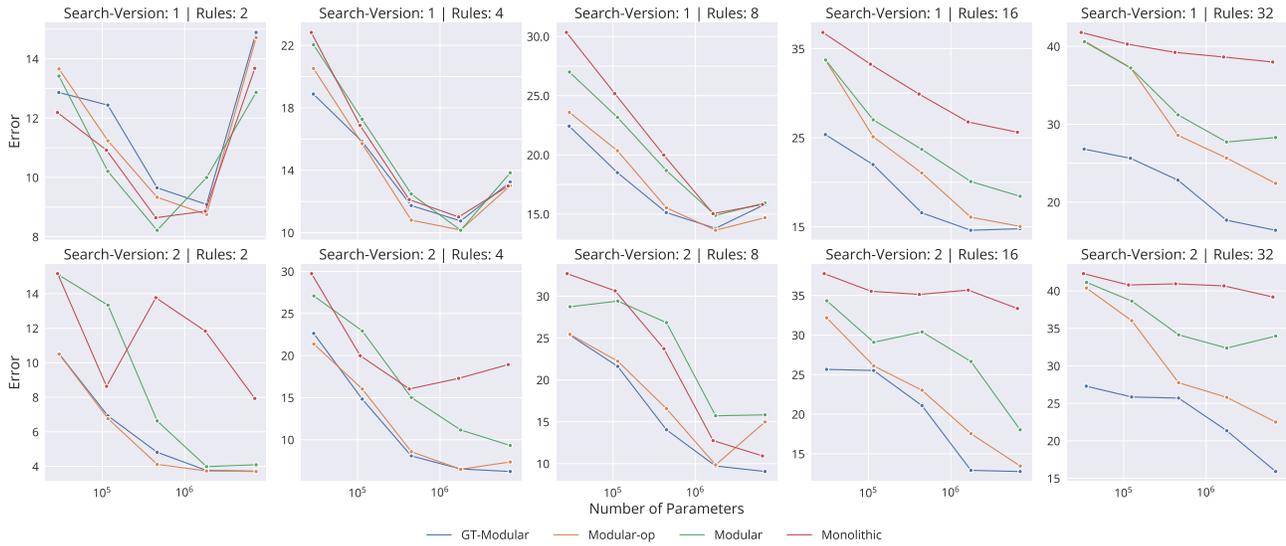


Figure 60. Out-of-Distribution (Sequence Length: 3 - Individual Token Sampling: Altered) Performance on MHA-Classification Models. Performance (*lower is better*) of different models of varying capacities, different search versions in data and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

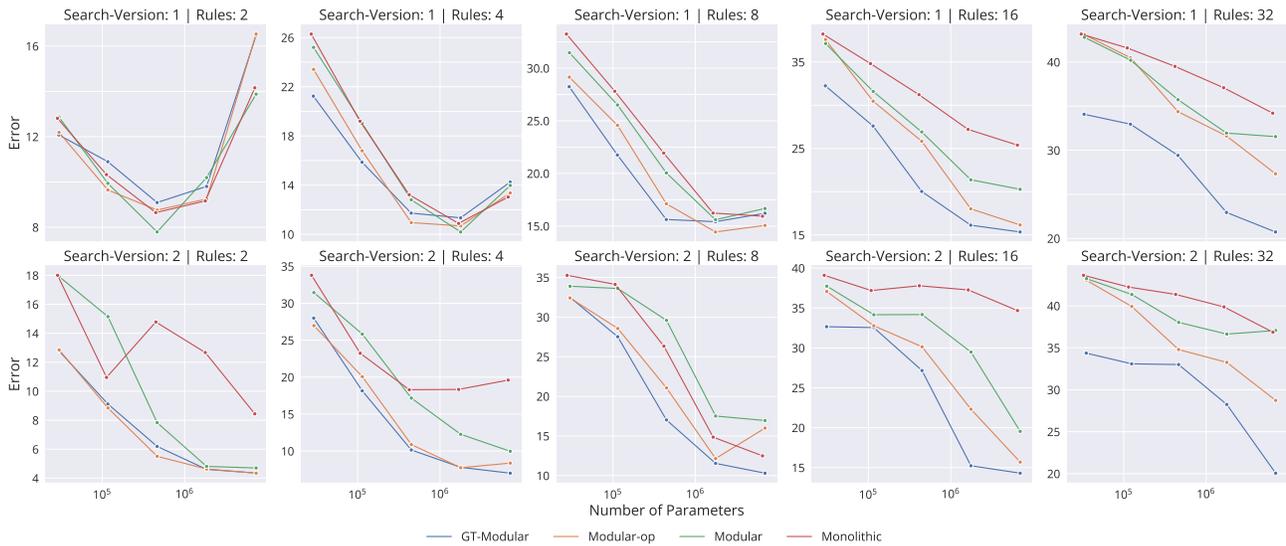


Figure 61. Out-of-Distribution (Sequence Length: 5 - Individual Token Sampling: Altered) Performance on MHA-Classification Models. Performance (*lower is better*) of different models of varying capacities, different search versions in data and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

Is a Modular Architecture Enough?

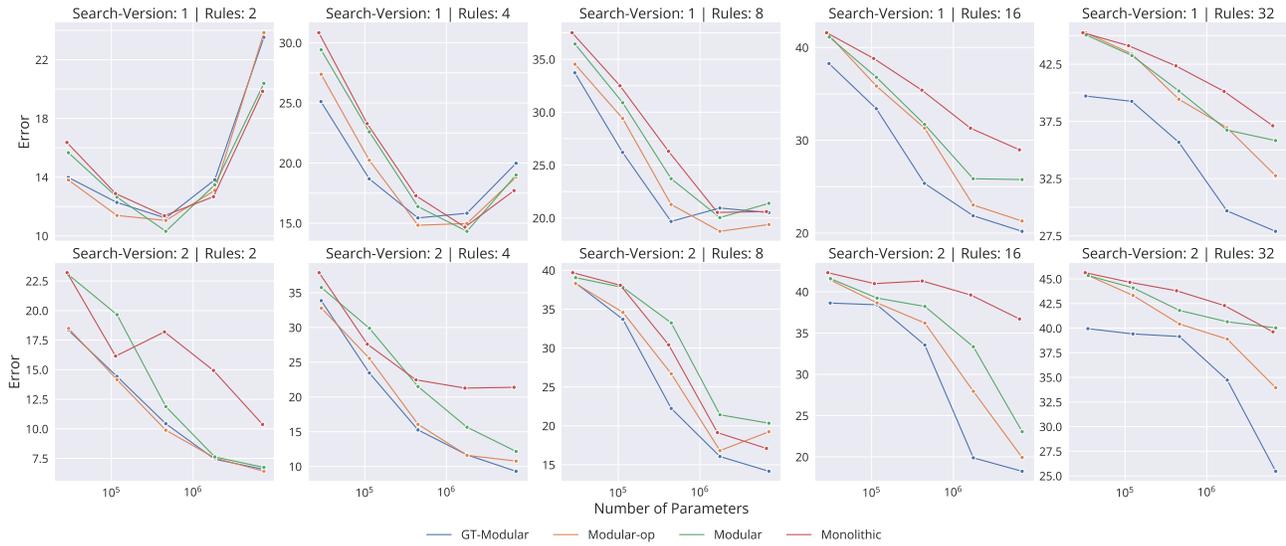


Figure 62. Out-of-Distribution (Sequence Length: 10 - Individual Token Sampling: Altered) Performance on MHA-Classification Models. Performance (*lower is better*) of different models of varying capacities, different search versions in data and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

