

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 TYPE-COMPLIANT ADAPTATION CASCADES: ADAPTING PRO- GRAMMATIC LM WORKFLOWS TO DATA

Anonymous authors

Paper under double-blind review

ABSTRACT

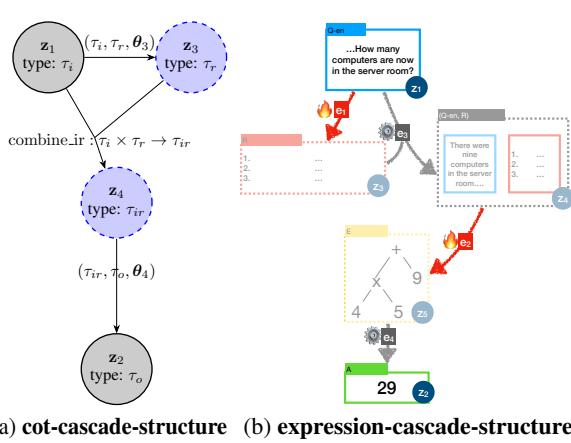
Reliably composing Large Language Models (LLMs) for complex, multi-step workflows remains a significant challenge. The dominant paradigm — optimizing discrete prompts in a pipeline — is notoriously brittle and struggles to enforce the formal compliance required for structured tasks. We introduce Type-Compliant Adaptation Cascades (TACs), a framework that recasts workflow adaptation as learning typed probabilistic programs. TACs treat the entire workflow, which is composed of parameter-efficiently adapted LLMs and deterministic logic, as an unnormalized joint distribution. This enables principled, gradient-based training even with latent intermediate structures. We provide theoretical justification for our tractable optimization objective, proving that the optimization bias vanishes as the model learns type compliance. Empirically, TACs significantly outperform state-of-the-art prompt-optimization baselines. Gains are particularly pronounced on structured tasks, improving FinQA from 12.0% to 24.7% for a Qwen 3.8B model, MGSM-SymPy from 57.1% to 75.9% for a Gemma 2.27B model, MGSM from 1.6% to 27.3%, and MuSR from 36.5% to 62.6% for a Gemma 7B model. TACs offer a robust and theoretically grounded paradigm for developing reliable, task-compliant LLM systems.

1 INTRODUCTION

The expressive power of Large Language Models (LLMs) has catalyzed the rapid development of programmatically composed workflows and agentic systems (Khattab et al., 2022; Chase, 2022; Yao et al., 2023; Wu et al., 2024). By chaining model calls and integrating deterministic logic, practitioners construct complex systems capable of multi-step reasoning and interaction. However, the dominant paradigm for adapting these systems — optimizing discrete prompts within the pipeline — is notoriously brittle (Cao et al., 2024) and struggles to enforce the formal compliance required for structured tasks. Optimization often devolves into a difficult discrete search problem (Pryzant et al., 2023; Yuksekgonul et al., 2025), relying on heuristics that are computationally expensive and difficult to scale.

In this paper, we propose a fundamental shift in perspective: we recast the entire LLM workflow as a **typed probabilistic program**. Instead of optimizing the *inputs* (prompts) to a fixed system, we optimize the *program parameters*. We treat the workflow as a parametric latent variable model where each step is a probabilistic transformation backed by a parameter-efficient fine-tuning (PEFT) adaptor. This transforms workflow adaptation from an ad-hoc, discrete search problem into a principled, gradient-based optimization task focused on maximizing data likelihood.

While probabilistic programming languages (PPLs) (Bingham et al., 2019; Tran et al., 2017) offer powerful tools for modeling complex distributions, they are generally designed to capture and conditionalize *normalized* models. LLM workflows present a unique challenge. Enforcing type constraints means restricting the support of an LLM — which naturally generates arbitrary strings — to only those strings representing valid typed objects. This restriction



(a) cot-cascade-structure (b) expression-cascade-structure

Figure 1: Two TAC workflow patterns experimented in this paper. We illustrate the more complicated Fig. 1b with example node values (we also explore additional patterns in §B). Dashed-boundary nodes indicate variables whose values are not available in annotated data, and solid-boundary nodes indicate nodes with training time observable values. A main message of this work is that **we can treat an entire typed workflow as a single probabilistic program, whose parameters are lightweight PEFT modules, allowing end-to-end training with latent variables**, instead of defining workflows imperatively as fixed-parameter systems.

renders the resulting (unconditional) distribution inherently *unnormalized* ($\mathcal{Z}_\theta \neq 1$), making the partition function required for standard maximum likelihood estimation intractable.

We introduce Type-Compliant Adaptation Cascades (TACs), a framework designed specifically for learning these typed, unnormalized probabilistic programs. TACs treat the entire workflow, composed of adapted LLMs and deterministic logic, as a joint unnormalized distribution. To enable tractable optimization, we propose the TACStaR algorithm, a generalization of the Self-Taught Reasoner (STaR) (Zelikman et al., 2022) formalized within an MC-EM framework.

Crucially, we provide theoretical justification for optimizing the tractable unnormalized likelihood directly. We prove that the bias introduced by ignoring the partition function gradient is bounded by the degree of type violation. As the model learns to comply with the workflow’s type constraints during training, the bias vanishes, and the optimization converges to the true maximum likelihood solution (Theorems 1 and 2). Furthermore, this probabilistic framing allows us to decouple inference from training, enabling advanced techniques such as amortized inference to improve the E-step during optimization.

Our primary contributions are:

- **Framework.** We formalize typed LM workflows as *unnormalized probabilistic programs*, where type contracts restrict the support of learned transformations.
- **Theory.** We propose TACStaR, a tractable optimization algorithm, and prove that its optimization bias vanishes as the model learns type compliance during training.
- **Practice.** Across reasoning-heavy tasks (MGSM, MGSM-SymPy, FinQA, MuSR) and model families (Gemma, Qwen), TACs consistently outperform strong DSPy prompt-optimization baselines. Gains are largest when (1) base models are smaller and (2) tasks require strict structure. For example, on MGSM-SymPy with a Gemma 27B model, TACs achieve **75.9** vs. **57.1**; on FinQA, **34.0** vs. **12.7** (Gemma 27B) and **24.7** vs. **12.0** (Qwen 3 8B). With a Gemma 7B model, MGSM improves from **1.6** to **27.3**, and MuSR from **36.5** to **62.6**.

