# ADAPTIVITY AND MODULARITY FOR EFFICIENT GENERALIZATION OVER TASK COMPLEXITY

**Samira Abnar, Omid Saremi, Laurent Dinh, Shantel Wilson,**
**Miguel Angel Bautista, Chen Huang, Vimal Thilak, Etai Littwin**
**Jiatao Gu, Josh Susskind, Samy Bengio**
Apple

arXiv:2310.08866v1 [cs.LG] 13 Oct 2023

## ABSTRACT

Can transformers generalize efficiently on problems that require dealing with examples with different levels of difficulty? We introduce a new task tailored to assess generalization over different complexities and present results that indicate that standard transformers face challenges in solving these tasks. These tasks are variations of pointer value retrieval previously introduced by Zhang et al. (2021). We investigate how the use of a mechanism for adaptive and modular computation in transformers facilitates the learning of tasks that demand generalization over the number of sequential computation steps (i.e., the depth of the computation graph). Based on our observations, we propose a transformer-based architecture called Hyper-UT, which combines dynamic function generation from hyper networks with adaptive depth from Universal Transformers. This model demonstrates higher accuracy and a fairer allocation of computational resources when generalizing to higher numbers of computation steps. We conclude that mechanisms for adaptive depth and modularity complement each other in improving efficient generalization concerning example complexity. Additionally, to emphasize the broad applicability of our findings, we illustrate that in a standard image recognition task, Hyper-UT's performance matches that of a ViT model but with considerably reduced computational demands (achieving over 70% average savings by effectively using fewer layers).

## 1 INTRODUCTION

Tackling many real-world problems such as scientific research, math problem solving (Saxton et al., 2019; Wang & Lu, 2023), and parsing scenes (Kong & Fowlkes, 2018), requires reasoning over a chain of steps (Baldock et al., 2021; Agarwal et al., 2022) where the sequence of steps and the length of the chain are not known immediately. We conjecture that the basic ingredient to solve these learning problems in a generalizable and efficient manner is the capability to break up the learning problem into reusable components and compose them systematically (structured replacement of operations) and productively (constructing more complex operations by composing simpler ones) (Szabó, 2008).

An emergent technique to deal with this setting in Large Language Models (LLMs) is to use chain of thought (Wei et al., 2022a;b; Lee & Kim, 2023) where the chain of thought adjusts how much and what type of compute the model applies to solve a given example. *Can NNs learn a generalized chain of thought reasoning process without the need to tie it explicitly to their input/output?* Bubeck et al. (2023) argues that one of the limitations of LLMs such as GPT-4 is their inability to perform multi-step computations without an explicit scratch-pad (e.g., if not prompted to output the intermediate steps). For example, GPT-4 can correctly answer the question of counting prime numbers in a range only if it is asked to first list these numbers.

Mechanisms for having adaptive depth (Graves, 2016; Dehghani et al., 2019; Banino et al., 2021) and modularity/sparsity Perez et al. (2018); Ha et al. (2017); Pfeiffer et al. (2023) offer natural solutions to compose and execute computation graphs dynamically. Modularity helps with composing operations systematically and adaptivity enables constructing operations with varying complexities. Hence

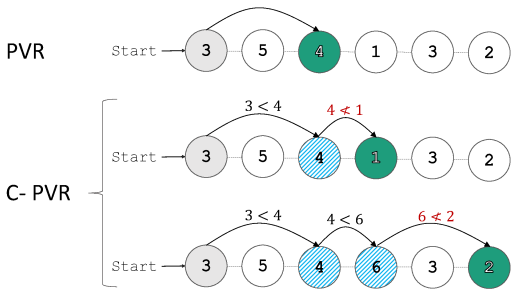


Figure 1: **Example of PVR and C-PVR Tasks.** In the regular PVR task, only one part of the input (in this case, the first element in the sequence) serves as the pointer. However, in the C-PVR, each element can act as a pointer. The task is to recursively retrieve the values until a condition is met. For this illustration, the condition is to continue as long as the sequence progresses forward. In this diagram, gray circles refer to initial steps, blue to intermediate steps, and green the final steps (the value of the green circle is returned).

incorporating these mechanisms simultaneously into machine learning systems should increase their efficiency and performance in multi-step reasoning scenarios.