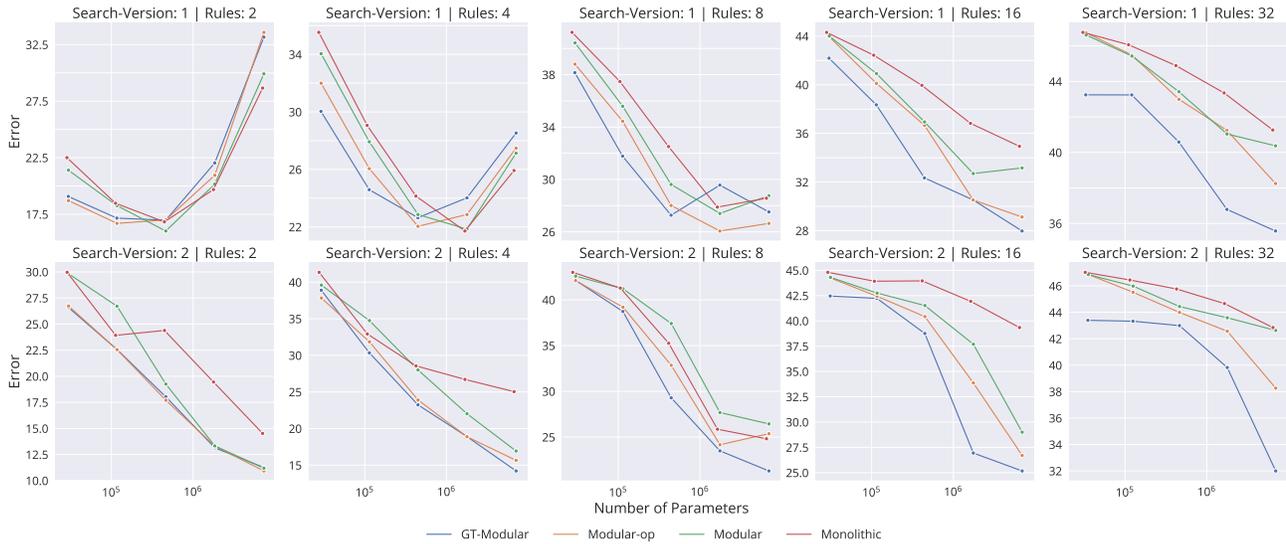


Figure 63. Out-of-Distribution (Sequence Length: 20 - Individual Token Sampling: Altered) Performance on MHA-Classification Models. Performance (*lower is better*) of different models of varying capacities, different search versions in data and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

Is a Modular Architecture Enough?

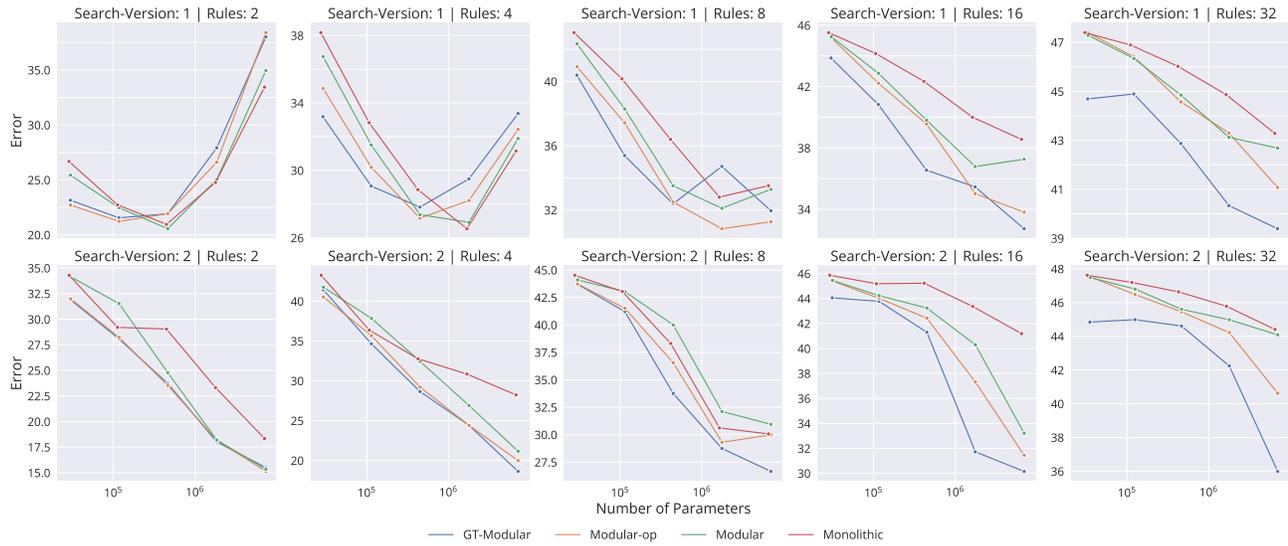


Figure 64. Out-of-Distribution (Sequence Length: 30 - Individual Token Sampling: Altered) Performance on MHA-Classification Models. Performance (*lower is better*) of different models of varying capacities, different search versions in data and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

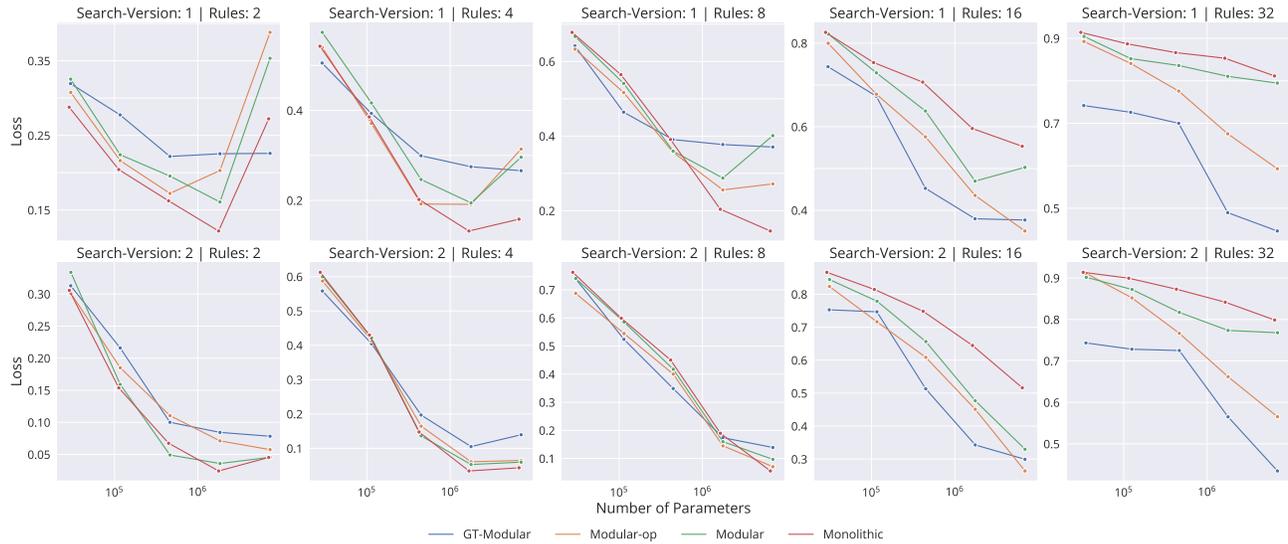


Figure 65. Out-of-Distribution (Sequence Length: 3 - Individual Token Sampling: Same) Performance on MHA-Regression Models. Performance (*lower is better*) of different models of varying capacities, different search versions in data and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

Is a Modular Architecture Enough?