094 **Summary of results.** (1) Gradient-based adaptation of typed probabilistic programs is markedly more effective
 095 and compute-efficient than discrete prompt search for structured tasks. (2) Flexible posterior inference, such as
 096 amortization, improves training stability and performance. (3) Empirically, the estimated type compliance mass
 097 \mathcal{Z}_θ rises rapidly during training, supporting our theoretical justification for the unnormalized objective.
 098

100 2 TYPE-COMPLIANT ADAPTOR CASCADES

102 The core idea of TACs is to decompose a task into a hypergraph of interconnected transformations. Formally, a TAC
 103 is represented as a directed acyclic hypergraph (DAH) $C = (\mathbf{Z}, \mathbf{E})$.¹ The acyclic constraint ensures that the
 104 workflow has a well-defined topological order for execution and guarantees termination of the generative process.
 105

106 **Nodes.** The nodes $\mathbf{Z} = \{\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_M\}$ in a TAC act as containers for typed data. Each node \mathbf{z}_m is associated
 107 with a specific data type $\tau \in \mathcal{T}$, and holds string representations $\in \Sigma^*$ for τ -typed objects. Special nodes are
 108 designated as the **input node** \mathbf{z}_1 and the **output node** \mathbf{z}_2 (e.g., holding the initial question of type Q_en and the
 109 final answer of type A in Fig. 1b, respectively).

110 **Hyperedges.** Hyperedges $\mathbf{E} = \{e_1, e_2, \dots, e_K\}$ define the transformations between nodes. A hyperedge e_k
 111 connects a set of source nodes $S_k \subseteq \mathbf{Z}$ (its inputs) to a set of target nodes $T_k \subseteq \mathbf{Z}$ (its outputs). Transformations
 112 in TACs can be either learnable (LM adaptors) or fixed (deterministic algorithms):
 113

- 114 • **LM adaptor hyperedges.** These are stochastic transformations implemented by PEFT-adapted LMs. An adaptor
 115 (τ_i, τ_o, θ) defines an unnormalized distribution over $\mathbf{y} \in \Sigma^*$ given input string \mathbf{x} :²

$$116 \quad \tilde{p}(\mathbf{y} \mid \mathbf{x}; \theta) = p_{LM}(\mathbf{y} \mid \mathbf{x}; \theta) \mathbb{I}(\mathbf{z}_t \in \text{valid}(\tau_o)), \quad (1)$$

118 where $p_{LM}(\cdot \mid \mathbf{x}; \theta)$ is a normalized distribution over strings, conditioned on τ_i -typed string representation \mathbf{x} ,
 119 and parametrized by adaptor parameters θ , and $\text{valid}(\tau_o) \subseteq \Sigma^*$ is the set of strings that represent valid τ_o -typed
 120 objects (we will further discuss them in §2.1).

- 121 • **Deterministic algorithm hyperedges.** These are fixed, non-learnable transformations, such as a self-contained
 122 Python function. A deterministic algorithm f maps an input object of type τ_i to an output object of type τ_o . Under
 123 the probabilistic view, we represent them as δ distributions:

$$124 \quad \tilde{p}(\mathbf{y} \mid \mathbf{x}; f) = \delta_{\text{canon}(f(\text{parse}(\mathbf{x}, \tau_i)))}(\mathbf{y}) \quad (2)$$

126 where `canon` (see §2.1) produces a canonicalized string for an object, and `parse` converts strings back to typed
 127 objects.
 128

129 2.1 INTERFACING LLMs WITH TYPED DATA: PARSING AND CANONICALIZATION

131 To integrate LLMs, which operate on strings (Σ^*), into a typed workflow, TACs require mechanisms to bridge the
 132 gap between strings and typed objects (\mathcal{O}). This bridge is typically handled by data validation libraries such as
 133

134 ¹We use a reasoning workflow that generates domain-specific code, illustrated in Fig. 1b, as a running example. The task is
 135 to take a math question in English (input type Q_en), generate a step-by-step rationale (intermediate type R), convert the
 136 rationale into a formal arithmetic expression (intermediate type E), and finally, have a deterministic function evaluate this
 137 expression to produce the answer (output type A). This section formalizes how such an intuitive sketch is realized within the
 138 TAC framework.

138 ²This distribution may be unnormalized because while p_{LM} is a distribution over all strings, Eq. (1) restricts the support to
 139 only strings that are valid instances of τ_o . Thus, the total probability mass may sum to less than 1 if the LM assigns probability
 140 to invalid strings.

141 Pydantic³ and PyGlove (Peng et al., 2020).⁴ We formalize this conversion using two essential operations: `parse`
 142 and `canon`.

144
 145 **Parsing (parse).** When an LM adaptor produces an output string y intended to represent an object of type τ_o ,
 146 this string is validated and converted into a usable typed object by the algorithm $\text{parse} : \Sigma^* \times \mathcal{T} \rightarrow \mathcal{O} \cup \{\text{error}\}$.⁵
 147 The expressivity of the TAC type system is determined by the implementation of `parse` and `canon`. In this
 148 work, we leverage LangFun, which supports primitive types, compound types (e.g., Python classes), and recursive
 149 data types (an example is listed in Listing 7).

150 For example, in Fig. 1b, \mathbf{z}_5 has the deterministic function e_4 as an outgoing edge. During execution of the
 151 probabilistic program, $\text{parse}(\mathbf{z}_5, \mathbb{E})$ attempts to convert \mathbf{z}_5 into a SymPy expression object (typed \mathbb{E}). If the
 152 conversion fails, an error is signaled. For convenience, we use $\text{valid}(\tau) = \{\text{parse}(y, \tau) \neq \text{error} \mid y \in \Sigma^*\}$ to
 153 denote valid string representations of τ .

154
 155 **Canonicalization (canon).** Conversely, inputs of LM adaptor hyperedges must be converted into a consistent
 156 string format that the adaptor expects. The $\text{canon} : \mathcal{O} \rightarrow \Sigma^*$ operation maps a typed object to a unique string
 157 representation — we call such strings *canonicalized*. The invertibility of `canon` (i.e., $\text{parse}(\text{canon}(o), \tau_o) = o$)
 158 in turn ensures that deterministic hyperedges have support over only one string given a valid input, eliminating
 159 spurious ambiguity (Cohen et al., 2012).

