

HINT: Hypernetwork Instruction Tuning for Efficient Zero- & Few-Shot Generalisation

Hamish Ivison ^α

Akshita Bhagia ^α

Yizhong Wang ^ω

Hannaneh Hajishirzi ^{αω}

Matthew Peters ^α

^αAllen Institute for AI

^ωPaul G. Allen School of Computer Science & Engineering, University of Washington
{hamishi, akshita, yizhongw, hanna, matthewp}@allenai.org

Abstract

Recent NLP models have shown the remarkable ability to effectively generalise ‘zero-shot’ to new tasks using only natural language instructions as guidance. However, many of these approaches suffer from high computational costs due to their reliance on concatenating lengthy instructions with every input example, resulting in costly reprocessing of the instruction. To avoid this, we introduce Hypernetworks for INstruction Tuning (HINT), which convert task instructions and examples into parameter-efficient modules inserted into an underlying model using a pretrained text encoder, eliminating the need to include instructions in the model input. The hypernetwork in HINT also produces an encoded instruction, which we concatenate with encoded inputs during decoding to further improve performance. HINT models outperform strong state-of-the-art baselines by over 10% when controlling for compute (measured in FLOPs). By converting instructions into modules, HINT models can effectively disregard the length of instructions and few-shot example inputs in terms of compute usage. As a result, HINT can enhance its performance by up to 25% by incorporating additional few-shot data, while utilizing only up to 5% more compute. This combines the strengths of parameter-efficient fine-tuning and in-context learning. We release our code publicly¹.

1 Introduction

Large pretrained language models have demonstrated a striking ability to perform new tasks through the use of *in-context* examples or instructions alone (Brown et al., 2020), or after training on input instances augmented with instructions (Weller et al., 2020; Mishra et al., 2022; Sanh et al., 2022; Wei et al., 2022; Chung et al., 2022; Wang et al., 2022b). This ability allows a single model to

¹Our code is available at:
<https://github.com/allenai/hyper-task-descriptions>.

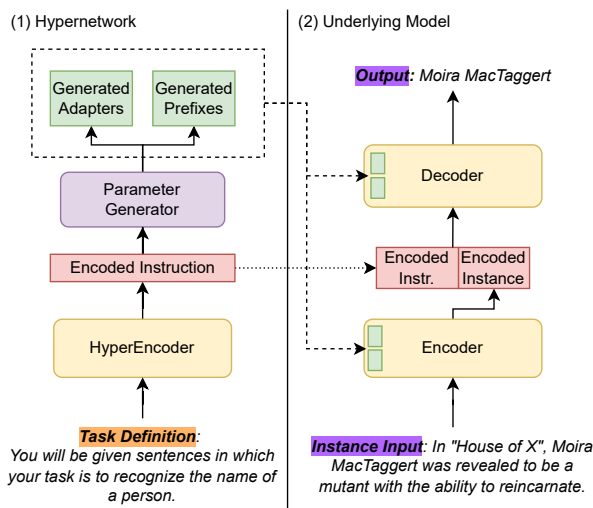


Figure 1: Overview of HINT. (1) We feed an instruction into a HyperEncoder to produce an encoded instruction, use it to generate prefix and adapter weights, and then insert them into the underlying model. (2) We run the underlying model encoder as usual and optionally concatenate the encoded input with the previously encoded instruction, before running the underlying model decoder to generate the answer. We only use the hypernetwork once per task.

adapt to many tasks where training data is difficult to collect or task-specific fine-tuning is impractical (i.e., ‘zero-shot’ settings): models trained on instructions need only a single instruction to achieve non-trivial performance on the task at hand. The most common method to achieve this zero-shot ability is to meta-train the model with task instructions concatenated with every input, allowing the model to learn to associate instructions with tasks. While empirically highly successful, this is inefficient and requires reprocessing lengthy task instructions and any additional task data (e.g., few-shot examples) with every input example.

In this paper, we introduce Hypernetworks² for Instruction Tuning (HINT), which directly

²Hypernetworks are neural networks trained to generate neural networks (Ha et al., 2017).

generate task-specific parameter-efficient modules given only an instruction, combining the benefits of instruction-based learning with parameter-efficient modules. HINT models convert instructions and other task data (e.g., few-shot examples) into efficient modules within a pretrained language model, enabling cheaper inference and better compute scaling with few-shot data for an underlying instruction-based meta-learning approach. Additionally, fusing hypernetwork-encoded instructions with the encoded input at the underlying model decoder greatly improves the performance while using minimal extra compute. An important benefit of HINT is that it processes instructions and other task information only once, making the compute used by our method almost independent of the amount of task data available, unlike both regular finetuning and input concatenation-based approaches (see Figure 3).

We find that our hypernetwork-based approach (‘HINT’), is able to achieve similar performance to baselines that receive the full instruction with every input example while using significantly less compute (as measured by FLOPs), due to the greatly reduced input length. When controlling for inference budget, we find that **HINT models outperform strong baselines in zero- and few-shot settings**. This validates our assumption that we can significantly reduce inference costs by avoiding reprocessing the instruction with every input, and instead saving it for repeated use. Furthermore, we find that including additional few-shot information alongside task instructions significantly improves HINT model performance while using minimal additional compute during inference. Ultimately, our work pushes towards directly generating cheap, customised models from task data, without requiring any expensive task-specific finetuning.

In summary, our findings are:

- We introduce HINT models, which make use of a text-conditioned hypernetwork to generate parameter-efficient modules based on task descriptions and few-shot examples.
- HINT models, by reducing input lengths, are able to achieve similar performance to strong full-input baselines while reducing inference cost (measured in FLOPs) by up to $4\times$.
- As the compute used by HINT models is effectively independent of the length of the instruction and amount of few-shot data provided

with the instruction, HINT models provided with additional few-shot data simultaneously outperform and use up to $4\times$ fewer FLOPs than baselines without few-shot data.

- HINT models outperform strong decoder-only baselines. While decoder-only models allow for input caching, we find that instruction-tuned GPT-2 models significantly underperform HINT models (8-9 point difference), matching prior work suggesting that encoder-decoder models work better for instruction-tuning (Wang et al., 2022a; Iyer et al., 2022).

2 Related Work

Instruction Following Further finetuning large pretrained language models on instructions has been found to greatly improve zero-shot generalisation, as the finetuned model learns to make use of the instructions to perform the given task (Weller et al., 2020; Wei et al., 2022; Mishra et al., 2022; Chung et al., 2022; Wang et al., 2022b). Additionally, Sanh et al. (2022) found that training models on multiple prompts per task also resulted in improved performance, suggesting that further increasing prompt diversity aids generalisation, even when using the same pool of underlying tasks. The majority of these popular instruction-tuning approaches involve concatenating the instruction with the input directly and training a text-to-text model on these combined inputs. As the instruction can be as long as, if not longer, than the input³, this can greatly increase the computation needed to process inputs compared to task-specific models.

In-Context Learning Similar to instruction-based models, in-context learning (Brown et al., 2020), where example instances are used in place of or in addition to instructions, also requires extremely long and expensive-to-process inputs for every test example, with Liu et al. (2022) showing that parameter-efficient finetuning (PEFT) can be cheaper and more effective when dealing with many test examples. In this work, we propose a halfway step between PEFT and instruction concatenation, where we train a model to predict parameter-efficient modules based on instructions, avoiding the few-shot training required by Liu et al. (2022) while also avoiding repeatedly processing lengthy inputs.