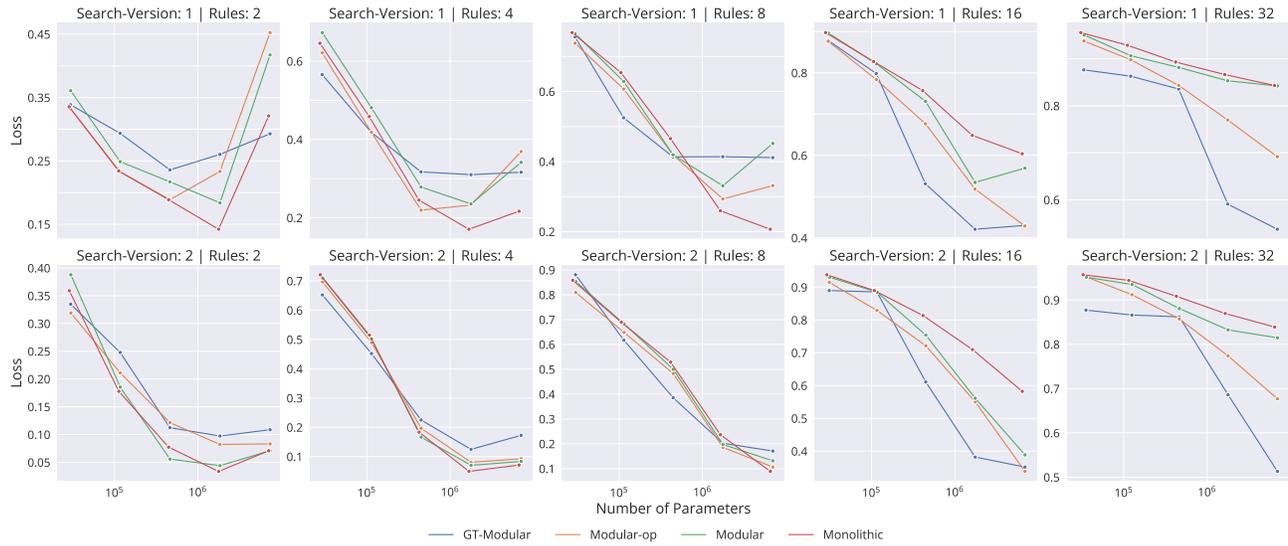


Figure 66. Out-of-Distribution (Sequence Length: 5 - Individual Token Sampling: Same) Performance on MHA-Regression Models. Performance (*lower is better*) of different models of varying capacities, different search versions in data and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

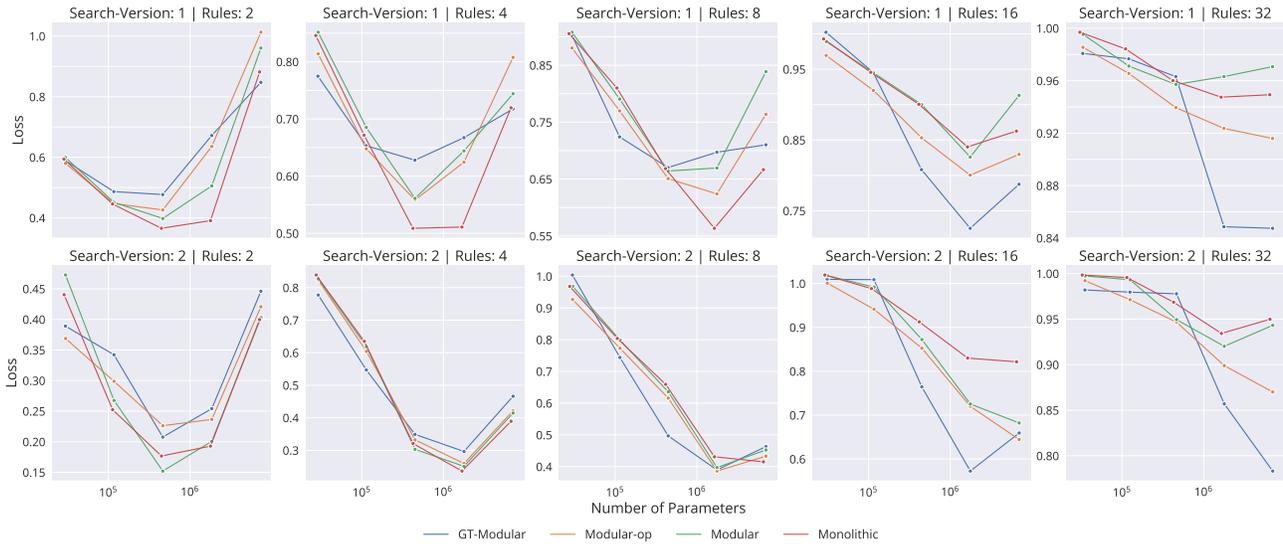


Figure 67. Out-of-Distribution (Sequence Length: 20 - Individual Token Sampling: Same) Performance on MHA-Regression Models. Performance (*lower is better*) of different models of varying capacities, different search versions in data and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

Is a Modular Architecture Enough?

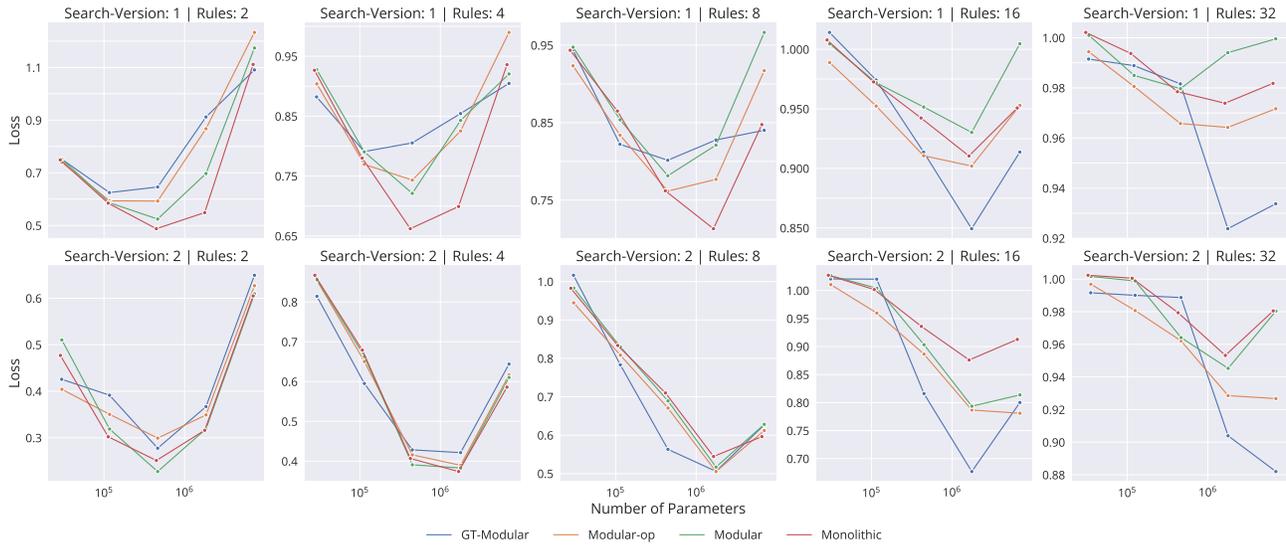


Figure 68. Out-of-Distribution (Sequence Length: 30 - Individual Token Sampling: Same) Performance on MHA-Regression Models. Performance (*lower is better*) of different models of varying capacities, different search versions in data and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.