160 2.2 TACS AS PROGRAMS AND DISTRIBUTIONS

161 TACs admit both a program view, and also a probabilistic view⁶:

- 162 • **TACs are probabilistic programs.** From an operational perspective, executing a TAC in the forward direction
 163 involves processing data through the hypergraph, respecting the topological order of nodes and hyperedges. Using
 164 our running example from Fig. 1b: the process traverses the hypergraph, starting at the input variable \mathbf{z}_1 (typed
 165 $Q\text{-en}$), and ending at the output variable \mathbf{z}_2 (typed \mathbb{A}). A general process is described in Algorithm 1.
- 166 • **TACs are also probability distributions.** From a statistical perspective, TACs define unnormalized joint probability
 167 distributions over all node assignments $\mathbf{Z}^* = (\mathbf{z}_1^*, \mathbf{z}_2^*, \dots, \mathbf{z}_M^*)$. This score reflects the plausibility of a complete
 168 execution trace according to the model’s components:

$$172 \log \tilde{p}_{\theta}(\mathbf{Z}^*) = \sum_k \log \tilde{p}_{\theta}(\{\mathbf{z}_t^*\}_{t \in T_k} \mid \{\mathbf{z}_s^*\}_{s \in S_k}; e_k), \quad (3)$$

173 where θ represent all adaptor parameters used in the TAC, and $\tilde{p}_{\theta}(\cdot \mid \cdot; e_k)$ is the conditional probability defined by
 174 the LM adaptor (Eq. (1)) or deterministic algorithm (Eq. (2)) associated with e_k . The unnormalized distribution
 175 view connects TACs to the broader family of language model cascades (Dohan et al., 2022), but with the key
 176 distinction that TACs are designed for end-to-end adaptation.

177
 178 ³<https://github.com/pydantic/pydantic>

179 ⁴Examples of type-specifying prompts generated by LangFun (which leverages PyGlove) are listed in §N. LangFun
 180 supports primitive types, compound types (e.g., Pyglove/Pydantic objects), and recursive types (e.g., Expressions in Listing 7).

181 ⁵We note that while primitive data types (e.g., Python types `str` and `list`) appear in common workflows, `parse` can
 182 be any computable function, and can be leveraged by a practitioner to implement complex business logic. For example,
 183 one can define a Python custom type `CoherentDialog` where valid objects are strings deemed coherent by an external
 184 LLM-backed classifier, and adapt LM adaptors in a TAC to generate and work with such objects. Implementation details are
 185 further discussed in §E.

186 ⁶These two views are also summarized in Table 1.

188

3 ADAPTING TACS

189
 190 The goal of adapting a TAC is to maximize the marginalized likelihood of the training data. Since TACs generally
 191 define distributions over unobserved (latent) intermediate variables, Monte Carlo Expectation-Maximization
 192 (MC-EM) algorithms (Wei & Tanner, 1990) provide a suitable training paradigm.⁷

193 However, adapting TACs presents a challenge. As TACs are generally unnormalized models, proper Maximum
 194 Likelihood Estimation (MLE) updates in the M-step require computing partition function gradients. Denoting the
 195 partition function summing all possible assignments as $\mathcal{Z}_\theta = \sum_{\mathbf{Z}'} \tilde{p}_\theta(\mathbf{Z}')$, the gradient of the log-likelihood
 196 $\mathcal{L}(\theta) = \log p(\mathbf{Z}^*)$ is:
 197

$$198 \quad \nabla_\theta \mathcal{L} = \nabla_\theta \log \tilde{p}_\theta(\mathbf{Z}^*) - \nabla_\theta \log \mathcal{Z}_\theta. \quad (4)$$

199 Estimation of the log partition function's gradients $\nabla_\theta \log \mathcal{Z}_\theta$ is typically intractable, expensive, and can have
 200 high variance (Goodfellow et al., 2016).

202

3.1 TRACTABLE OPTIMIZATION VIA COMPLIANCE

203 To overcome the intractable partition function gradient, we propose optimizing for the unnormalized log-likelihood
 204 $\mathcal{L}'(\theta) = \log \tilde{p}_\theta(\mathbf{Z}^*)$ instead, effectively dropping the $\nabla_\theta \log \mathcal{Z}_\theta$ term from Eq. (4).

205 While ignoring the partition function gradient generally leads to biased estimation, the TAC formalism ensures this
 206 strategy is both tractable and robust. This becomes evident as we rewrite $\mathcal{L}'(\theta) = \mathcal{L}(\theta) + \log \mathcal{Z}_\theta$: optimizing
 207 the unnormalized likelihood $\mathcal{L}'(\theta)$ is equivalent to jointly maximizing the normalized likelihood $\mathcal{L}(\theta)$ and the
 208 model's type compliance (the partition function $\log \mathcal{Z}_\theta$ is maximized at $\log \mathcal{Z}_\theta = 0$ when θ is well-specified).
 209 We now provide theoretical justification for this approach, under the assumption that the adapted models can
 210 perfectly model type-valid outputs (*i.e.*, the model family is well-specified).⁸

211 **Theorem 1.** *Let Θ be the entire parameter space and let $\Theta' \subseteq \Theta$ be the subset of well-specified parameters. Assume θ^* uniquely maximizes the normalized likelihood $p_\theta(\mathbf{z}_{2..M} | \mathbf{z}_1)$ and resides in Θ' . Then, $\hat{\theta} = \arg \max_{\theta \in \Theta} \tilde{p}_\theta(\mathbf{z}_{2..M} | \mathbf{z}_1) \implies \hat{\theta} = \theta^*$.*

212 Moreover, while optimizing $\mathcal{L}'(\theta)$ introduces a bias by ignoring the gradient term $\nabla_\theta \log \mathcal{Z}_\theta$, this bias is bounded
 213 below a constant multiplicative factor of $(1 - \mathcal{Z}_\theta)$ under the common assumption that $\|\nabla_\theta p_{LM}(\cdot | \mathbf{x}; \theta)\|$ is
 214 uniformly bounded:

215 **Theorem 2.** *Let $\theta = \{\theta_1 \dots \theta_K\}$ be the union of a K -adaptor TAC's LM adaptor parameters. If $\forall \mathbf{z}_{k,1} \in \Sigma^*, \mathbf{z}_{k,2} \in \Sigma^*, \|\nabla_\theta (\sum \log p_{LM}(\mathbf{z}_{k,2} | \mathbf{z}_{k,1}; \theta))\|_\infty \leq G$, then $\nabla_\theta \log \mathcal{Z}_\theta \leq 2G(1 - \mathcal{Z}_\theta)$.*

216 Theorems 1 and 2 provide theoretical assurance that if the model achieves high type compliance as we optimize for
 217 $\mathcal{L}'(\theta)$, the optimization bias vanishes, and the update approaches the true MLE update. Empirically, we observe
 218 that training rapidly drives \mathcal{Z}_θ towards 1 (§4.4).

219

3.2 TACSTAR

220 We introduce the TACSTAR algorithm (Algorithm 3), which generalizes the Self-Taught Reasoner (STaR) algorithm
 221 (Zelikman et al., 2022) to the TAC framework. TACSTAR employs an iterative MC-EM approach to optimize the
 222 tractable objective $\mathcal{L}'(\theta)$ (§3.1). It alternates between E- and M-steps:

223 ⁷We acknowledge that another reasonable approach for training TACs is reinforcement learning, and note the connection
 224 between TACSTAR and RL in §A.