³As is the case for Super-Natural Instructions, see Appendix A.1.

Hypernetworks In NLP Hypernetworks (Ha et al., 2017; Schmidhuber, 1992) in NLP have recently gained popularity in multitask and multilingual setups due to their ability to softly share parameters while avoiding negative interference through the use of shared parameter generation module. Several approaches (Karimi Mahabadi et al., 2021; Tay et al., 2021; He et al., 2022b) learn per-task embeddings along with a shared hypernetwork to generate task-specific adapter or prefix modules. This means making the model perform new tasks requires at least few-shot learning to learn a task embedding. Recent work has explored using text-conditioned hypernetworks for parameter-efficient multitasking (Iverson and Peters, 2022) or improving out-of-domain generalisation (Volk et al., 2022), removing the need to train task-specific embeddings. Hypernetwork-based methods have also been highly successful in multilingual settings, where generating language-specific models via shared hypernetworks often results in improved performance across various tasks (Platanios et al., 2018; Baziotis et al., 2022; Ustun et al., 2022, *inter alia*)

Our work primarily builds on Ye and Ren (2021), which explored generating adapters from task descriptions. We expand their approach to larger models and datasets and find that pretraining and a significantly different hypernetwork architecture are important for achieving strong performance.

Our work is also similar to the concurrently developed Phang et al. (2022) and Deb et al. (2022), both of which examine how well hypernetwork-based meta-learning can improve model performance in zero- and few-shot settings. Deb et al. (2022) examine hypernetworks and model-agnostic meta-learning for instruction-finetuning and find that they can yield improved performance on difficult unseen tasks in Super-Natural Instructions. However, they still struggle to achieve overall good zero-shot performance and do not investigate eliminating task descriptions from the model input itself. Phang et al. (2022) find that training a hypernetwork to produce model adaptations provides an initialisation better than pretraining for parameter-efficient adaptations and that this initialisation improves with more few-shot examples provided to the hypernetwork. They also explore eliminating the instruction from the underlying model input but find this severely underperforms baseline approaches. We have similar findings, but find

that our novel hypernetwork design and use of instruction fusion closes the gap with baseline approaches. We also perform further analysis of the hypernetwork-based models and show that when controlling for inference compute budgets, our hypernetwork-based model still outperforms strong baselines.

3 HINT Model

Here, we introduce the main elements of our HINT model. The model has two core parts: the **hypernetwork**, which takes in text instructions and outputs parameter-efficient modules, and the **underlying model**, into which we insert the generated parameter weights. The underlying model is simply an encoder-decoder⁴ transformer model (Vaswani et al., 2017) with additional parameter-efficient adaptations inserted in, while the hypernetwork has a more complex architecture which we describe below. Figure 1 provides a visual overview.

3.1 Hypernetwork

The first step in our model is to make use of a hypernetwork to convert an instruction to parameter-efficient modules. Our hypernetwork consists of three core elements: an **encoder** (or ‘**hyperencoder**’) to transform instruction and few-shot text into continuous (contextual) representations, saving the encoded instructions for **instruction fusion** during decoding, and a **parameter generator** to then convert these embeddings into the parameter efficient modules.

HyperEncoder To encode our text, we use a pre-trained language model encoder. We initially experimented with using different encoder configurations, and find that re-using the encoder from the underlying model we wish to augment works well, and tying the hypernetwork and underlying encoder model weights works best.

Instruction Fusion We save the instruction representations produced by the hyperencoder and allow the decoder of the underlying model access to them by concatenating them with input examples during inference and training. This is inspired by the fusion-in-decoder method used in open-domain QA (Izacard and Grave, 2021).

⁴We use encoder-decoder models as they generally outperform decoder-only models for zero-shot generalisation – see Section 5.1.

3.1.1 Efficient Parameter Generators

Parameter Generators Our generator design consists of two parts. First, we use a trainable set of embeddings and perform multi-head cross-attention with the encoded instruction and these embeddings. Each embedding represents a unique column or token in each parameter-efficient module (e.g., prefix, adapter - see below) for each layer. This allows us to effectively collect the information required for different parameters in different embeddings via cross-attention with the instruction:

$$\begin{aligned} \text{embed} &= [\alpha_{e_1^1}, \alpha_{e_2^1}, \dots, \alpha_{e_1^2}, \alpha_{e_2^2}, \dots, \pi_{e_1^1}, \dots] \\ \text{embed}' &= \text{Cross-Attention}(\text{embed}, \text{instr.}) \end{aligned}$$

Where $\alpha_{e_1^1}$ refers to an embedding we will use as the first column of the first layer adapter weight, $\alpha_{e_2^1}$ is the second column, $\alpha_{e_1^2}$ is the first column for the second layer adapter, $\pi_{e_1^1}$ is the first token of the first layer prefix, etc. We then take the subset of the embedding representing all columns/tokens for a particular model adaptation and pass it through a two-layer MLP to generate parameters. We use a unique network for each adaptation and share between layers (i.e. one network for prefixes for all layers, one for all adapter weights for all layers, etc.).

$$\begin{aligned} \text{Adapter}_1 &= \text{reshape}[\text{MLP}_a(\alpha'_{e_1^1}); \text{MLP}_a(\alpha'_{e_2^1}); \dots] \\ \text{Prefix}_1 &= \text{reshape}[\text{MLP}_p(\pi'_{e_1^1}); \text{MLP}_p(\pi'_{e_2^1}); \dots] \end{aligned}$$

Where Adapter_1 and Prefix_1 are the first layer adapter and prefix, respectively.

Generated Parameters We generate two types of parameter-efficient modules: **adapters** (Houlsby et al., 2019) and **prefixes** (Li and Liang, 2021). **Adapters** (Houlsby et al., 2019) are small bottleneck networks inserted into a transformer model. We follow He et al. (2022a) in placing our adapters parallel to the feed-forward layer:

$$\mathbf{x} = \text{FFN}(\mathbf{x}) + f_1(\text{GELU}(f_2(\mathbf{x}))) \quad (1)$$

Where f_1, f_2 are linear layers that project an input \mathbf{x} to a small bottleneck size n_a and then back up to the hidden size of the model respectively. **Prefixes** (Li and Liang, 2021) are short continuous sequences concatenated with the key and values in the self- and cross-attention modules in every layer of the underlying model.




Figure 2: Overview of our proposed pretraining scheme.

Scaling Down Parameters A naïve hypernetwork implementation may suffer from poor scaling with the size of the parameter-efficient modules. Consider a case where we wish to convert a single embedding of size n_e to an adapter weight matrix of size $n_d \times n_a$ (the model hidden dimension size by adapter bottleneck size). Our hypernetwork generator will have $n_d * n_e * n_a$ parameters, and so increasing the adapter bottleneck n_a quickly becomes extremely expensive, especially if n_d is large - as is the case for large language models. We address this by decomposing the adapter weight into columns and *assigning an embedding per column*. Thus, our hypernetwork now has to convert a sequence of embeddings with size $n_a \times n_e$ to an adapter weight of size $n_a \times n_d$, meaning that the network only needs $n_e * n_d$ parameters. This means that the size of our parameter-efficient modules is independent of the size of the hypernetwork, and we can effectively scale the size of our adapters or number of prefixes *without extreme parameter blowup*. Note that we set $n_e = n_d$ in our experiments for simplicity.