We design a synthetic task to probe the capability of models to generalize across complexity of examples in multi-step reasoning tasks. We aim to abstract away perceptual and other task-specific factors that may otherwise lead to spurious correlations and other complexities not related to the central question of our work, and reduce task complexity to a number of sequential computation steps (referred to as hops) required to solve the task. This task, **conditional PVR** (C-PVR), is an extension of the pointer value retrieval task (PVR) (Zhang et al., 2021). We employ a symbolic sequential version of the task and introduce conditional multi-hop, meaning the number of steps required to solve an example is unpredictable until processed. In order to tackle this class of problems involving sequential reasoning steps, we extend the depth-adaptive framework of Universal Transformer (Dehghani et al., 2019) through the use of hyper-modules (Ha et al., 2017), resulting in what we call a **Hyper-UT** model. Hyper module is a hyper network (Ha et al., 2017) that composes a set of new weights via linear interpolation of weights from a bank of weight embeddings at each layer. Instead of sharing parameters across layers the hyper-network and the pool of weight embeddings are shared. This enables increasing the capacity of the model through increasing the number of modules while still benefiting from inductive biase of parameter sharing, e.g., modular reuse.

Interestingly, we find that the inductive biases of adaptivity and modularity encoded in Hyper-UT are not only helpful in iterative reasoning tasks like C-PVR, but also in standard System-1 problems (Bengio et al., 2021) like Imagenet1K classification. In particular, we show that introducing the hyper-module component remedies the capacity problem of the Universal Transformer, and with this architecture, we can achieve the same level of performance on Imagenet1K benchmark while being more compute efficient.

As a summary, in this paper we make the following contributions:

- Introduce a new task to probe the capability of models to generalize across reasoning steps when the total number of steps is undetermined;
- We explore the interplay between adaptive depth mechanism and modularity and how they can synergize for efficient generalization in the context of example complexity;
- Examine the generality of the improvements in accuracy and efficiency of adaptive-depth, modular transformers on a standard image classification task;

## 2 Multi-Step Reasoning

Multi-step reasoning is a process where a system breaks down the task into a sequence of steps and each step builds upon the previous ones. This type of problem solving is crucial in cases where there is ambiguity in the solution space, i.e., where the sequence of steps cannot be determined immediately. Scientific research is a prominent example of tasks that require such processes for forming a hypothesis and iteratively adjusting it based on the observations until a conclusion can be reached. A building block to enable a system to learn and execute this type of multi-step reasoning processes, is the capability to learn tasks that require iterations of dynamic length. We introduce a synthetic task, Conditional Pointer Value Retrieval (C-PVR) to probe this in neural networks (NNs).

### 2.1 CONDITIONAL POINTER VALUE RETRIEVAL

We introduce C-PVR to investigate the capability and limits of NNs in performing tasks that require generalization over the number of sequential steps. The number of sequential steps can be viewed as a notion of complexity of examples. Generalization to length is a special case of this and it has been studied in both synthetic settings (Zhang et al., 2022; Abbe et al., 2023; Jelassi et al., 2023) and large scale setting (Anil et al., 2022). However, the number of sequential computation steps needed to solve the examples can be independent of their length (e.g., we can have shorter sequences that require more steps or longer sequences that require less). Additionally, we are interested in cases where there are no external or immediate clues for the model to know the sequence of steps or its length before solving the example.