225 ⁸We refer the reader to §D for proofs of formal statements in this section.

- **E-step: Sampling Latent Variables.** The E-step aims to sample complete, valid execution traces \mathbf{Z}^* consistent with the training data. We first try to execute the TAC C as a probabilistic program under the `forward` algorithm (Algorithm 1). If `forward` succeeds, we have a complete assignment of values $\mathbf{Z}^* = (\mathbf{z}_1^*, \mathbf{z}_2^*, \dots, \mathbf{z}_M^*)$ and can proceed to M-step. Otherwise, we attempt a **rationalization heuristic** step. Inspired by the original STaR algorithm which conditions on the correct answer in the second attempt, we construct a ‘fallback’ TAC, whose input node takes (x^*, y^*) as input, with the rest of the workflow unchanged. This essentially asks ‘*what intermediate steps would lead from x^* to y^* ?*’, analogous to the inverse rendering problem (Ritchie et al., 2023). A forward pass is then executed on this new TAC to sample $(\mathbf{z}_2, \dots, \mathbf{z}_M)$, now conditioned on both the original input x^* and the desired output y^* . This encourages the generation of latent intermediate steps that are consistent with the correct final answer.
- **M-step: Parameter Optimization.** In the M-step, we update the adaptor parameters θ by maximizing the unnormalized likelihood $\mathcal{L}'(\theta)$ of the samples collected in the E-step.⁹

247 3.3 AMORTIZED TACSTAR

249 The basic TACStaR algorithm relies on a fixed ‘fallback’ heuristic during the E-step, which may be inefficient.
250 Amortized TACStaR (Algorithm 4) addresses this by generalizing the heuristic using parametric inference
251 networks (Kingma & Welling, 2014; Mnih & Gregor, 2014), jointly trained to approximate the true posterior given
252 observed input and outputs. By learning to propose better, task-adapted latent variable configurations, Amortized
253 TACStaR can lead to more efficient training and potentially better performance.

254 For model TAC C with nodes $\mathbf{z}_1 \dots \mathbf{z}_M$, we construct an inference network TAC C' with nodes $\mathbf{z}'_1 \dots \mathbf{z}'_M$, which
255 is trained alongside with C . In this work, we construct $\mathbf{z}'_2 \dots \mathbf{z}'_M$ to have the same types as $\mathbf{z}_2 \dots \mathbf{z}_M$, except for
256 its input node \mathbf{z}'_1 , which has a type to represent the input-output pair (x^*, y^*) . Moreover, we construct C' so that
257 every adaptor hyperedge e_k in C has a counterpart e'_k in C' that is additionally conditioned on \mathbf{z}'_1 . We train C'
258 alternately with C , with the goal of making the unnormalized distribution of C' approximate the posterior over
259 C ’s intermediate nodes, conditioning on (x^*, y^*) observations. Denoting the unnormalized distribution of C' as
260 \tilde{q}_ϕ parametrized by adaptors’ parameters ϕ , we hope to learn ϕ such that $\tilde{q}_\phi(\mathbf{z}'_m \mid \mathbf{z}'_1 = \text{canon}((x^*, y^*))) \approx$
261 $p_\theta(\mathbf{z}_m \mid \mathbf{z}_1 = x^*, \mathbf{z}_2 = y^*)$, where $x_c^* = \text{canon}(x^*)$, $y_c^* = \text{canon}(y^*)$, $\forall m \in [2..M]$. Approximating
262 the posterior $p_\theta(\mathbf{z}_m \mid \mathbf{z}_1 = \text{canon}(x^*), \mathbf{z}_2 = \text{canon}(y^*))$ as \hat{p} using self-normalized multiple importance
263 sampling (Veach & Guibas, 1995), we optimize ϕ to minimize $\text{KL}[\hat{p} \parallel q_\phi]$ following Bornschein & Bengio (2014);
264 Lin & Eisner (2018). Empirically, we verify that the learned inference network C' significantly reduces the KL
265 divergence to the true posterior compared to the fixed fallback, confirming it provides a better approximation for
266 training (§P).

267 4 EXPERIMENTS

270 To empirically validate TAC models, we conduct QA, code-like structured generation, and classification experiments
271 on subsets of MGSM (Shi et al., 2023), FinQA (Chen et al., 2021), and MuSR (Sprague et al., 2024b) datasets,¹⁰
272 adapting both instruction-tuned Gemma 7B and Gemma 2 27B (referred to as `gemma-1.1-7b-it` and
273 `gemma-2-27b-it`) (Team et al., 2024), and Qwen 3 8B models (Qwen3-8B) (Yang et al., 2025). We aim to
274 answer the following research questions:

275
276 ⁹**Remark on efficiency.** Since gradients of the log unnormalized probability decompose linearly as $\nabla_\theta (\log \tilde{p}_\theta(\mathbf{Z}^*)) =$
277 $\sum_k \nabla_\theta \log \tilde{p}_\theta(\{\mathbf{z}_t^*\}_{t \in T_k} \mid \{\mathbf{z}_s^*\}_{s \in S_k}; e_k)$, computation of adaptors’ gradients can be parallelized easily. This embarrassingly
278 parallel structure ensures computational scalability, allowing the M-step to be efficiently distributed across available compute
279 resources. Algorithm 2 computes $\log \tilde{p}_\theta(\mathbf{Z}^*)$ and its gradients $\nabla_\theta \log \tilde{p}_\theta(\mathbf{Z}^*)$. These gradients are then used in a standard
280 gradient-based optimization algorithm to update θ .

281 ¹⁰We defer the study of how different TAC patterns affect performance to §B, where we expand our experiments to include
282 HotPotQA tasks (Yang et al., 2018).