3.2 Underlying Model

Once our hypernetwork has produced a set of parameter-efficient modules, we then insert these into our underlying network, and can then perform training and inference as normal. The underlying model can be any pretrained encoder-decoder model that works with our parameter-efficient modules. We make use of T5 (Raffel et al., 2022) as our underlying model in our experiments.

3.3 Training and Inference

Hypernetwork Pretraining To help better generalization, we pretrain the hypernetwork on a large corpus (C4; Raffel et al., 2022) before finetuning on multitask prompted datasets. Given a single input string, we split our input string into random-length chunks a, b, c , and feed a to the hypernetwork, b to the main network, and predict c . This resembles the input in used in instruction finetuning (as the instruction precedes the input in the default prompt used for Super-Natural Instructions). We fully fine-

tune all parameters during pretraining.

Training HINT HINT training looks similar to pretraining, except we replace a with the task instruction (and any few-shot examples), b with the main input, and c with the gold generation. We used mixed-task batches such that a unique adaptor and prefix set is generated for every input in each batch.

This means, for every batch, we first generate a set of adaptors, prefixes, and encoded instructions from a batch of tasks using the hypernetwork. The adaptors and prefixes are placed within the underlying model to act as parameter-efficient modules (i.e., insert them into the model), and the encoded instruction is concatenated with encoded inputs during decoding. We then perform a forward pass of the underlying model with the inputs associated with each task in the batch and perform back-propagation using cross-entropy loss as standard for text-to-text models. As we fully finetune all parameters, the parameter generator will produce different weights for the same task inputs after a gradient step, meaning that we have to rerun the hypernetwork for every batch. This means that HINT requires more compute to train than a baseline transformer - although it provides significant compute reductions during inference, as we will see.

Inference The inference process is similar to training, but we do not use mixed-batch inputs: instead, we generate the parameters for one task, insert them into the underlying model, and then process all test-time inputs for that task. This prevents redundant processing of the instruction.

We also consider the cost of HINT models during inference. We consider a case where we have to process n samples from a single task. Assume each sample has length i and the task instruction has length t . We will ignore the cost of processing (typically short) output sequences. Following prior work (Kaplan et al., 2020; Liu et al., 2022), we use FLOPs as an estimate of the amount of compute required to run particular models and estimate that processing a token with an encoder-decoder model takes N FLOPs to process a single token, where N is the total number of model parameters.

In this scenario, a standard instruction-trained model which concatenates every input with the instruction (e.g., Tk-Instruct) uses $Nn(t + i)$ FLOPs to process all examples. Meanwhile, HINT models

process the task instruction only once and so use roughly $N(t + ni)$ FLOPs⁵. This makes clear that HINT models (a) **scale better with more same-task inference examples** than input concatenation approaches (increasing n), and (b) **require relatively few extra FLOPs to process long instructions** (large t), allowing them to benefit from adding more few-shot examples without incurring significant compute increases.

4 Experimental Details

We evaluate our approach on two popular instruction-based datasets: **Super-Natural Instructions (SNI)** (Wang et al., 2022b) and the **T0 split of P3** (Sanh et al., 2022; Bach et al., 2022). We use t5x and seqio (Roberts et al., 2022) to handle data preprocessing and model training. We use T5 v1.1 + LM adaptation (Lester et al., 2021) as our base models, using the 3B size unless otherwise stated. Unless otherwise stated, the hypernetwork generates prefixes of length 30 and adaptors with a bottleneck size of 512, matching the sizes recommended by He et al. (2022a). We use the Adafactor optimizer (Shazeer and Stern, 2018) with a constant learning rate of 0.001. Unless otherwise stated, we report results from single runs.

Pretraining We pretrain all models for 10,000 steps (7,000 for 11B size models) using C4 (Raffel et al., 2022) with a batch size of 1,024 samples and sequences of length 512.

Super-Natural Instructions (SNI) For SNI, we examine two settings: providing the hypernetwork with the task definition and the underlying network with the instance input only (**‘Def’**), and providing the hypernetwork with the task definition and two few-shot task examples (**‘Def + 2 Pos.’**). To train, we finetune our pretrained HINT models for 1,000 steps with a batch size of 1,024, with a maximum sequence length of 1,024 for both the underlying model and the hypernetwork input. We then evaluate the final checkpoint on the test split of SNI, which is a set of 119 unseen tasks. We use v2.6 of Super-Natural Instructions.

P3 For P3, we explore two settings: (a) **‘joint’**, where we give the hypernetwork a templated form of the prompt with instance information removed and give the underlying model the full prompted

⁵When reporting FLOPs, we use a more detailed formula described in Appendix C.1 that takes into account extra (albeit small) hypernetwork costs.




Figure 3: (Left) SNI RougeL against FLOPs for varying model sizes using task definitions with two positive examples; (Centre) SNI RougeL against FLOPs for differing numbers of additional few-shot information; (Right) FLOPs against instruction and task data length (in number of tokens). FLOPs calculations are based on processing 100 examples from the same task during inference.

input. (b) ‘split’, where we give the hypernetwork the templated prompt without instance information and give the underlying model only the instance information *without* the prompt. In both cases, we fully finetune our model for 10,000 steps with a batch size of 2,048. We use a maximum input length of 1,024 for the underlying model and 512 for the hypernetwork. We train and evaluate on the same tasks and splits as T0 (Sanh et al., 2022).

Baselines We primarily compare against **T0** and **Tk-Instruct**, models fully-finetuned on P3 and SNI respectively with all task information concatenated with the input. We replicate these models, matching the finetuning settings used for HINT models, and find that our replications significantly outperform previously reported results, making these baselines extremely strong. We note where results are our replications or reported from prior work. We additionally compare against ‘**X + PEFT**’, the prior models with adapters and prefixes added in before finetuning, **HyperTune** (Phang et al., 2022), a concurrent work that primarily makes use of a pretrained hypernetwork but without instruction fusion, **Hypter** (Ye and Ren, 2021), a prior hypernetwork-based model that does not use pre-training, **GPT-2** (Radford et al., 2019), a strong decoder-only model, which we fully finetune, and **OPT** (Zhang et al., 2022), another strong decoder-only model, which we also fully finetune.

5 Model Performance and Efficiency

5.1 Super-Natural Instructions

We report the performance and inference costs of HINT models and baselines in Table 1 and Figure 3 and find that:

HINT models outperform baselines when

FLOPs-matched. As seen in Figure 3 (left), when FLOPs-matched, HINT models outperform Tk-Instruct, a strong baseline that fully concatenates the instruction with every input. This holds for both ‘Def’ and ‘Def + 2 Pos.’ settings.

HINT models are up to 4× more efficient than similarly-sized baselines. we find that HINT models use 2–4x fewer FLOPs than similarly-sized state-of-the-art Tk-Instruct baselines (Table 1). While other hypernetwork-based models are able to achieve similar compute savings, their performance is significantly worse than HINT (≥ 8 points). HINT has similar cost to a model trained without including instructions in the input (‘No-Instruct’), while performing over 30 points better.

HINT models improve performance with few-shot examples, but do not cost more FLOPs. When introducing additional few-shot data (‘Def + 2 Pos.’), HINT models improve dramatically (5–8 points) but the compute used barely increases (Figure 3, centre), as HINT models only need to encode the task data (instruction and few-shot examples) once per task. In contrast, while Tk-Instruct similarly improves with few-shot examples, the compute needed during inference increases dramatically, usually costing around 1.5× more. Overall, we find that HINT models require much less compute to deal with longer instruction and few-shot data inputs than Tk-Instruct (Figure 3, right).