Zhang et al. (2021) introduced PVR as a benchmark to study if transformers are capable of human-style reasoning. In PVR tasks, a specific portion of the input serves as a pointer, offering instructions that pertain to a particular input location. This location is then processed to generate the output. In a basic symbolic version of PVR tasks, inputs are sequences of numbers, where the first element in the sequence is the pointer to another element in the sequence, and the output value is computed based on the pointed element. In the simplest case, the output is the value where the pointer is pointing. In this task, the second step is conditioned on the output of the first step. A multi-hop version of this task can be constructed by allowing all the elements in the sequence to be interpreted as both pointers and values (directly or by applying some transformations) and specifying the number of hops at the example level or at the data-set level. We focus on the symbolic sequential version of this task and extend the notion of multi-hop PVR such that the number of steps is not given but depends on the input sequence. In simple terms, the task is defined as "Continue the retrieval steps recursively until a certain condition is met".

In designing the C-PVR task we follow two goals: (1) The minimum number of sequential steps should vary for different examples; and (2) The number of steps should not be identifiable before processing the sequence. The task comprises of the following components: (a) the input, which is a sequence of length $L$ of integers in the range of $(1, K)$, (b) the output which is an integer in the range of $(1, K)$, (c) the operation for getting the pointer from a given element, (d) the operation for getting the value from a given element, and (e) the halting condition.

In our experiments we set the halting condition to be get_pointer$(a[i]) >$ i (keep retrieving until we are not moving forward in the sequence). We compare two variants of our task: (1) **C-PVR (plain)** where both pointer and value functions are identity and $K = L$, and (2) **C-PVR (modulus)** where the value function is get_value$(x) = x\%L$. Figure 1 shows examples of simple PVR and C-PVR (plain).

## 3 ADAPTIVE AND MODULAR TRANSFORMERS

**Adaptive Compute** Given the aforementioned varying difficulty of examples many learning problems must tackle, a possible approach to address it is augmenting the model with an ability to dynamically allocate computational budget to different examples (or part of an example) accordingly. In this approach, a module, e.g., an LSTM in Graves (2016) or a transformer layer in Dehghani et al. (2019), is provided with an additional module (often a fully connected network) predicting a halting score that determines how many times this module will be repeated on the current input before a scoring threshold is met. To encourage the model to limit the amount of computation it will use, Graves (2016) proposes to not only truncate the maximum amount of updates, i.e., the number of repeated computations, but regularize the model using an additional Ponder Cost (for more details, see Graves, 2016).

**Modularity** Modularity is typically enforced in neural networks by breaking them into smaller sub-networks which can then be composed differently to solve different examples. Modular neural networks are interesting both for their potential generalization ability through a systematic composition of modules, their efficiency by reducing the cost of each compute step while allowing the total capacity of the model to grow, and sometimes their interpretability. There are different ways to make neural networks modular, e.g., mixture of expert transformer models (Fedus et al., 2022; Lepikhin et al., 2021), using hyper networks (Ha et al., 2017), modulating the activations (Perez et al., 2018), etc. We explore the following modularity techniques:

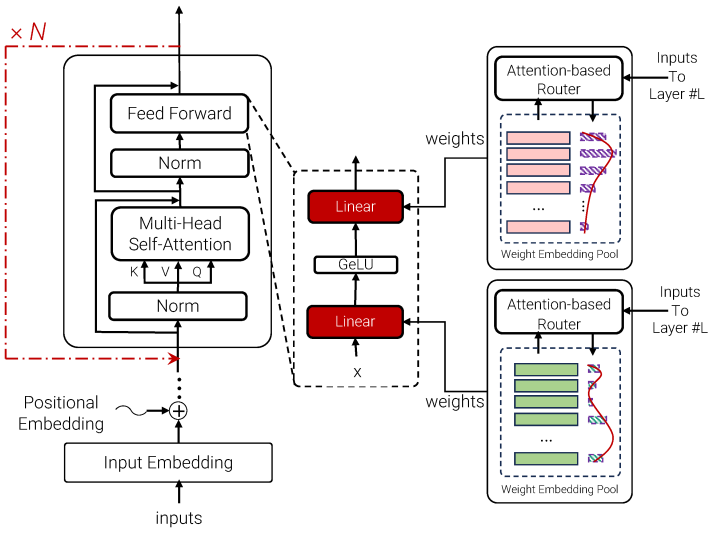


Figure 2: **Overview of the Hyper-UT Architecture.** In each Hyper-UT layer, the linear projection weights are dynamically generated by the "Module Selector", based on the input for each layer. Although the illustration primarily highlights the linear layers in the FeedForward block for simplicity, the same principle applies to the query, key, value, and output projections within the attention block.