- **(§4.2) Are TACs competitive against existing approaches?** TACs differ from existing LM adaptation approaches in two major ways: 1) TACs support gradient-based learning in a unified probabilistic programming framework (when compared against prior prompt optimization-focused LM programming frameworks such as DSPy); and 2) TACs support structured workflows by design (when compared to the original STaR algorithm). We hypothesize that such difference translates into meaningful performance improvements.
- **(§4.3) Is exploiting TACs' probabilistic flexibility effective?** Probability models (such as TACs) benefit from the decoupling of probabilistic modeling and inference procedures, allowing conditioning on additional observations *a posteriori*. We evaluate whether exploiting this flexibility is effective in two scenarios: 1) We compare Amortized TACSTaR (§3.3), which conditions on the output variable to learn a better proposal distribution for training, against the standard (unconditioned) TACSTaR; and 2) We evaluate TACs on a classification task, comparing the performance of unconstrained generation against a renormalized classifier that evaluates and normalizes the conditional probability of each possible output.
- **(§4.4) Does the model achieve high type compliance?** A key theoretical result (§3.2) is that the soundness and near-optimality of the TACSTaR optimization strategy rely on the model learning to comply with the workflow's type constraints (*i.e.*, driving the partition function $\mathcal{Z}_\theta \rightarrow 1$). As type compliance increases, the gap between the tractable unnormalized likelihood and the true normalized likelihood ($\log \mathcal{Z}_\theta$) closes. We estimate how \mathcal{Z}_θ over TACSTaR epochs to verify that this gap is negligible after training.

300 4.1 EXPERIMENT SETUP

302 We provide an overview of our TAC and baseline DSPy setups below:

- **TACs.** We parametrize TAC adaptors to take the form of rank-1 LoRA models (Hu et al., 2022) on the attention weights, with 573, 440; 1, 413, 120; and 958, 464 parameters per adaptor for `gemma-1.1-7b-it`, `gemma-2-27b-it` and `Qwen3-8B` respectively. For `parse` and `canon` implementations (§2.1), we leverage the LangFun library, which prompts LLMs to generate Python classes and objects, and parses their responses. LoRA weights are initialized ('zero-init') following Hu et al. (2022).
- **DSPy.** We conduct prompt-optimizing baseline experiments under DSPy, with base models served on vLLM. We subclass `dspy.Signature` to represent training examples, with property names and types identical to their TAC counterparts (some examples are listed in §G.2). We employ XGrammar (Dong et al., 2024) for schema-based constrained decoding for all experiments. We implement two types of reasoning workflows for all tasks: 1) the native `dspy.ChainOfThought` module, and 2) an explicitly two-step composite module that resembles **cot-cascade-structure** patterns under TACs. We experiment with various prompt optimization configurations under `dspy.MIPROv2` (Opsahl-Ong et al., 2024) and `dspy.BootstrapFewShotWithRandomSearch` (Khattab et al., 2024).

316 We conduct experiments of 5 reasoning-heavy tasks, on subsets from datasets MGSM¹¹ (Shi et al., 2023), FinQA
 317 (Yang et al., 2018), HotPotQA (Yang et al., 2018) and MuSR (Sprague et al., 2024b) respectively. Details of
 318 experiment setup are described in §G.

320 4.2 COMPARISON AGAINST PROMPT-OPTIMIZING AND UNTYPED STAR BASELINES.

323 Figure 2 lists MGSM, MGSM-SymPy, FinQA, and MuSR results from best-performing TACs and DSPy models.
 324 In addition, we compare the untyped (original) STaR against typed TAC results on MGSM on Gemma models.

326 ¹¹The MGSM-SymPy task uses the same problems of MGSM, but additionally restrict the outputs to be rational expressions
 327 under SymPy. This variant was specifically included to test the framework's ability to generate and comply with highly
 328 structured, code-like output.

TACs are competitive against prompt-optimizing baseline methods. We observe that TACs consistently and significantly outperform DSPy baselines in every setting. The performance gap is especially wide when 1) the base model is smaller, and 2) the task involves structured inputs (FinQA) or structured outputs (MGSM-SymPy).¹²

Base Model	DSPy	TAC
gemma-1.1-7b-it	0.7%	9.7%
gemma-2-27b-it	12.7%	34.0%
Qwen3-8B	12.0%	24.7%

Base Model	DSPy	TAC
gemma-1.1-7b-it	36.5%	62.6%
gemma-2-27b-it	51.5%	65.0%
Qwen3-8B	61.5%	63.7%

Base Model	DSPy	TAC
gemma-2-27b-it	57.1%	75.9%

Base Model	DSPy	TAC
gemma-1.1-7b-it	1.6%	27.3%
gemma-2-27b-it	81.9%	82.2%

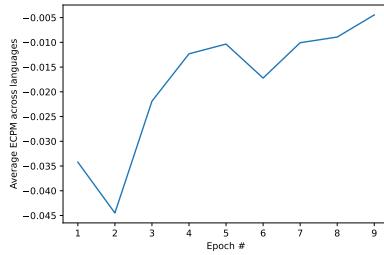
(a) FinQA

(b) MuSR

(c) MGSM

(d) MGSM-SymPy

Figure 2: Comparison between best performing prompt-optimizing methods under DSPy and TACs (full results can be found in Sections H to L). We report the best DSPy result for each task.



(a) Average estimate $\log \mathcal{Z}_\theta$ over validation set inputs versus # of TACStaR epochs over MGSM languages. Note that later epochs (as early as epoch 5) do not have samples from all languages, as some languages early-stopped.

At the end of epoch	Failure rate
1	83.0%
2	1.0%
3	1.6%
4	0.4%

(b) Average MGSM training data parsing failure rate vs # of epochs of TACStaR on gemma-1.1-7b-it. The pattern is **cot-cascade-structure**.

Figure 3: Type compliance during TAC training.

TACStaR compares favorably against the original STaR algorithm on unstructured data. On the MGSM task (Fig. 2c), the original (untyped) STaR algorithm scored an average accuracy of 76.9 and 10.5 (from gemma-2-27b-it and gemma-1.1-7b-it respectively), lower than variants of reasoning TAC patterns on the same dataset. This demonstrates that the structured, typed approach of TACs improves performance over the untyped STaR baseline.

4.3 FLEXIBLE POSTERIOR INFERENCE HELPS TAC PERFORMANCE.

Amortized inference at training time is effective. The Amortized TACStaR algorithm (§3.3) brings consistent improvement over vanilla TACStaR on 3 tasks (Fig. 4a). Notably, the gains are most substantial on FinQA (+5.7

¹²We also compare between TACStaR-adapted and un-adapted models on the same LangFun prompts in §B.2, and find that TACStaR consistently outperforms the un-adapted counterparts.

Task	TACSTaR	Amortized TACSTaR	Base Model	Cl. Gen.
MGSM	82.2	82.4	gemma-1.1-7b-it	62.6 62.1
FinQA	36.0	41.7	gemma-2-27b-it	65.0 51.6
HotPotQA	32.0	34.0		

(a) Comparison between TACSTaR and Amortized TACSTaR on **cot-cascade-structure** / gemma-2-27b-it.

(b) Comparison between classification and unconstrained generation results on MuSR.

Figure 4: Comparison between ‘default’ and more informative inference methods.

points). This suggests that amortized inference is particularly valuable for complex tasks where the initial sampling or fixed rationalization heuristics struggle to find valid latent traces, allowing the model to learn a more effective inference strategy.