HINT models outperform a strong decoder-only baseline. HINT significantly outperforms GPT-2 and OPT-13B, in line with prior work that shows encoder-decoder models often significantly outperform even much larger decoder-only equivalents (Wang et al., 2022a; Iyer et al., 2022). In particular, 11B-size HINT outperforms OPT-13B by 5 points or more despite using a similar number

Model	Model Size	Def		Def + 2 Pos.	
		RougeL	Rel. FLOPs	RougeL	Rel. FLOPs
Tk-Instruct (our replication)	250 mil.	35.3	×1.0	42.9	×1.5
Tk-Instruct (Wang et al., 2022b)	250 mil.	-	-	42.1	×1.5
Tk-Instruct + PEFT	250 mil.	33.3	×1.1	42.9	×1.6
Hypter (our replication)	250 mil.	12.1	×0.4	10.6	×0.4
HINT (ours)	250 mil.	33.3	× 0.4	41.8	× 0.4
Tk-Instruct (our replication)	3B	48.9	×12.0	56.6	×17.9
Tk-Instruct (Wang et al., 2022b)	3B	45.0	×12.0	54.3	×17.9
Tk-Instruct + PEFT	3B	49.8	×12.4	56.2	×18.5
No-Instruct	3B	12.4	× 3.9	-	-
GPT-2 XL	1.5B	38.2	×4.1	45.3	× 4.2
Hypter (our replication)	3B	16.8	×4.3	14.2	×4.4
HyperTune (Phang et al., 2022)	3B	38.9	×4.1	48.6	×4.3
HINT (ours)	3B	47.2	×4.5	53.2	×4.6
Tk-Instruct (our replication)	11B	53.6	×44.0	60.5	×65.7
Tk-Instruct (Wang et al., 2022b)	11B	-	-	62.0	×65.7
Tk-Instruct + PEFT	11B	54.6	×44.0	60.3	×65.7
OPT-13B	13B	44.8	×15.9	51.5	× 16.4
Hypter (our replication)	11B	15.5	× 15.3	13.4	×15.7
HINT (ours)	11B	51.1	×16.1	56.4	×16.5

Table 1: Super-Natural Instructions RougeL and relative number of FLOPs used when given task definition only (‘Def’), and task definition along with 2 labelled examples (‘Def + 2 Pos.’). Where noted, results are taken directly from Phang et al. (2022) and Wang et al. (2022b). Relative FLOPs cost is calculated relative to the base-size Tk-Instruct with task definition only. We calculate the values using the number of FLOPs required to process 1 task with 100 examples for each model. Model size is given by the number of parameters.

of FLOPs. This highlights the utility of improving efficiency for encoder-decoder-based models. We also note that caching key/value attention pairs, the simplest way to reduce inference costs with decoder-only models, scales worse than HINT. The size of cached key/value pairs for GPT-2 is $\propto lds$, where l is the number of layers, d is the size of the model hidden dimension, and s is the cached sequence length. In contrast, the size of the saved PEFT parameters for HINT is $\propto sd + ld$, which scales better with respect to sequence length (larger s) and model size (larger d, l)⁶.

5.2 P3

We report results on the T0 evaluation set in Table 2, with full results in Appendix B. We find that:

Our T0-3B replication significantly outperforms the results reported by Sanh et al. (2022). This matches prior suggestions that T0 is under-trained (Phang et al., 2022; Wu et al., 2022). We provide further details in Appendix E.

HINT outperforms hypernetwork baselines. The HINT model consistently outperforms Hypter, a prior hypernetwork-based approach, and learns to make use of the P3 prompts as evidenced by

⁶We provide details for these calculations in Appendix C.2.

Model	Avg	Rel. FLOPs
T0-3B	54.9	×1.0
T0-3B (our replication)	64.4	×1.0
T0-3B + PEFT	65.5	×1.0
No Prompt	57.5	× 0.8
Hypter (Joint)	64.6	×1.0
HINT (Joint)	65.4	×1.1
Hypter (Split)	56.2	× 0.8
HINT (Split)	60.3	× 0.8

Table 2: Avg performance over T0 evaluation tasks after training on the T0 P3 train set. FLOPs are calculated assuming we are processing 100 examples of a single task. The ‘Joint’ and ‘Split’ HINT variants refer to the two input formats for P3 described in Section 4.

its improved performance over a baseline model trained without prompts (‘No Prompt’).

HINT remains cheaper than T0 for inference. HINT uses significantly less flops than T0-3B, albeit with smaller savings compared to SNI, likely due to the different style of prompts: P3 prompts tend to be shorter, and interleave task inputs (e.g. ‘Does <sentence 1> entail <sentence 2>?’). Despite this, HINT still provides reasonable FLOPs savings. We suggest that the performance of HINT could be greatly improved by leveraging additional few-

# Shots	T0-3B (our repl.)	HINT (split)	HINT FLOPs Reduction
1	66.4	66.4	$\times 2.3$
2	67.1	66.6	$\times 3.2$
4	67.1	67.2	$\times 5.1$
5	67.9	67.1	$\times 6.0$

Table 3: P3 performance with differing numbers of few-shot examples using 3B size models and the FLOPs reduction when using HINT instead of T0-3B for that number of shots. In few-shot settings, HINT always remains within 1 point of T0-3B despite the greatly reduced FLOPs cost.

Model	Pretraining	SNI RougeL
HINT	None	44.0
HINT	Ours	46.3
HINT	CACLM	45.8
HINT - No Instr. Fus.	None	27.4
HINT - No Instr. Fus.	Ours	32.1
HINT - No Instr. Fus.	CACLM	30.4

Table 4: SNI performance for HINT models with and without instruction fusion after 10,000 steps of the given pretraining scheme and 1,000 steps of finetuning on SNI. CACLM is the pretraining scheme proposed by Phang et al. (2022).

shot information, further exploiting the efficiency of HINT models in encoding task data.

We investigate if HINT models can provide benefits even when the input and instruction are concatenated through training and evaluating in the ‘joint’ setting of P3, and find that HINT performs similarly to T0-3B with additional parameter-efficient modules, which suggests that the hypernetwork is unable to improve on the baseline model through additional customisation, and so is primarily useful as a mechanism for reducing inference costs and cheaply incorporating few-shot data.

HINT performs similarly to the baseline in few-shot settings. In Table 3, we show that HINT remains within 1 point performance of T0 in few-shot settings, despite the large reductions in FLOPs cost, using up to $6\times$ fewer FLOPs. This makes HINT especially useful in few-shot scenarios.

6 Analysis

6.1 Pretraining

We compare using no pretraining, our pretraining scheme, and the pretraining scheme proposed by Phang et al. (2022) (‘CACLM’) in Table 4. As the pretraining scheme is primarily for improving the parameter generators, we evaluate its effect both

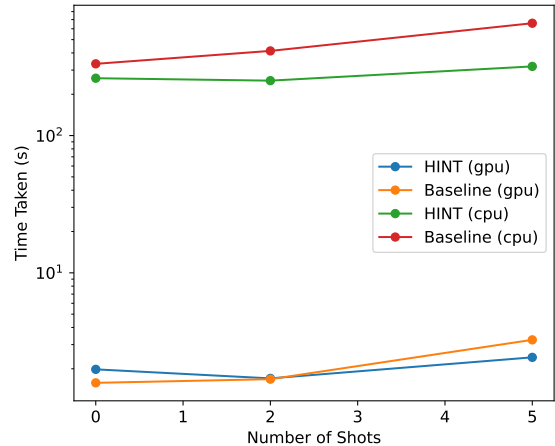


Figure 4: Time taken to process 100 examples with average SNI lengths with varying numbers of shots on CPU and GPU for 3B HINT and baseline (i.e., vanilla T5) models.

with and without using instruction fusion (‘HINT’ and ‘HINT - No Instr. Fus.’, respectively).