- FiLM layers (Perez et al., 2018): The activations are modulated by shifting and scalar values that are predicted with an MLP block given varying input representations at each layer.
- Perceiver (Jaegle et al., 2021) style layers: At each step, there is an additional cross attention to a set of latent vectors where each vector can be interpreted as a module.
- Hyper-Module: The values of parameters of each layer are predicted by a hyper-module (that is shared across layers) given a representation of the input at each layer (Figure 2).

### 3.1 HYPER-UT

Universal transformers have been shown to generalize better than transformers in certain settings (Dehghani et al., 2019; Csordás et al., 2021); however, when applied to tasks which need higher capacity they often cannot match the performance of their non-adaptive counterparts (Xue et al., 2023). There are two main components for universal transformers: the adaptive depth module and the parameter sharing in depth. If the goal for applying adaptive depth is merely compute efficiency, one does not need to share parameters across depth and sacrifice capacity. But if the goal is to improve generalization, we want to jointly benefit from inductive biases of both adaptivity and parameter sharing (modular reuse). More importantly, if the goal is generalization, we expect models to be able to deal with examples of higher complexity at inference time.

The challenge of sharing parameters across the layers of the transformer model is that the only knob to increase the capacity of the model in terms of the number of parameters is to increase the width of the shared layer. This results in more flops per step or layer.

By incorporating the hyper-module into the UT architecture, we remedy this problem. While there is still an overhead for the modularization, we hypothesize that the capacity of the model can grow faster than the extra cost of compute per layer. Additionally, modularity can also be viewed as a source of inductive bias that can enhance both efficiency and generalization. Our experiments in Section 4.3 verify this to some extent.

In Hyper-UT, we replace the dense layers in the self-attention module and the MLP after the self-attention in the transformer block with hyper-modules (the hyper-modules for different components of the layers are not shared but they are shared across layers). At every layer, the hyper-module composes a set of new weights via linear interpolation of weights from a bank of weight embeddings, conditioned on some representation of the input at that layer. To avoid having to store weight embeddings of the size of the parameters, there is a projection layer that maps the module embedding to a vector which is the size of the weights needed to be predicted.

The Hyper-Module consists of the following components.

- Weight embedding pool: Each embedding in the weight embedding pool can be loosely interpreted as a module or expert. We split each embedding into key and value where the