Classification with renormalized posterior at inference time is effective. We leverage the probabilistic nature of TACs to estimate the output label posterior $p_\theta(\mathbf{z}_2 \mid \mathbf{z}_1)$ for the MuSR classification task. We achieve this by first estimating the unnormalized probability \tilde{p} for each label using importance sampling, and then renormalizing these estimates over the finite label space (Self-Normalized Importance Sampling). The detailed formulation is described in §R. We output the label with the highest estimated probability. Figure 4b shows that the renormalized-posterior classifier outperforms unconstrained generation on both gemma-1.1-7b-it and gemma-2-27b-it.

4.4 TAC MODELS RAPIDLY ACHIEVE HIGH TYPE COMPLIANCE.

We argued in §3.2 that optimizing the unnormalized likelihood drives the model towards structural compliance. The average MGSM parsing error rate during training (Fig. 3b) suggests that TACs learn compliance fast. We further empirically verify this by estimating the partition function \mathcal{Z}_θ — which represents the total probability mass the model assigns to type-compliant outputs (the Estimated Compliant Probability Mass, ECPM) — throughout training. We estimate $\log \mathcal{Z}_\theta$ on the validation sets of the MGSM benchmark during training of the **cot-cascade-structure** pattern on gemma-1.1-7b-it. We sample 100 generations of entire traces without type-compliant masking per input with temperature = 1, top-p = 1, and top-k set to the vocabulary size. Figure 3a shows that the model rapidly learns to comply with the type constraints. The average $\log \mathcal{Z}_\theta$ approaches -0.005 by epoch 9, corresponding to an ECPM of $\exp(-0.005) \approx 99.5\%$, and thus confirms that the degree of misspecification $(1 - Z_\theta)$ is negligible. Since the difference between unnormalized and normalized likelihood gradients is bounded by a multiplicative factor of $(1 - Z_\theta)$ (Theorem 2), our empirical estimates imply that the difference is indeed small at the end of training, and TACSTaR M-step (§3.2) approaches the true MLE update. Moreover, since $\log \mathcal{Z}_\theta$ is the difference between normalized and unnormalized likelihoods, the small magnitude suggests it is practical to do model selection with unnormalized likelihood directly, after a few epochs of training.

5 CONCLUSION

We have presented Type-Compliant Adaptation Cascades (TACs), a novel probabilistic programming framework designed to empower ML practitioners to design trainable workflows that adapt to data. Our findings demonstrate that TACs’ gradient-based learning paradigm is highly effective, consistently outperforming strong prompt-optimization baselines. Moreover, we also find flexible posterior inference of TACs at both training and inference time help with performance. We also find that empirically, the model learns to comply with type constraints fast in training, justifying the assumptions in our theoretical results. These results underscore the versatility and efficacy of TACs as a scalable paradigm for adapting to complex, reasoning-heavy tasks.

423 REFERENCES
424

- 425 David Belanger and Andrew McCallum. Structured prediction energy networks. In *Proceedings of the 33rd*
426 *International Conference on International Conference on Machine Learning - Volume 48*, ICML'16, pp.
427 983–992. JMLR.org, 2016.
- 428 Luca Beurer-Kellner, Marc Fischer, and Martin Vechev. Prompting is programming: A query language for
429 large language models. *Proc. ACM Program. Lang.*, 7(PLDI), June 2023. doi: 10.1145/3591300. URL
430 <https://doi.org/10.1145/3591300>.
- 431 Luca Beurer-Kellner, Marc Fischer, and Martin Vechev. Guiding llms the right way: fast, non-invasive constrained
432 generation. In *Proceedings of the 41st International Conference on Machine Learning*, ICML'24. JMLR.org,
433 2024.
- 434 Eli Bingham, Jonathan P. Chen, Martin Jankowiak, Fritz Obermeyer, Neeraj Pradhan, Theofanis Karaletsos, Rohit
435 Singh, Paul A. Szerlip, Paul Horsfall, and Noah D. Goodman. Pyro: Deep universal probabilistic programming.
436 *J. Mach. Learn. Res.*, 20:28:1–28:6, 2019. URL <http://jmlr.org/papers/v20/18-403.html>.
- 437 Jörg Bornschein and Yoshua Bengio. Reweighted wake-sleep. *CoRR*, abs/1406.2751, 2014. URL <https://api.semanticscholar.org/CorpusID:10872458>.
- 438 Bowen Cao, Deng Cai, Zhisong Zhang, Yuexian Zou, and Wai Lam. On the worst prompt performance of large
439 language models. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
440 URL <https://openreview.net/forum?id=Mi853QaJx6>.
- 441 Harrison Chase. LangChain, October 2022. URL <https://github.com/langchain-ai/langchain>.
- 442 Zhiyu Chen, Wenhui Chen, Charese Smiley, Sameena Shah, Iana Borova, Dylan Langdon, Reema Moussa, Matt
443 Beane, Ting-Hao Huang, Bryan Routledge, and William Yang Wang. FinQA: A dataset of numerical reasoning
444 over financial data. In Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (eds.),
445 *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 3697–3711,
446 Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics.
447 doi: 10.18653/v1/2021.emnlp-main.300. URL <https://aclanthology.org/2021.emnlp-main.300>.
- 448 Shay B. Cohen, Carlos Gómez-Rodríguez, and G. Satta. Elimination of spurious ambiguity in transition-based
449 dependency parsing. *ArXiv*, abs/1206.6735, 2012. URL <https://api.semanticscholar.org/CorpusID:15438603>.
- 450 David Dohan, Winnie Xu, Aitor Lewkowycz, Jacob Austin, David Bieber, Raphael Gontijo Lopes, Yuhuai Wu,
451 Henryk Michalewski, Rif A. Saurous, Jascha Sohl-dickstein, Kevin Murphy, and Charles Sutton. Language
452 model cascades, 2022. URL <https://arxiv.org/abs/2207.10342>.
- 453 Yixin Dong, Charlie F Ruan, Yaxing Cai, Ruihang Lai, Ziyi Xu, Yilong Zhao, and Tianqi Chen. Xgrammar:
454 Flexible and efficient structured generation engine for large language models. *Proceedings of Machine Learning
455 and Systems* 7, 2024.
- 456 Saibo Geng, Martin Josifoski, Maxime Peyrard, and Robert West. Grammar-constrained decoding for structured
457 NLP tasks without finetuning. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Proceedings of the*
458 *2023 Conference on Empirical Methods in Natural Language Processing*, pp. 10932–10952, Singapore,
459 December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.674. URL
460 <https://aclanthology.org/2023.emnlp-main.674/>.