Figure 69. Out-of-Distribution (Sequence Length: 3 - Individual Token Sampling: Altered) Performance on MHA-Regression Models. Performance (*lower is better*) of different models of varying capacities, different search versions in data and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

Is a Modular Architecture Enough?

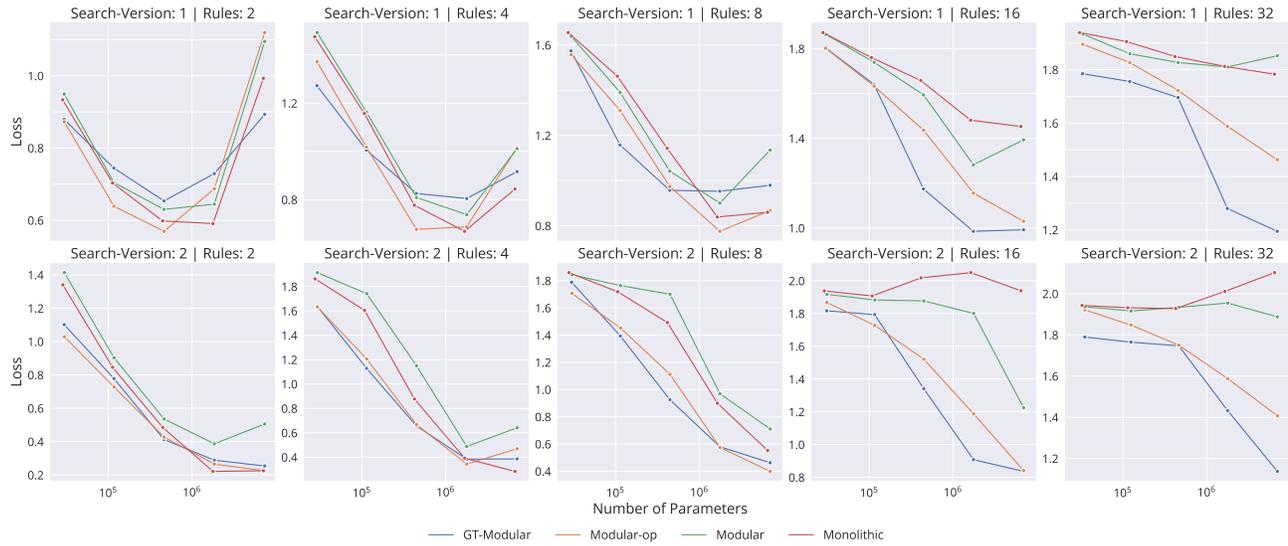


Figure 70. Out-of-Distribution (Sequence Length: 5 - Individual Token Sampling: Altered) Performance on MHA-Regression Models. Performance (*lower is better*) of different models of varying capacities, different search versions in data and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

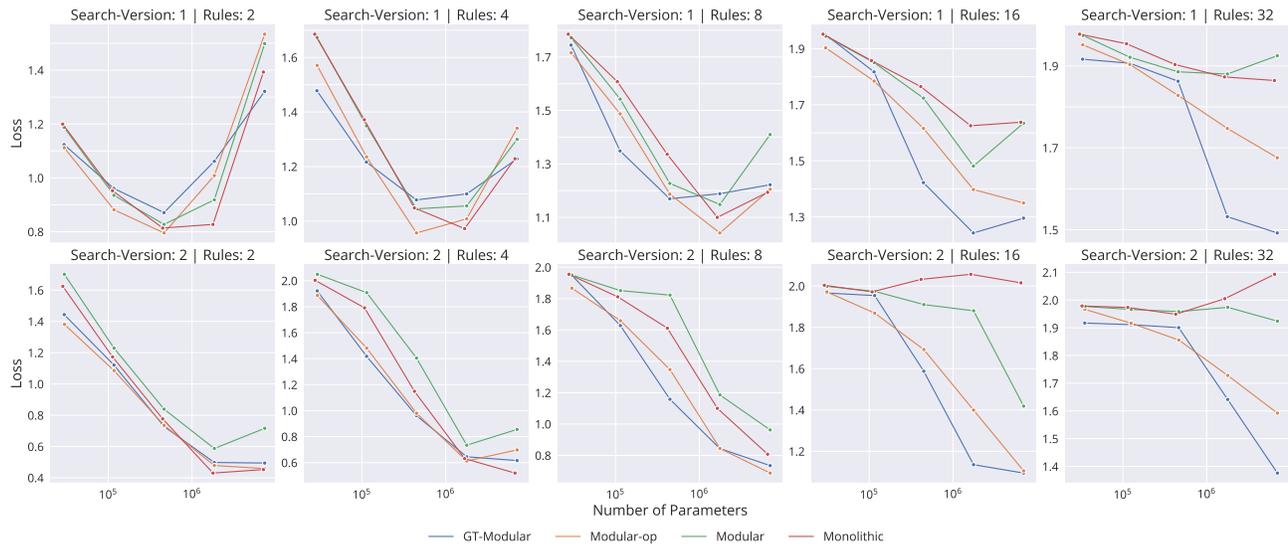


Figure 71. Out-of-Distribution (Sequence Length: 10 - Individual Token Sampling: Altered) Performance on MHA-Regression Models. Performance (*lower is better*) of different models of varying capacities, different search versions in data and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

Is a Modular Architecture Enough?

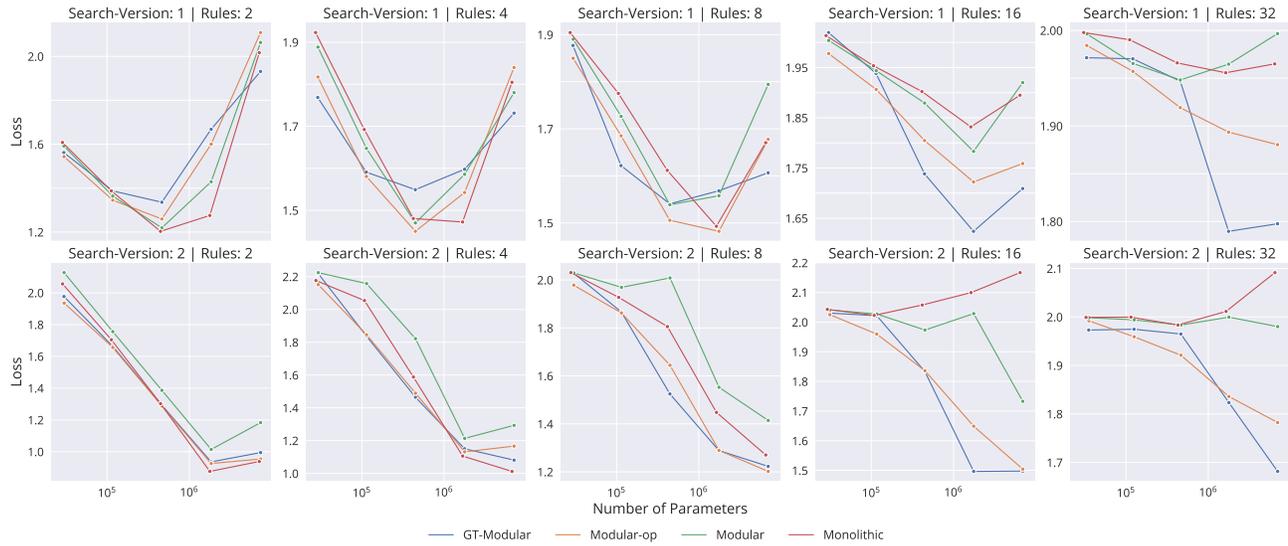


Figure 72. Out-of-Distribution (Sequence Length: 20 - Individual Token Sampling: Altered) Performance on MHA-Regression Models. Performance (*lower is better*) of different models of varying capacities, different search versions in data and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