We find that: (a) **using pretraining gives a large boost in performance for hypernetwork-only models**, showing that pretraining is essential to good hypernetwork performance, and (b) **using our pretraining scheme works best overall**. We hypothesise this reflects the fact that our scheme is closer to the Super-Natural Instructions format than CACLM. Unlike Phang et al. (2022), we found that further pretraining did not aid performance. This is likely due to the fact that we tie the underlying model encoder and hypernetwork encoder weights together, meaning that the model weights must balance between acting as the hypernetwork and underlying model encoder.

6.2 Inference Speed

While HINT provides significant FLOPs reductions compared to baselines, these do not necessarily translate to real-world inference speedups. We examine this by measuring the average speed of HINT to process 100 samples of the same task, assuming the average input lengths given in Appendix A.1.

As seen in Figure 4, while baseline decoding remains faster for small input lengths on GPU⁷, it lags compared to HINT for longer sequences. In fact, HINT’s inference latency increases at a much slower rate compared to the baseline as the input size increases (with the number of shots), highlight-

⁷This is likely due to the small additional overhead of running the HyperEncoder, which must be run before the rest of the model.

Model	SNI RougeL
Adapters + Prefixes	32.1
Adapters ($a = 512$)	30.1
Prefixes ($l = 30$)	12.1
Prefixes ($l = 512$)	15.1
LoRA ($r = 128$)	12.1
LoRA ($r = 512$)	12.6

Table 5: SNI performance of different parameter-efficient modules in a HINT model without instruction fusion. a , l , r are the bottleneck size, number of tokens, and rank used for each experiment respectively.

Model	SNI RougeL
HINT	47.2
+ Decoder	42.6
- Instr. Fus.	32.1
- PEFT Gen.	40.9

Table 6: HINT model ablations. All models are pre-trained for 10,000 steps, except for - PEFT Gen., which contains no new parameters and requires no pretraining.

ing that HINT is especially effective in few-shot scenarios and scenarios with lengthy inputs.

6.3 Architecture Ablations

We experiment with a series of ablations to determine the best architecture for HINT, and find that:

Adapters and prefixes work best together. We consider alternatives to using adapters and prefixes together: using adapters alone, using prefixes alone, and using LoRA (Hu et al., 2022) instead of either. In order to isolate the effect of these choices, we test without using instruction fusion. We find that adapters and prefixes provide the best overall performance, with prefixes-only and LoRA-only performance substantially worse, even when increasing the number of parameters generated. This suggests that our hypernetwork approach is more adept at generating certain types of PEFT modules.

PEFT and instruction fusion are complementary. We find that using just the generated parameter-efficient modules or the encoded instruction alone (‘-Instr. Fus.’ and ‘-PEFT Gen.’ in Table 6) perform significantly worse than using both methods together, suggesting that these methods provide complementary improvements.

Cross-Attention Layer wins over Full Decoder. We compare using a full T5 decoder (with self-attention removed) as the hypernetwork weight gen-

erator as in Phang et al. (2022) with our approach, and find that our single multi-head cross-attention layer performs better at a much cheaper cost than using the full decoder (‘+ Decoder’ in Table 6).

7 Conclusion

We introduce Hypernetworks for INstruction Tuning (HINT) models and show that they consistently outperform strong full-input baselines when controlling for inference compute. This is primarily due to the fact that HINT models process their task instructions once per task, while current state-of-the-art models re-encode instructions with every task input. We show that the success of HINT models relies on a pretrained hypernetwork, which converts task instructions into parameter-efficient modules and an encoded instruction, both of which we insert into the underlying model.

Future work could investigate how HINT aids in few-shot settings, further building on HINT’s strong few-shot efficiency and taking advantage of the improved initialisation provided by hypernetworks (Phang et al., 2022). Overall, HINT models combine the benefits of parameter-efficient learning with the benefits of instruction-based learning, allowing one to easily turn pretrained language models into efficient, task-customised models.

Limitations

While promising, HINT comes with several drawbacks related to its ease of use. First, HINT takes advantage of the fact that (a) instructions are often long, and (b) often we want to perform inference over a larger (> 100) amount of examples with the same instruction. If either of these items are not true in a setup, then HINT is unlikely to provide a large benefit over simply including the instruction with the input text. This can be seen in the smaller compute savings provided by HINT for P3 in Table 2. Second, while HINT is compute-efficient at inference time, it is far more costly to train, as it effectively requires running the underlying model together with the hypernetwork for every batch. This means that while HINT may be useful for practitioners with limited compute budgets, it may be difficult to train HINT models with the same limited budget. Finally, we train and test on English data only, and do not explore the generalisation of our approach to multilingual setups. Considering the success of hypernetworks in multilingual settings (Platanios et al., 2018; Baziotis

et al., 2022; Ustun et al., 2022), we believe this is a promising direction for future research. As such, while promising, HINT is limited by certain assumptions made about the length and format of instruction-augmented data, and we hope further improvements of the method work towards loosening these assumptions.

Ethics Statement

We believe that the impact of our work is largely beneficial, examining a novel method to make instruction-based models cheaper to use. This may aid in reducing the carbon footprint of large language models running in inference (Schwartz et al., 2019) and in making these models more accessible to people with limited compute budgets. However, we also note that our approach requires unsupervised pretraining on a large corpus, making it difficult to document exactly the data it has seen during training and making it likely to reflect problematic or even dangerous biases within the corpus (Bender et al., 2021). We believe that future research could investigate reducing the need for hypernetwork pretraining and further investigate the behaviour of hypernetwork-augmented language models.

Acknowledgements

Research supported with Cloud TPUs from Google’s TPU Research Cloud (TRC). We thank AllenNLP members and Jonas Pfeiffer for encouraging and useful conversations in earlier stages of this project.

References

Stephen H. Bach, Victor Sanh, Zheng-Xin Yong, Albert Webson, Colin Raffel, Nihal V. Nayak, Abheesht Sharma, Taewoon Kim, M Saiful Bari, Thibault Fevry, Zaid Alyafeai, Manan Dey, Andrea Santilli, Zhiqing Sun, Srulik Ben-David, Canwen Xu, Gunjan Chhablani, Han Wang, Jason Alan Fries, Maged S. Al-shaibani, Shanya Sharma, Urmish Thakker, Khalid Almubarak, Xiangru Tang, Xiangru Tang, Mike Tian-Jian Jiang, and Alexander M. Rush. 2022. [Promptsources: An integrated development environment and repository for natural language prompts](#).

Christos Baziotis, Mikel Artetxe, James Cross, and Shruti Bhosale. 2022. [Multilingual machine translation with hyper-adapters](#). *EMNLP*.

Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. [On the](#)

[dangers of stochastic parrots: Can language models be too big?](#) In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT ’21, page 610–623, New York, NY, USA. Association for Computing Machinery.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.

Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2022. [Scaling instruction-finetuned language models](#).

Budhaditya Deb, Guoqing Zheng, and Ahmed Hassan Awadallah. 2022. [Boosting natural language generation from instructions with meta-learning](#).