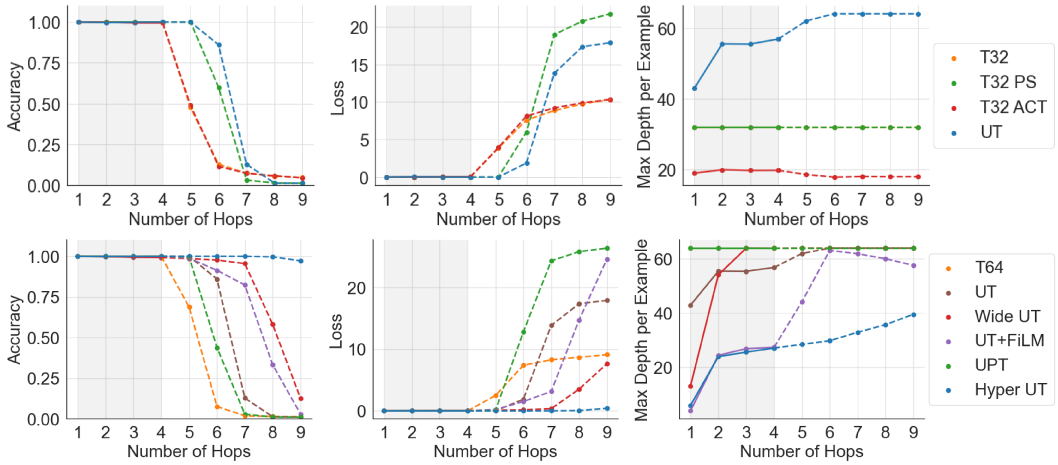


Figure 3: Performance of transformer variants on C-PVR (modulus) on test sets with different number of hops. The shaded area is for in-distribution and the un-shaded area is out-of-distribution. In the first row we examine the impact of adaptive depth and parameter sharing separately and when combined together. In the second row, we compare how different modularisation techniques improve the performance and efficiency of a model with adaptive depth. Note that in the first row, in the left most plot, T32 and T32 ACT overlap and in the right most plot, T32 and T32 PS overlap making T32 invisible in these two plots.

key is used for module selection and the values are linearly interpolated based on a score generated by the router to compute the target module embedding.

- Attention-based router: The attention-based router predicts the scores for each module to compute the embedding of the target module based on the attention between the representation of the input at the current layer and the keys in the weight embedding pool.
- Weight generator: A linear layer that projects the weight embedding to the weights of the target module.

Since the weight generator needs to predict a vector with the size of the full matrix of the dense layers, it can become a bottleneck as we scale up the model size and increase the dimensions of these dense layers. To address this challenge, a potential solution is to predict factorized versions of the weight matrices. We defer the investigation of scalable implementations for the hyper-module to future research.

## 4 Experiments

We present empirical results to show how incorporating adaptive compute and modular compute simultaneously yields better generalization and efficiency.

First, we show on the C-PVR task that generalizing to more complex examples (examples that require a higher number of sequential steps), is not trivial for standard transformers. Under a setting where the diversity of the complexity of examples seen during training is limited, we investigate the effects of adaptivity and modularity on the performance of transformers on this task.

Second, we investigate how pre-training on a language modelling task can help with learning and generalizing on the C-PVR task, and how augmenting the model with an explicit scratch-pad impacts the results.

Third, we look into a standard image classification task where models with adaptive depth do not match the performance of their non-adaptive counterparts despite being more efficient (Xue et al., 2023). In this setting, we show that introducing modularity, as we do in Hyper-UT remedies this problem, with Hyper-UT matching the performance of standard ViT (Dosovitskiy et al., 2021) with less compute in terms of the number of layers.

### 4.1 C-PVR: EFFICIENT GENERALIZATION OVER EXAMPLE COMPLEXITY

In our experiments on the C-PVR task, we look into performance and efficiency of different adaptive and non-adaptive transformers with different modularization mechanisms. Our focus is on the generalization performance of these models with respect to the number of hops or the sequential computation steps needed to solve examples, as well as their efficiency in terms of correlation of the amount of compute and the complexity of examples.

We study the setting where the models are trained on examples of lower complexity (1-4) and evaluated on examples of higher complexity (5-9). Figure 3 shows the performance of the models on test sets with different numbers of hops in this setting. In this experiments, the sequence, $L$, is $100$, and the values of elements are in the range of $[1, 1000]$. In the train set we have around $450k$ of each number of hops ($\sim 2M$ training examples) and the size of the test sets are $50K$.