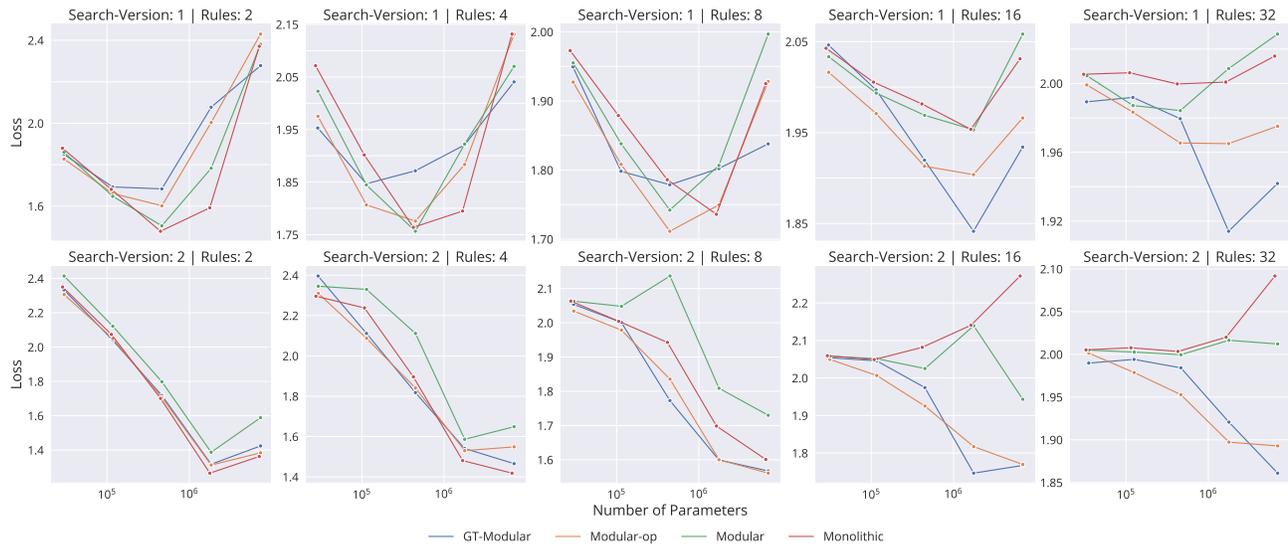


Figure 73. Out-of-Distribution (Sequence Length: 30 - Individual Token Sampling: Altered) Performance on MHA-Regression Models. Performance (*lower is better*) of different models of varying capacities, different search versions in data and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

Is a Modular Architecture Enough?

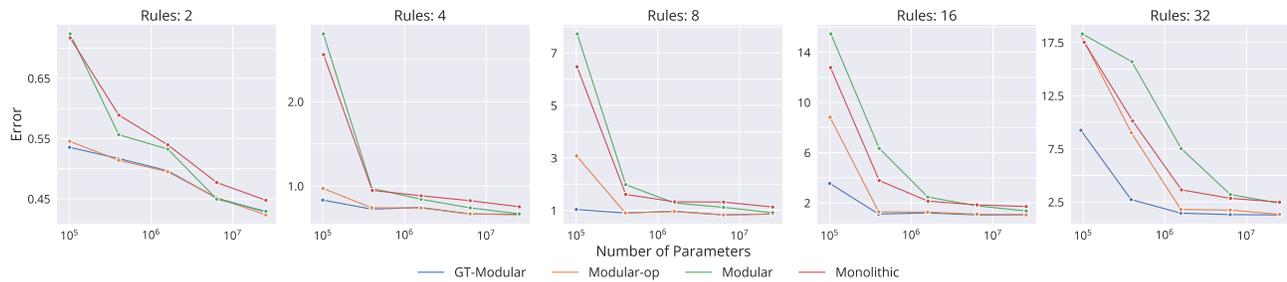


Figure 74. Out-of-Distribution (Sequence Length: 3 - Individual Token Sampling: Same) Performance on RNN-Classification Models. Performance (*lower is better*) of different models of varying capacities and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

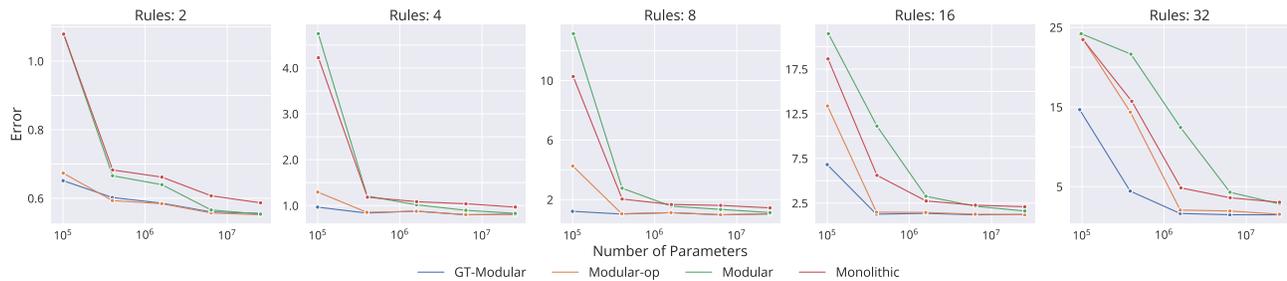


Figure 75. Out-of-Distribution (Sequence Length: 5 - Individual Token Sampling: Same) Performance on RNN-Classification Models. Performance (*lower is better*) of different models of varying capacities and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

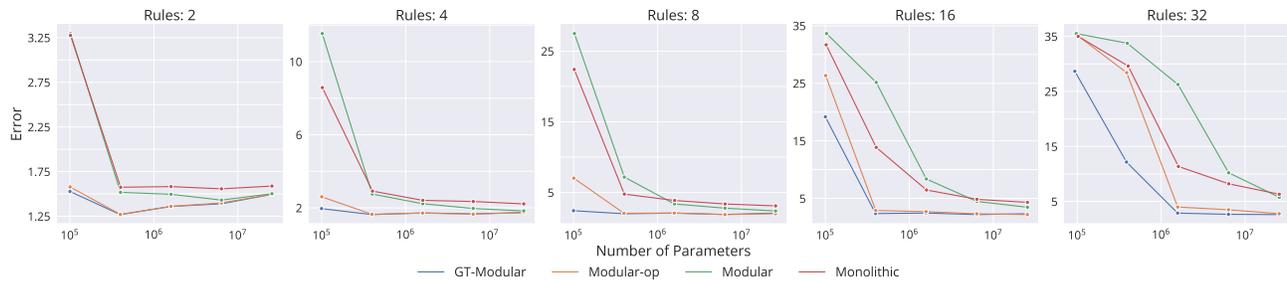


Figure 76. Out-of-Distribution (Sequence Length: 20 - Individual Token Sampling: Same) Performance on RNN-Classification Models. Performance (*lower is better*) of different models of varying capacities and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