David Ha, Andrew M. Dai, and Quoc V. Le. 2017. [Hypernetworks](#). In *International Conference on Learning Representations*.

Junxian He, Chunting Zhou, Xuezhe Ma, Taylor Berg-Kirkpatrick, and Graham Neubig. 2022a. [Towards a unified view of parameter-efficient transfer learning](#). In *International Conference on Learning Representations*.

Yun He, Steven Zheng, Yi Tay, Jai Gupta, Yu Du, Vamsi Aribandi, Zhe Zhao, Yaguang Li, Zhao Chen, Donald Metzler, Heng-Tze Cheng, and Ed H. Chi. 2022b. [HyperPrompt: Prompt-based task-conditioning of transformers](#). In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 8678–8690. PMLR.

Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. [Parameter-efficient transfer learning for NLP](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 2790–2799. PMLR.

- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [LoRA: Low-rank adaptation of large language models](#). In *International Conference on Learning Representations*.
- Hamish Ivison and Matthew E. Peters. 2022. [Hyperdecoders: Instance-specific decoders for multi-task NLP](#). *Findings of EMNLP*.
- Srinivasan Iyer, Xi Victoria Lin, Ramakanth Pasunuru, Todor Mihaylov, Daniel Simig, Ping Yu, Kurt Shuster, Tianlu Wang, Qing Liu, Punit Singh Koura, Xian Li, Brian O’Horo, Gabriel Pereyra, Jeff Wang, Christopher Dewan, Asli Celikyilmaz, Luke Zettlemoyer, and Ves Stoyanov. 2022. [Opt-impl: Scaling language model instruction meta learning through the lens of generalization](#).
- Gautier Izacard and Edouard Grave. 2021. [Leveraging passage retrieval with generative models for open domain question answering](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 874–880, Online. Association for Computational Linguistics.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. [Scaling laws for neural language models](#). *CoRR*, abs/2001.08361.
- Rabeeh Karimi Mahabadi, Sebastian Ruder, Mostafa Dehghani, and James Henderson. 2021. [Parameter-efficient multi-task fine-tuning for transformers via shared hypernetworks](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 565–576, Online. Association for Computational Linguistics.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. [The power of scale for parameter-efficient prompt tuning](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3045–3059, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Xiang Lisa Li and Percy Liang. 2021. [Prefix-tuning: Optimizing continuous prompts for generation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4582–4597, Online. Association for Computational Linguistics.
- Haokun Liu, Derek Tam, Mohammed Muqeeth, Jay Mohhta, Tenghao Huang, Mohit Bansal, and Colin Raffel. 2022. Few-shot parameter-efficient fine-tuning is better and cheaper than in-context learning. *arXiv preprint arXiv:2205.05638*.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *International Conference on Learning Representations*.
- Swaroop Mishra, Daniel Khoshabi, Chitta Baral, and Hannaneh Hajishirzi. 2022. [Cross-task generalization via natural language crowdsourcing instructions](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3470–3487, Dublin, Ireland. Association for Computational Linguistics.
- Jason Phang, Yi Mao, Pengcheng He, and Weizhu Chen. 2022. [Hypertuning: Toward adapting large language models without back-propagation](#).
- Emmanouil Antonios Platanios, Mrinmaya Sachan, Graham Neubig, and Tom Mitchell. 2018. [Contextual parameter generation for universal neural machine translation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 425–435, Brussels, Belgium. Association for Computational Linguistics.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2022. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(1).
- Adam Roberts, Hyung Won Chung, Anselm Levskaya, Gaurav Mishra, James Bradbury, Daniel Andor, Sharan Narang, Brian Lester, Colin Gaffney, Afroz Mohiuddin, Curtis Hawthorne, Aitor Lewkowycz, Alex Salcianu, Marc van Zee, Jacob Austin, Sebastian Goodman, Livio Baldini Soares, Haitang Hu, Sasha Tsvyashchenko, Aakanksha Chowdhery, Jasmijn Bastings, Jannis Bulian, Xavier Garcia, Jianmo Ni, Andrew Chen, Kathleen Kenealy, Jonathan H. Clark, Stephan Lee, Dan Garrette, James Lee-Thorp, Colin Raffel, Noam Shazeer, Marvin Ritter, Maarten Bosma, Alexandre Passos, Jeremy Maitin-Shepard, Noah Fiedel, Mark Omernick, Brennan Saeta, Ryan Sepassi, Alexander Spiridonov, Joshua Newlan, and Andrea Gesmundo. 2022. [Scaling up models and data with t5x and seqio](#).
- Victor Sanh, Albert Webson, Colin Raffel, Stephen Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Arun Raja, Manan Dey, M Saiful Bari, Canwen Xu, Urmish Thakker, Shanya Sharma Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal Nayak, Debajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang, Matteo Manica, Sheng Shen, Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang, Trishala Neeraj, Jos Rozen, Abheesht Sharma, Andrea Santilli, Thibault Fevry, Jason Alan Fries, Ryan Teehan, Teven Le Scao, Stella Biderman, Leo Gao,

- Thomas Wolf, and Alexander M Rush. 2022. [Multi-task prompted training enables zero-shot task generalization](#). In *International Conference on Learning Representations*.
- Jürgen Schmidhuber. 1992. [Learning to control fast-weight memories: An alternative to dynamic recurrent networks](#). *Neural Computation*, 4(1):131–139.
- Roy Schwartz, Jesse Dodge, Noah A. Smith, and Oren Etzioni. 2019. [Green AI](#). *CoRR*, abs/1907.10597.
- Noam Shazeer and Mitchell Stern. 2018. [Adafactor: Adaptive learning rates with sublinear memory cost](#). In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 4596–4604. PMLR.
- Yi Tay, Zhe Zhao, Dara Bahri, Donald Metzler, and Da-Cheng Juan. 2021. [Hypergrid transformers: Towards a single model for multiple tasks](#). In *International Conference on Learning Representations*.
- A. Ustun, Arianna Bisazza, Gosse Bouma, Gertjan van Noord, and Sebastian Ruder. 2022. [Hyper-x: A unified hypernetwork for multi-task multilingual transfer](#). *EMNLP*, abs/2205.12148.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Tomer Volk, Eyal Ben-David, Ohad Amosy, Gal Chechik, and Roi Reichart. 2022. [Example-based hypernetworks for out-of-distribution generalization](#). *arXiv preprint arXiv:2203.14276*.
- Thomas Wang, Adam Roberts, Daniel Hesslow, Teven Le Scao, Hyung Won Chung, Iz Beltagy, Julien Launay, and Colin Raffel. 2022a. [What language model architecture and pretraining objective works best for zero-shot generalization?](#) In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 22964–22984. PMLR.
- Yizhong Wang, Swaroop Mishra, Pegah Alipoor-molabashi, Yeganeh Kordi, Amirreza Mirzaei, Anjana Arunkumar, Arjun Ashok, Arut Selvan Dhanasekaran, Atharva Naik, David Stap, et al. 2022b. [Super-NaturalInstructions: Generalization via Declarative Instructions on 1600+ Tasks](#). In *EMNLP*.
- Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V Le. 2022. [Finetuned language models are zero-shot learners](#). In *International Conference on Learning Representations*.
- Orion Weller, Nicholas Lourie, Matt Gardner, and Matthew E. Peters. 2020. [Learning from task descriptions](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1361–1375, Online. Association for Computational Linguistics.
- Zhaofeng Wu, Robert L. Logan IV, Pete Walsh, Akshita Bhagia, Dirk Groeneveld, Sameer Singh, and Iz Beltagy. 2022. [Continued pretraining for better zero- and few-shot promptability](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Qinyuan Ye and Xiang Ren. 2021. [Learning to generate task-specific adapters from task description](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 646–653, Online. Association for Computational Linguistics.
- Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, Todor Mihaylov, Myle Ott, Sam Shleifer, Kurt Shuster, Daniel Simig, Punit Singh Koura, Anjali Sridhar, Tianlu Wang, and Luke Zettlemoyer. 2022. [Opt: Open pre-trained transformer language models](#).