- 470 Saibo Geng, Hudson Cooper, Michał Moskal, Samuel Jenkins, Julian Berman, Nathan Ranchin, Robert West, Eric
 471 Horvitz, and Harsha Nori. Generating structured outputs from language models: Benchmark and studies, 2025.
 472 URL <https://arxiv.org/abs/2501.10868>.
 473
- 474 Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*, chapter 18. MIT Press, 2016.
 475 <http://www.deeplearningbook.org>.
 476
- 477 Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text degeneration. In
 478 *ICLR*. OpenReview.net, 2020. URL <http://dblp.uni-trier.de/db/conf/iclr/iclr2020.html#HoltzmanBDFC20>.
 479
- 480 Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo,
 481 Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for NLP. In Kamalika Chaudhuri
 482 and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning*,
 483 volume 97 of *Proceedings of Machine Learning Research*, pp. 2790–2799. PMLR, 09–15 Jun 2019. URL
 484 <https://proceedings.mlr.press/v97/houlsby19a.html>.
 485
- 486 Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and
 487 Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *ICLR*, 2022. URL <https://openreview.net/forum?id=nZeVKeeFYf9>.
 488
- 489 Yasaman Jafari, Dheeraj Mekala, Rose Yu, and Taylor Berg-Kirkpatrick. MORL-prompt: An empirical analysis of
 490 multi-objective reinforcement learning for discrete prompt optimization. In Yaser Al-Onaizan, Mohit Bansal,
 491 and Yun-Nung Chen (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2024*, pp. 9878–
 492 9889, Miami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/
 493 2024.findings-emnlp.577. URL <https://aclanthology.org/2024.findings-emnlp.577>.
 494
- 495 Omar Khattab, Keshav Santhanam, Xiang Lisa Li, David Hall, Percy Liang, Christopher Potts, and Matei Zaharia.
 496 Demonstrate-search-predict: Composing retrieval and language models for knowledge-intensive NLP. *arXiv
 497 preprint arXiv:2212.14024*, 2022.
- 498
- 499 Omar Khattab, Arnav Singhvi, Paridhi Maheshwari, Zhiyuan Zhang, Keshav Santhanam, Sri Vardhamanan, Saiful
 500 Haq, Ashutosh Sharma, Thomas T. Joshi, Hanna Moazam, Heather Miller, Matei Zaharia, and Christopher Potts.
 501 Dspy: Compiling declarative language model calls into self-improving pipelines. 2024.
- 502 Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*,
 503 2014.
- 504
- 505 Diederik P. Kingma and Max Welling. Auto-Encoding Variational Bayes. In *2nd International Conference on
 506 Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*,
 507 2014.
- 508
- 509 Vijay Konda and John Tsitsiklis. Actor-critic algorithms. In S. Solla, T. Leen, and
 510 K. Müller (eds.), *Advances in Neural Information Processing Systems*, volume 12. MIT Press,
 511 1999. URL https://proceedings.neurips.cc/paper_files/paper/1999/file/6449f44a102fde848669bdd9eb6b76fa-Paper.pdf.
 512
- 513 John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. Conditional random fields: Probabilistic
 514 models for segmenting and labeling sequence data. In *Proceedings of the Eighteenth International Conference
 515 on Machine Learning*, ICML '01, pp. 282–289, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers
 516 Inc. ISBN 1558607781.

- 517 Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. In
 518 Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (eds.), *Proceedings of the 2021*
 519 *Conference on Empirical Methods in Natural Language Processing*, pp. 3045–3059, Online and Punta Cana,
 520 Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.
 521 emnlp-main.243. URL <https://aclanthology.org/2021.emnlp-main.243/>.
- 522 Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. In Chengqing Zong,
 523 Fei Xia, Wenjie Li, and Roberto Navigli (eds.), *Proceedings of the 59th Annual Meeting of the Association for*
 524 *Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*
 525 (*Volume 1: Long Papers*), pp. 4582–4597, Online, August 2021. Association for Computational Linguistics. doi:
 526 10.18653/v1/2021.acl-long.353. URL <https://aclanthology.org/2021.acl-long.353/>.
- 527 Chu-Cheng Lin and Jason Eisner. Neural particle smoothing for sampling from conditional sequence models. In
 528 Marilyn Walker, Heng Ji, and Amanda Stent (eds.), *Proceedings of the 2018 Conference of the North American*
 529 *Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long*
 530 *Papers*), pp. 929–941, New Orleans, Louisiana, June 2018. Association for Computational Linguistics. doi:
 531 10.18653/v1/N18-1085. URL <https://aclanthology.org/N18-1085/>.
- 532 Xiao Liu, Kaixuan Ji, Yicheng Fu, Weng Tam, Zhengxiao Du, Zhilin Yang, and Jie Tang. P-tuning: Prompt tuning
 533 can be comparable to fine-tuning across scales and tasks. In Smaranda Muresan, Preslav Nakov, and Aline
 534 Villavicencio (eds.), *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*
 535 (*Volume 2: Short Papers*), pp. 61–68, Dublin, Ireland, May 2022. Association for Computational Linguistics.
 536 doi: 10.18653/v1/2022.acl-short.8. URL <https://aclanthology.org/2022.acl-short.8/>.
- 537 Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri Alon, Nouha
 538 Dziri, Shrimai Prabhumoye, Yiming Yang, Shashank Gupta, Bodhisattwa Prasad Majumder, Katherine
 539 Hermann, Sean Welleck, Amir Yazdanbakhsh, and Peter Clark. Self-refine: Iterative refinement with
 540 self-feedback. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural*
 541 *Information Processing Systems*, volume 36, pp. 46534–46594. Curran Associates, Inc., 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/91edff07232fb1b55a505a9e9f6c0ff3-Paper-Conference.pdf.
- 542 Arya McCarthy, Hao Zhang, Shankar Kumar, Felix Stahlberg, and Ke Wu. Long-form speech translation through
 543 segmentation with finite-state decoding constraints on large language models. In Houda Bouamor, Juan Pino,
 544 and Kalika Bali (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2023*, pp. 247–257,
 545 Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-emnlp.
 546 19. URL <https://aclanthology.org/2023.findings-emnlp.19/>.
- 547 Andriy Mnih and Karol Gregor. Neural variational inference and learning in belief networks. In Eric P. Xing and
 548 Tony Jebara (eds.), *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of
 549 *Proceedings of Machine Learning Research*, pp. 1791–1799, Beijing, China, 22–24 Jun 2014. PMLR. URL
 550 <https://proceedings.mlr.press/v32/mnih14.html>.
- 551 Krista Opsahl-Ong, Michael J Ryan, Josh Purtell, David Broman, Christopher Potts, Matei Zaharia, and
 552 Omar Khattab. Optimizing instructions and demonstrations for multi-stage language model programs.
 553 In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Proceedings of the 2024 Conference on*
 554 *Empirical Methods in Natural Language Processing*, pp. 9340–9366, Miami, Florida, USA, November
 555 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-main.525. URL <https://aclanthology.org/2024.emnlp-main.525/>.
- 556 Daiyi Peng, Xuanyi Dong, Esteban Real, Mingxing Tan, Yifeng Lu, Gabriel Bender, Hanxiao Liu, Adam Kraft,
 557 Chen Liang, and Quoc Le. Pyglove: Symbolic programming for automated machine learning. In *Advances in*
 558 *Neural Information Processing Systems (NeurIPS)*, volume 33, pp. 96–108, 2020.