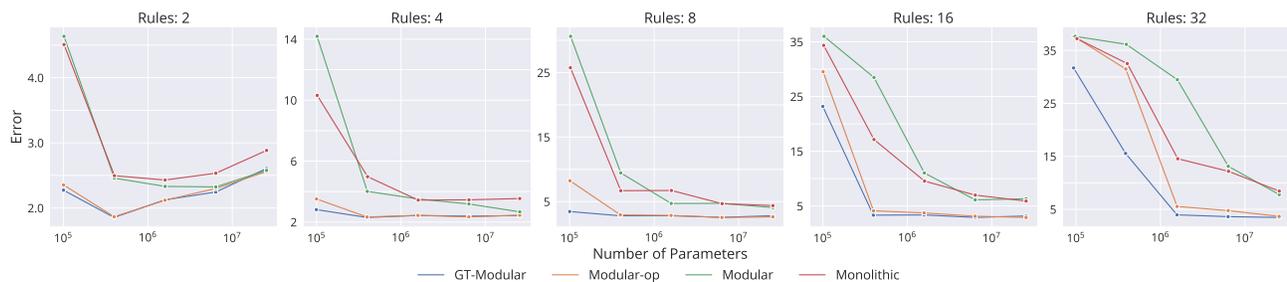


Figure 77. Out-of-Distribution (Sequence Length: 30 - Individual Token Sampling: Same) Performance on RNN-Classification Models. Performance (*lower is better*) of different models of varying capacities and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

Is a Modular Architecture Enough?

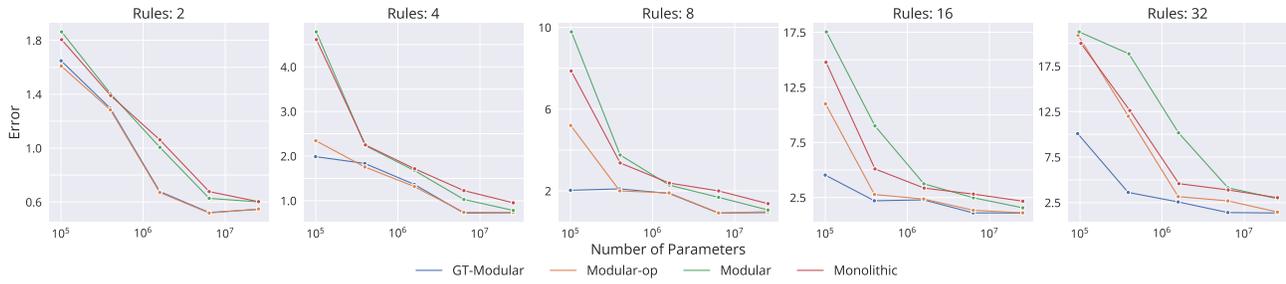


Figure 78. Out-of-Distribution (Sequence Length: 3 - Individual Token Sampling: Altered) Performance on RNN-Classification Models. Performance (*lower is better*) of different models of varying capacities and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

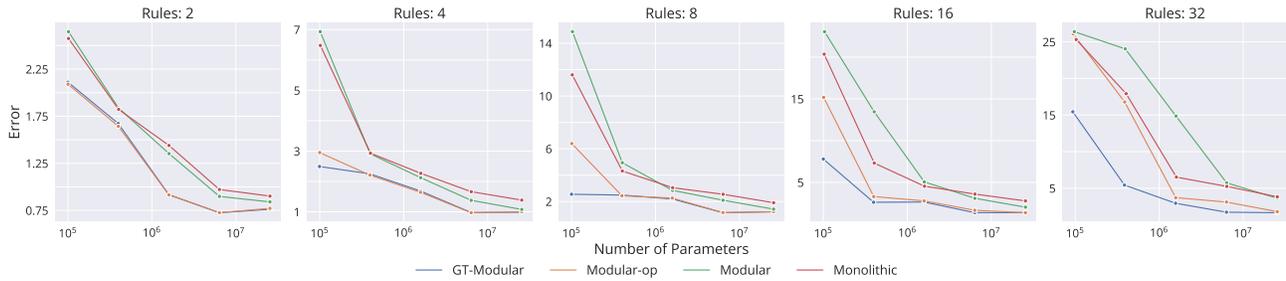


Figure 79. Out-of-Distribution (Sequence Length: 5 - Individual Token Sampling: Altered) Performance on RNN-Classification Models. Performance (*lower is better*) of different models of varying capacities and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

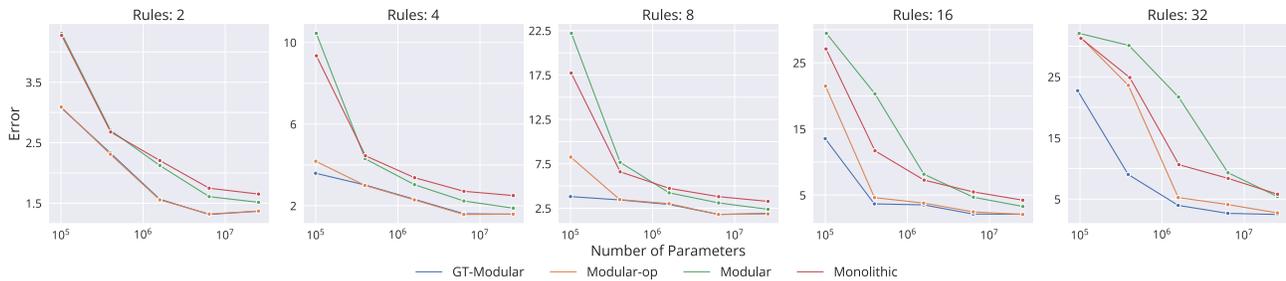


Figure 80. Out-of-Distribution (Sequence Length: 10 - Individual Token Sampling: Altered) Performance on RNN-Classification Models. Performance (*lower is better*) of different models of varying capacities and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

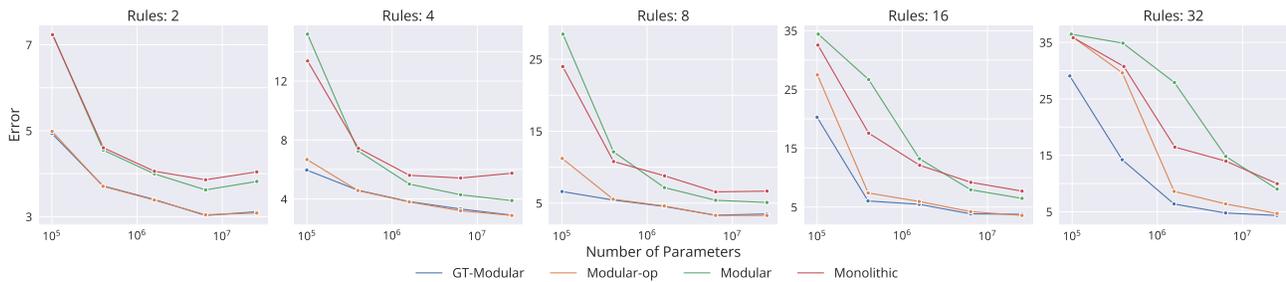


Figure 81. Out-of-Distribution (Sequence Length: 20 - Individual Token Sampling: Altered) Performance on RNN-Classification Models. Performance (*lower is better*) of different models of varying capacities and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

Is a Modular Architecture Enough?