A Dataset Details

A.1 Input Lengths

When calculating FLOPs estimates, we use the median sequence length of the inputs and outputs to calculate inference costs. We compute the median over the train split of Super-Natural Instructions and over 10,000 random samples from the T0 train split of P3. We calculate the medians for each format separately, rather than adding the instance and instruction-only values together (hence the mismatch in values). We provide the calculated values in Table 7. We find that P3 inputs mostly consist of the instance, with prompt templates consisting of relatively few tokens, while SNI inputs consist mostly of instructions. This explains why HINT models are much cheaper than Tk-Instruct models, but not that much cheaper than T0 models, as HINT models reduce FLOPs by avoiding reprocessing the instruction with every input.

A.2 Split Sizes

We report the sizes of splits here. For Super-Natural Instructions, we use the default setting from Wang et al. (2022b) where 100 examples are provided for each task in train and test splits. We also note that we follow the sampling procedure used by Sanh et al. (2022), where we “treat any dataset with over 500,000 examples as having 500,000 /

Text Sequence	Median # Tokens	
	SNI	P3
Instance only	44	81
Instruction only	69	24
Instruction + Instance	133	103
Instruction + 2 positives	197	-
Instruction + 2 pos. + instance	199	-
Output	1	6

Table 7: Median sequence length, given in number of T5 tokens, for Super-Natural Instructions and P3.

Dataset	Train	Test
Super-Natural Instructions	75,417	11,810
P3	90,897,454	2,940,068
P3 (adjusted for sampling)	17,277,532	2,940,068

Table 8: Number of samples in given splits for each dataset.

num templates examples” during training. Taking this sampling into account results in the much smaller dataset size seen in Table 8. We refer readers to Sanh et al. (2022) for more details on P3.

B Full P3 Results

We report the full results of models on the P3 dataset from Table 2 in Table 9.

C Model Compute Calculations

We provide a more thorough description of the compute and memory costs associated with various models we discuss here.

C.1 Compute Costs

We will let i be the sample length, t be the task instruction length, o be the output sequence length, n be the number of same-task samples we wish to process, and N be the number of parameters in the model. We will assume we are only processing examples from the same task.

Tk-Instruct As Tk-Instruct concatenates instruction and sample together as input, it uses roughly $Nn(i + t + o)$ FLOPs.

HINT The cost of the HINT model is more complicated. Let N' be the cost of the hypernetwork generator, and A be the cost of the parameter-efficient modules inserted into the underlying model. The cost of running the hypernetwork is

$t(N + N')$ (since the hypernetwork encoder is the same size as the underlying model). The cost of then running the underlying model with parameter-efficient modules is $n(N + A)(i + o)$. We sum these two terms to get the total cost of HINT: $t(N + N') + n(N + A)(i + o)$. We do not consider the additional cost of inserting the instruction in the decoder as this only affects the few (usually 1-2) output tokens and the decoder cross-attention only, and so is negligible. We can simplify the HINT compute cost down by observing that in most cases $N \gg A$ and $N \gg N'$, resulting in the cost of HINT being roughly $tN + nN(i + o)$. This simpler formulation highlights the main benefit of HINT: the instruction no longer is processed with every sample, and so compute cost is $\propto t + n$ as opposed to $\propto tn$.

C.2 Memory Costs

Here, we will let l be the number of layers, d the model hidden dimension, h the number of heads, k the size of the keys/values, and s be the length of the sequence we want to save. We ignore bias terms for simplicity.

Decoder-only Models If we want to cache the key/value pairs for a given sequence, we will store $2l h k s$ values - a key and value for every head in every layer, for each item in the sequence. We note that typically $kh = d$ in models, and so in the main text we simplify this to $2lds \propto lds$.

HINT In the default HINT setup, we save three elements: the processed instruction sequence, which contains ds values (one vector per token); the adapter weights, $2 * 512ld$ values (one adapter comprising of two weight matrices per layer, where each weight matrix has size $512 \times d$); the prefix values, $2 * 30lhk$ values (a 30-length prefix and key per layer per head). This sums to give a total memory cost of $ds + 1024ld + 60lhk$. Note that in the default HINT settings, we use prefixes of length 30 and adapters with bottleneck size 512, but these settings could be adjusted to reduce memory costs. Applying the simplification $kh = d$, we get that the HINT memory cost is $\propto ds + ld$.

D GPT-2 Instruction Finetuning

When finetuning GPT-2 for Table 1, we trained for [3, 5, 10] epochs with a batch size of 32. We use AdamW (Loshchilov and Hutter, 2019) and swept learning rates of $[1 \times 10^{-5}, 2 \times 10^{-5}, 5 \times 10^{-5}]$,

Model	Avg	Rel. FLOPs	ANLI	HellaSwag	StoryCloze	CB	COPA	RTE	WiC	WSC	WinoGrande
T0-3B	54.9	×1.0	33.4	27.3	84.0	45.4	72.8	64.6	50.6	64.9	50.9
T0-3B (our replication)	64.4	×1.0	41.7	30.1	96.9	72.7	89.1	81.2	51.7	57.2	59.2
T0-3B + PEFT	65.5	×1.0	41.5	30.1	96.6	76.9	92.2	82.1	54.2	56.6	59.2
No Prompt	57.5	× 0.8	34.4	27.7	88.8	69.4	66.3	56.5	52.5	61.3	60.6
Hypter (Joint)	64.6	×1.0	41.1	29.4	96.7	76.3	87.4	79.6	52.1	58.3	60.9
HINT (Joint)	65.4	×1.1	41.6	30.3	96.6	76.0	88.8	84.2	51.4	59.5	60.1
Hypter (Split)	56.2	× 0.8	34.2	28.1	86.8	58.0	67.3	65.0	50.5	60.0	55.7
HINT (Split)	60.3	× 0.8	37.0	29.1	85.6	67.6	71.0	77.2	51.0	64.2	60.0

Table 9: Results on T0 evaluation tasks after training on the T0 P3 train set. We report averaged accuracy across prompts for each task, and the overall average performance (‘avg’). FLOPs are calculated assuming we are processing 100 examples of a single task. The ‘Joint’ and ‘Split’ HINT variants refer to the two input formats for P3 described in Section 4.

using a linear warmup and decay schedule with 1,000 steps of warmup. We report the highest overall results in Table 1. Following Iyer et al. (2022), we minimise the loss only over the target tokens (with EOS token added after the target answer), not the inputs, since these are always provided during test time. Note that we calculate the FLOPs used by GPT-2 during inference based on the estimates provided by Kaplan et al. (2020) that GPT-series models use $2N$ FLOPs per token, where N is the number of parameters in the model.