We compare transformers with the following mechanisms:

- Adaptive Computation Time (ACT) (Graves, 2016) as the mechanism for having adaptive depth. We apply ACT per token.
- Different ways of parameter sharing and modularisation across layers with or without ACT:
  - No parameter sharing: The main problem with this approach is that if the model has adaptive depth and needs to be deeper during inference time, it needs to rely on parameters that are not well trained: T12 (Transformer with 12 layers), T32, T64, T32 ACT (Adaptive depth Transformer with max 32 layers)
  - Plain parameter sharing: All parameters are shared across all layers. We have one block that is called in a loop $N$ times, where $N$ is the number of layers: UT (Universal Transformer (Dehghani et al., 2019)), T32 PS (32 layer transformer with parameter sharing)
  - Parameter sharing with additional Film layers (UT+FiLM).
  - Sharing latents across layers of a perceiver style Transformer (UPT).
  - Hyper-Modules: Parameters of the dense layers are predicted by hyper-networks, which can be shared across layers: Hyper-UT

Details of the hyper-parameters of these models are presented in Appendix C.

**When do models break** Figure 3 shows how the accuracy of the models trained on examples with a smaller number of hops (1-4) drops as we increase the number of hops in the test set (from 1 to 9). We observe that the performance of the transformer models with no adaptive and modular compute drops more rapidly. Increasing the number of layers of a fixed-depth transformer model does not lead to any significant improvement in accuracy despite the increase in capacity (Figure 7 in Appendix A). On the contrary, parameter sharing in depth improves the generalization performance significantly, despite reducing the capacity in terms of number of parameters. In this particular setting adding an ACT mechanism without parameter sharing does not impact performance at all but it improves efficiency in terms of number of layers. UT, incorporating both parameter sharing and ACT at the same time, generalizes better than a standard transformer and a transformer with only parameter sharing. Adding modularity to UT in different forms, further improves both performance and efficiency. Among all, Hyper-UT achieves the best overall results. Moreover, increasing the width of the UT model, also leads to better generalization in terms of accuracy, verifying that a major problem for UT to generalize on this task is the capacity.

**Compute Efficiency and Fair Allocation of Compute** To investigate the compute efficiency of the different transformers we try on the C-PVR task, first, we look into the average depth of the models. Figure 3 shows how the average number of layers applied on examples varies for the test sets that require different numbers of hops. We expect a model to be efficient during inference not only if in general, it achieves better accuracy with less compute, but also if it can allocate compute to different examples not uniformly but based on their complexity.

For non-adaptive transformers average depth is always equal to the maximum number of layers. To make the comparison meaningful we train multiple instances of non-adaptive transformers with different numbers of layers (Figure 7) in Appendix A. We observe, in Figure 3, that UT+FiLM and Hyper-UT are more fair and efficient in terms of allocating compute compared to the UT model (HyperUT being the most efficient). By more fair in allocating compute, we mean the number of

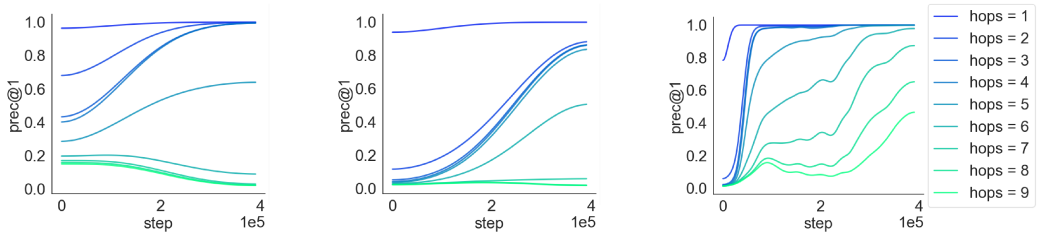


(a) Transfomer (32L) (b) UT (c) Hyper-UT

Figure 4: Implicit curriculum of transformers on C-PVR task. The y-axis is the accuracy and the x-axis is the training steps. This plot shows that transformers (a Hyper-UT instance) is learning examples in order of their complexity.