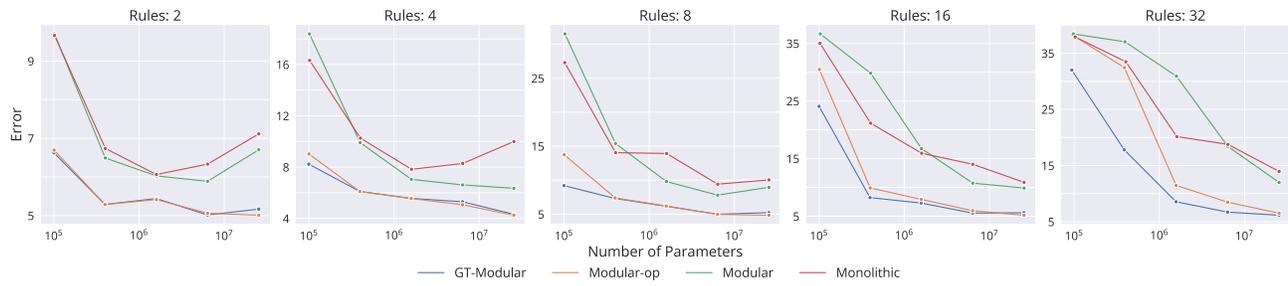


Figure 82. Out-of-Distribution (Sequence Length: 30 - Individual Token Sampling: Altered) Performance on RNN-Classification Models. Performance (*lower is better*) of different models of varying capacities and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

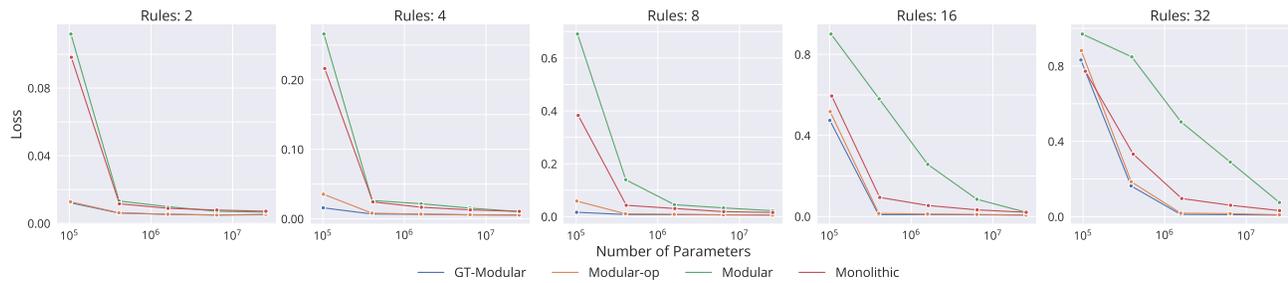


Figure 83. Out-of-Distribution (Sequence Length: 3 - Individual Token Sampling: Same) Performance on RNN-Regression Models. Performance (*lower is better*) of different models of varying capacities and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

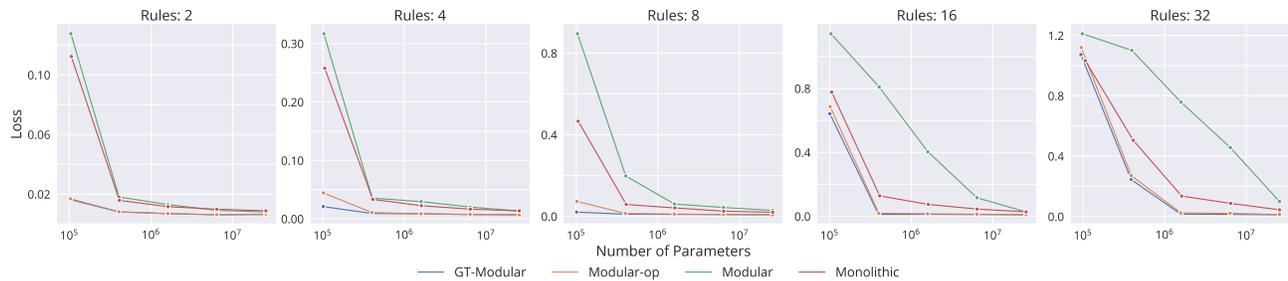


Figure 84. Out-of-Distribution (Sequence Length: 5 - Individual Token Sampling: Same) Performance on RNN-Regression Models. Performance (*lower is better*) of different models of varying capacities and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

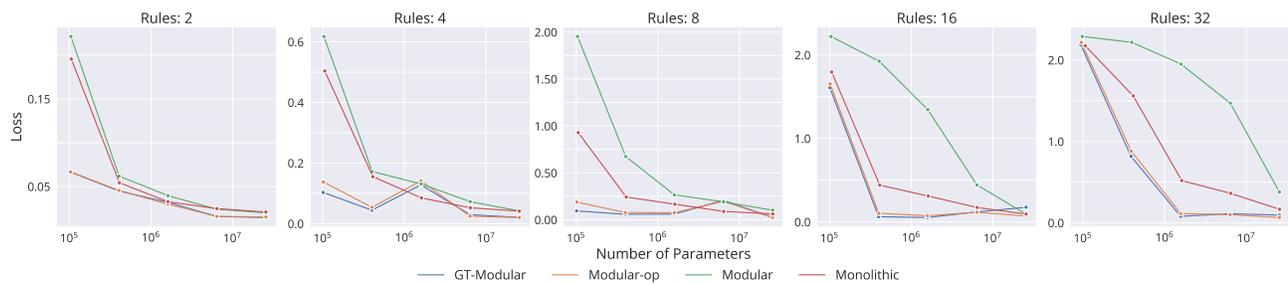


Figure 85. Out-of-Distribution (Sequence Length: 20 - Individual Token Sampling: Same) Performance on RNN-Regression Models. Performance (*lower is better*) of different models of varying capacities and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

Is a Modular Architecture Enough?

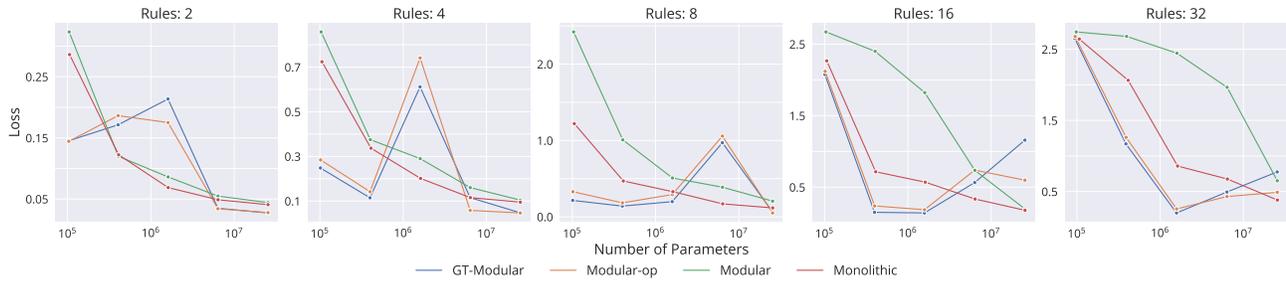


Figure 86. Out-of-Distribution (Sequence Length: 30 - Individual Token Sampling: Same) Performance on RNN-Regression Models. Performance (*lower is better*) of different models of varying capacities and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