E T0 Replication

During initial experiments, we replicated the T0 training in the t5x framework, using the same training set and mixing proportions as Sanh et al. (2022). We found that our replications performed significantly better than the reported T0 performance when trained for longer. We train 3B and 11B size models on the T0 training mixture for 20,000 steps using a batch size of 2048, a maximum input sequence length of 1024, a maximum output sequence length of 256, and the Adafactor optimizer with a constant learning rate of 0.001. We start from the T5 v1.1 + LM adaptation checkpoints and fully finetune the model. As seen in Table 10, our replications significantly outperform both T0 models, suggesting that T0 was undertrained. We also compare the variances in prompt performance in Figure 5.

Model	ANLI	HellaSwag	StoryCloze	CB	COPA	RTE	WiC	WSC	WinoGrande	AVG
T0-3B	33.4	27.2	84.0	45.4	75.9	64.6	50.7	65.1	51.0	55.2
T0-3B (ours)	41.7	30.1	96.9	72.7	89.1	81.2	51.7	57.2	59.2	64.4
T0-11B	41.0	33.6	92.4	70.1	91.5	81.0	56.1	61.1	59.9	65.2
T0-11B (ours)	46.8	34.1	98.2	81.2	96.6	84.0	52.1	62.6	64.8	68.9

Table 10: T0 evaluation task accuracy, comparing results using models from Sanh et al. (2022) and our own replications. We report accuracy averaged across prompts for the same dataset. For ANLI, we first average performance across prompts for ANLI R1/2/3 separately, and then average the three results.

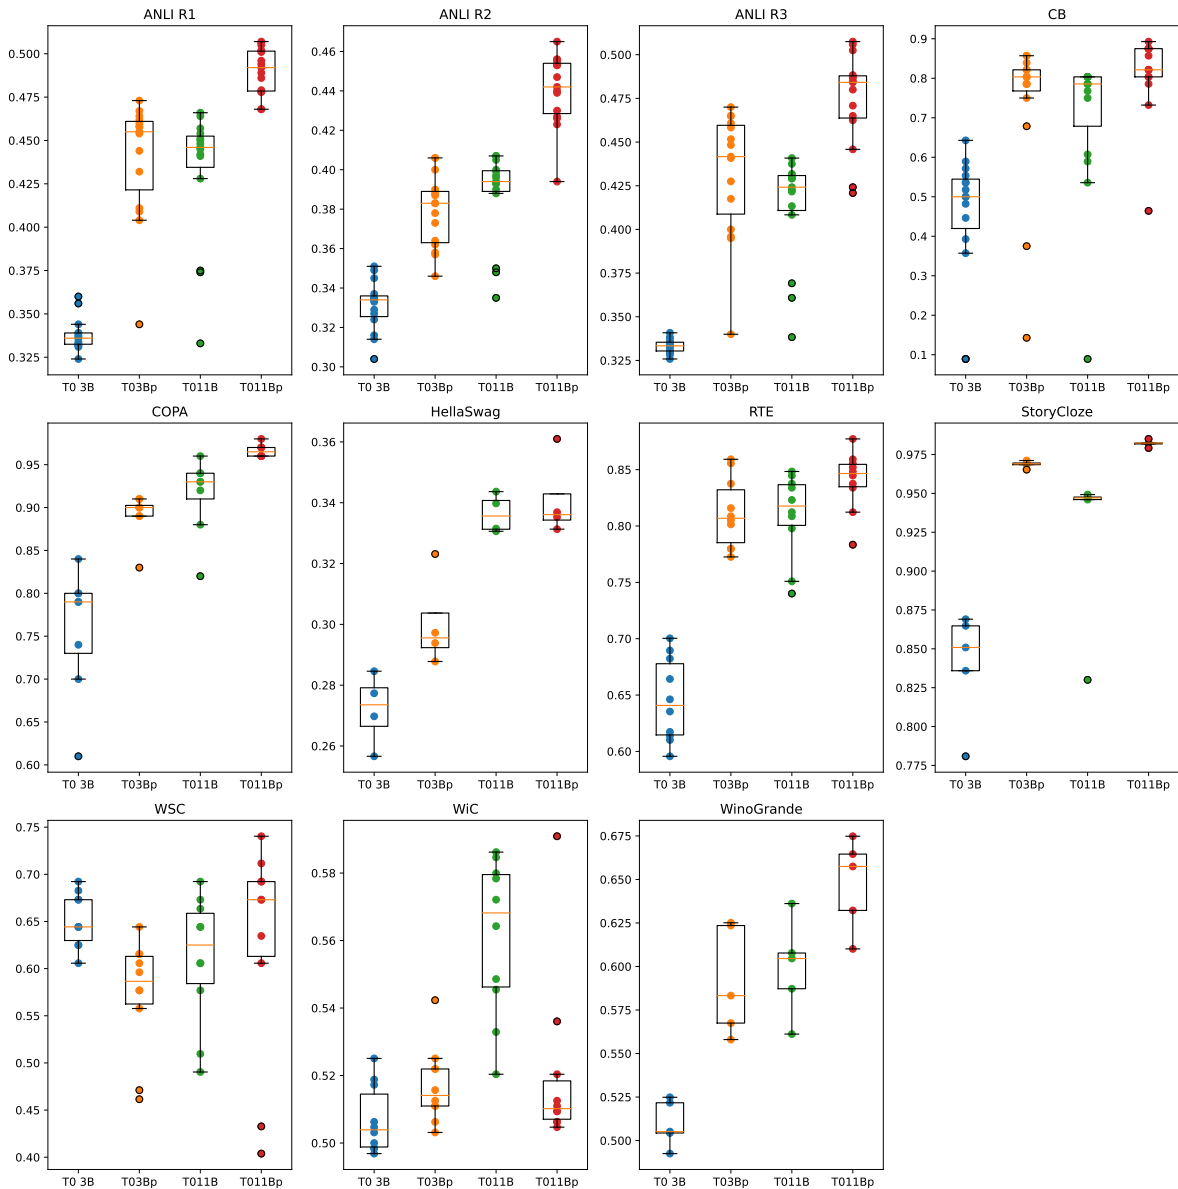


Figure 5: T0 evaluation set accuracy across tasks. Each dot represents the performance from a different prompt. ‘T03Bp’ and ‘T011Bp’ refer to our T0 replications, while ‘T03B’ and ‘T011B’ are the models released by Sanh et al. (2022).

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Limitations section after conclusion (non-numbered).
- A2. Did you discuss any potential risks of your work?
Ethics statement section after conclusion (non-numbered).
- A3. Do the abstract and introduction summarize the paper's main claims?
Abstract and section 1 (introduction)
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Primarily section 4 and 5, where we discuss models trained and datasets used.

- B1. Did you cite the creators of artifacts you used?
Section 4
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
All artefacts used were created and shared for research purposes, which we use them for.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
We use these artefacts only for the research in this work, for which they were intended.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
We do not create any new datasets or significantly alter the datasets we use, and make use only of extremely popular existing research datasets. We refer to the papers introducing the data we use for these details.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
Our work is focused on modelling, rather than data, and so we point readers to the papers proposing the datasets we use for this information.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
Section 4 and Appendix A (in particular, A.2).

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a [question on AI writing assistance](#).

C Did you run computational experiments?

Section 4 explains the experimental setup and Section 5 reports the main experiments. Section 6 contains ablation experiments.

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Section 4 and Appendix C provide details.
- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?
Section 4 and Appendix D (for GPT-2 experiment).
- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?
Section 4. Appendices B and E provides additional information on P3, and Appendix D on GPT-2 experiments.
- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?
Section 4.

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?
No response.
- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?
No response.
- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?
No response.
- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?
No response.
- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?
No response.