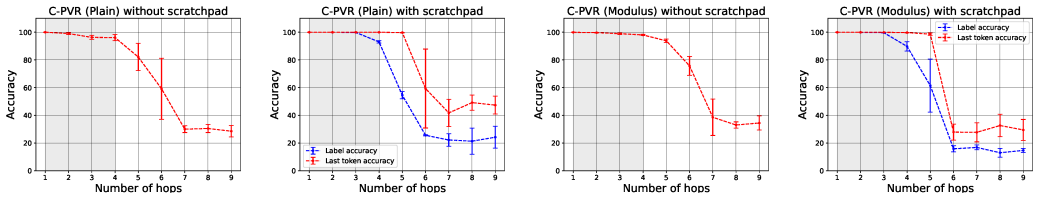


(a) (b) (c) (d)

Figure 5: Performance of pre-trained T5 models on C-PVR (plain), (a) and (b) and C-PVR (modulus), (c) and (d). To generate each data point on the plots, we did three rounds of training. The reported values are the mean while the spread is just standard deviation of the sample mean.

effective sequential compute steps (e.g., number of layers), is correlated with the complexity of the examples. It is intriguing that complementing adaptive depth with modularity leads the adaptive models to use fewer number of layers (comparing UT with UT+FiLM and HyperUT).

**Implicit Curriculum** As depicted in Figure 4, the order of examples learned is correlated with their complexity (with respect to the number of hops). This means the notion of complexity we employ here is aligned with the family of solutions these models are learning.

### 4.2 Pre-trained language models

One natural question to ask is whether pre-trained language models are capable of generalizing to higher number of hops than those seen during training on the C-PVR (plain) and C-PVR (modulus) tasks using scratch-pad. To examine this question, we fine-tuned small pretrained T5 models on C-PVR (plain) and C-PVR (modulus) with or without scratch-pad. Note that scratch-pad could in principle allow models to allocate different amounts of compute to each input instance and potentially help them generalize to higher number of hops at inference time.

Figure 5 shows accuracy on C-PVR (plain)/C-PVR (modulus) tasks versus number of hops in two settings with and without scratch-pad. The T5 models are fine-tuned on an equal-sized mixture of examples with number of hops ranging from 1 to 4 (shown as the grey area in the plots). The models are tested on all number of hops from 1 through 9 (1 through 4 hops would be the in-domain test cases, while 5 through 9 hops would constitute out of distribution in our setting). The training set sizes are 100K/160K for C-PVR (plain)/C-PVR (modulus) respectively, while test size was fixed at 5K. All models are trained to convergence (C-PVR (plain)/C-PVR (modulus) for 9/15 epochs). The input size is fixed at $L = 30$. For C-PVR (modulus), the maximum of array elements value is 300. The problem is formulated as text-to-text. For runs with scratch-pad, the target has the general format of LABEL # SCRATCH-PAD, where scratch-pad is the string representing the array values at intermediate hops including the the label at the end. The character # is just a separator for convenience.

As we can see in Figure 5, fine-tuned T5 shows non-trivial generalization (e.g, well beyond chance-level) on number of hops higher than 4, which diminishes as the number of hops increases. Including a scratch-pad enhances the results if the accuracy is computed based on the last token of the scratch-

incorporating adaptive compute introduce a form of recurrence in depth which can be a source of inductive bias that facilitates generalization in certain settings (Dehghani et al., 2019; Banino et al., 2021; Abnar et al., 2021; Csordás et al., 2021). RNNs, CNNs, and transformers are adaptive models that allocate compute to input based on the size of the input. While in principle all these models can deal with inputs of variable size their ability to generalize to inputs of varying shape is limited in their vanilla version (specific variants of RNNs, i.e., LSTMs are better at generalizing to unseen input sizes and in transformers using specific types of positional encodings could improve their ability to generalize to inputs of different size). However, the length/size of the input is not always a good proxy for its complexity. As a naive example, a given data point can be arbitrarily padded and resized but this does not increase its complexity. Graves (2016) introduced the concept of pondering and adaptive time compute and applied it to LSTMs. The main idea here is to allow the network to ponder as much as needed by having a module that predicts when to halt based on the current state, and add a regularization term to the objective to minimize the number of steps. Later, Dehghani et al.