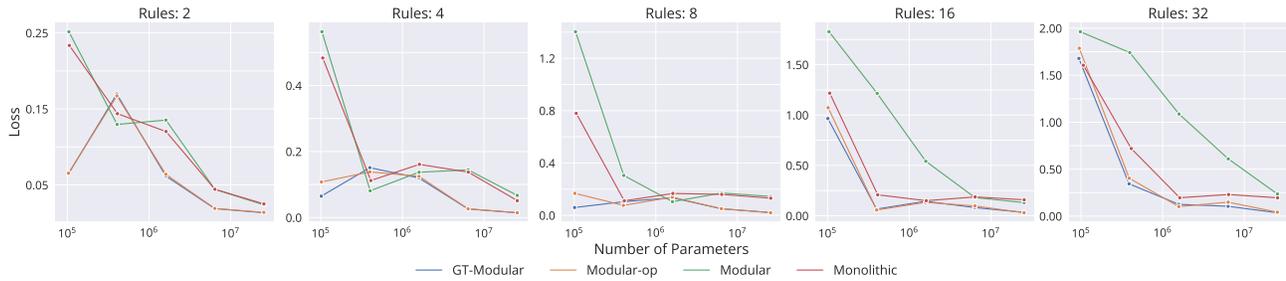


Figure 87. Out-of-Distribution (Sequence Length: 3 - Individual Token Sampling: Altered) Performance on RNN-Regression Models. Performance (*lower is better*) of different models of varying capacities and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

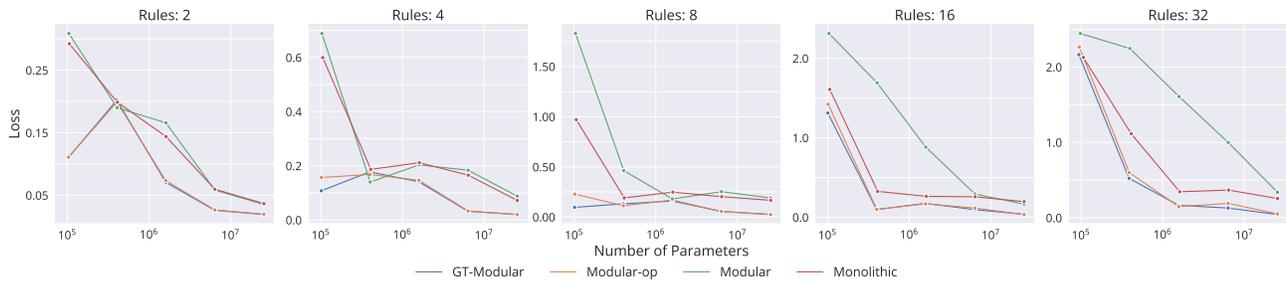


Figure 88. Out-of-Distribution (Sequence Length: 5 - Individual Token Sampling: Altered) Performance on RNN-Regression Models. Performance (*lower is better*) of different models of varying capacities and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

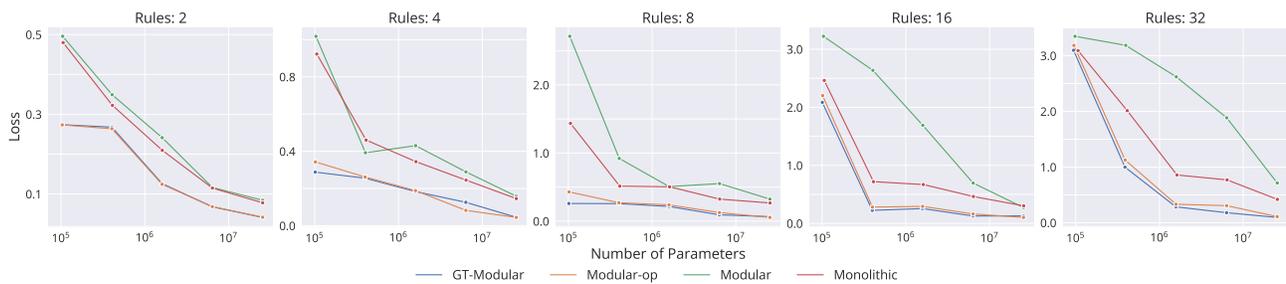


Figure 89. Out-of-Distribution (Sequence Length: 10 - Individual Token Sampling: Altered) Performance on RNN-Regression Models. Performance (*lower is better*) of different models of varying capacities and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

Is a Modular Architecture Enough?

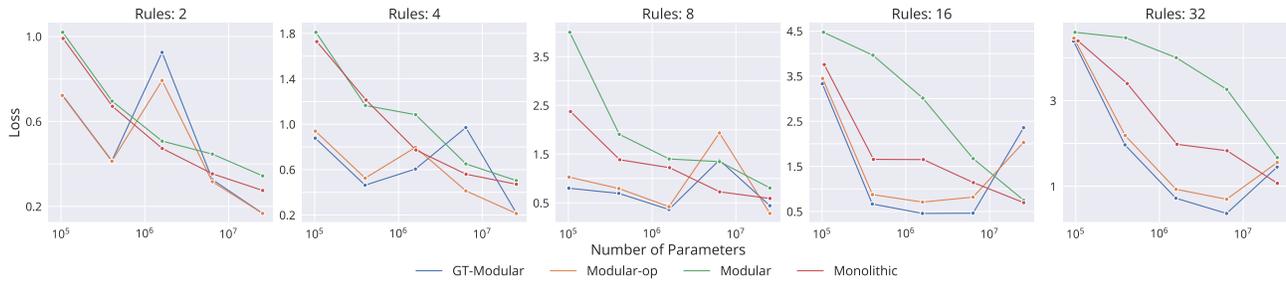


Figure 90. Out-of-Distribution (Sequence Length: 20 - Individual Token Sampling: Altered) Performance on RNN-Regression Models. Performance (*lower is better*) of different models of varying capacities and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.

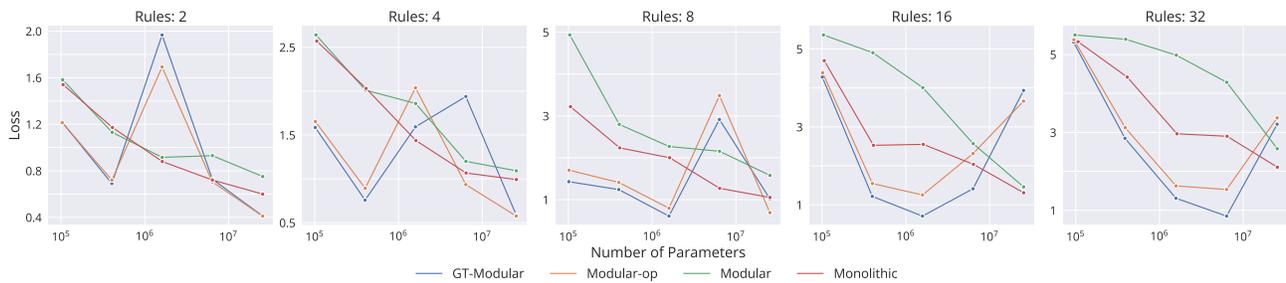


Figure 91. Out-of-Distribution (Sequence Length: 30 - Individual Token Sampling: Altered) Performance on RNN-Regression Models. Performance (*lower is better*) of different models of varying capacities and trained across different number of rules. Each point on the graph is obtained from an average over five tasks, each with five seeds, totaling 25 runs.