- 564 Gabriel Poesia, Alex Polozov, Vu Le, Ashish Tiwari, Gustavo Soares, Christopher Meek, and Sumit Gulwani.
 565 *Synchromesh: Reliable code generation from pre-trained language models.* In *International Conference on*
 566 *Learning Representations*, 2022. URL <https://openreview.net/forum?id=KmtVD97J43e>.
- 567 Reid Pryzant, Dan Iter, Jerry Li, Yin Lee, Chenguang Zhu, and Michael Zeng. Automatic prompt optimization
 568 with “gradient descent” and beam search. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Proceedings of*
 569 *the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 7957–7968, Singapore,
 570 December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.494. URL
 571 <https://aclanthology.org/2023.emnlp-main.494/>.
- 572 Daniel Ritchie, Paul Guerrero, R. Kenny Jones, Niloy J. Mitra, Adriana Schulz, Karl D. D. Willis, and Jiajun
 573 Wu. Neurosymbolic models for computer graphics. *Computer Graphics Forum*, 42(2):545–568, 2023. doi:
 574 <https://doi.org/10.1111/cgf.14775>. URL <https://onlinelibrary.wiley.com/doi/10.1111/cgf.14775>.
- 575 John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization
 576 algorithms, 2017. URL <https://arxiv.org/abs/1707.06347>.
- 577 Freda Shi, Mirac Suzgun, Markus Freitag, Xuezhi Wang, Suraj Srivats, Soroush Vosoughi, Hyung Won Chung,
 578 Yi Tay, Sebastian Ruder, Denny Zhou, Dipanjan Das, and Jason Wei. Language models are multilingual
 579 chain-of-thought reasoners. In *ICLR*, 2023.
- 580 Dilara Soylu, Christopher Potts, and Omar Khattab. Fine-tuning and prompt optimization: Two great steps that
 581 work better together. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Proceedings of the 2024*
 582 *Conference on Empirical Methods in Natural Language Processing*, pp. 10696–10710, Miami, Florida, USA,
 583 November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-main.597. URL
 584 <https://aclanthology.org/2024.emnlp-main.597/>.
- 585 Zayne Sprague, Fangcong Yin, Juan Diego Rodriguez, Dongwei Jiang, Manya Wadhwa, Prasann Singhal, Xinyu
 586 Zhao, Xi Ye, Kyle Mahowald, and Greg Durrett. To cot or not to cot? chain-of-thought helps mainly on math
 587 and symbolic reasoning, 2024a. URL <https://arxiv.org/abs/2409.12183>.
- 588 Zayne Rea Sprague, Xi Ye, Kaj Bostrom, Swarat Chaudhuri, and Greg Durrett. MuSR: Testing the limits
 589 of chain-of-thought with multistep soft reasoning. In *The Twelfth International Conference on Learning*
 590 *Representations*, 2024b. URL <https://openreview.net/forum?id=jenyYQzue1>.
- 591 Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju,
 592 Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, Johan Ferret, Peter Liu, Pouya
 593 Tafti, Abe Friesen, Michelle Casbon, Sabela Ramos, Ravin Kumar, Charline Le Lan, Sammy Jerome, Anton
 594 Tsitsulin, Nino Vieillard, Piotr Stanczyk, Sertan Girgin, Nikola Momchev, Matt Hoffman, Shantanu Thakoor,
 595 Jean-Bastien Grill, Behnam Neyshabur, Olivier Bachem, Alanna Walton, Aliaksei Severyn, Alicia Parrish, Aliya
 596 Ahmad, Allen Hutchison, Alvin Abdagic, Amanda Carl, Amy Shen, Andy Brock, Andy Coenen, Anthony
 597 Laforge, Antonia Paterson, Ben Bastian, Bilal Piot, Bo Wu, Brandon Royal, Charlie Chen, Chintu Kumar,
 598 Chris Perry, Chris Welty, Christopher A. Choquette-Choo, Danila Sinopalnikov, David Weinberger, Dimple
 599 Vijaykumar, Dominika Rogozińska, Dustin Herbison, Elisa Bandy, Emma Wang, Eric Noland, Erica Moreira,
 600 Evan Senter, Evgenii Eltyshev, Francesco Visin, Gabriel Rasskin, Gary Wei, Glenn Cameron, Gus Martins,
 601 Hadi Hashemi, Hanna Klimczak-Plucińska, Harleen Batra, Harsh Dhand, Ivan Nardini, Jacinda Mein, Jack
 602 Zhou, James Svensson, Jeff Stanway, Jetha Chan, Jin Peng Zhou, Joana Carrasqueira, Joana Iljazi, Jocelyn
 603 Becker, Joe Fernandez, Joost van Amersfoort, Josh Gordon, Josh Lipschultz, Josh Newlan, Ju yeong Ji, Kareem
 604 Mohamed, Kartikeya Badola, Kat Black, Katie Millican, Keelin McDonell, Kelvin Nguyen, Kiranbir Sodhia,
 605 Kish Greene, Lars Lowe Sjoesund, Lauren Usui, Laurent Sifre, Lena Heuermann, Leticia Lago, Lilly McNealus,
 606 Livio Baldini Soares, Logan Kilpatrick, Lucas Dixon, Luciano Martins, Machel Reid, Manvinder Singh, Mark
 607 Iverson, Martin Görner, Mat Velloso, Mateo Wirth, Matt Davidow, Matt Miller, Matthew Rahtz, Matthew
 608

- 611 Watson, Meg Risdal, Mehran Kazemi, Michael Moynihan, Ming Zhang, Minsuk Kahng, Minwoo Park, Mof
 612 Rahman, Mohit Khatwani, Natalie Dao, Nenshad Bardoliwalla, Nesh Devanathan, Neta Dumai, Nilay Chauhan,
 613 Oscar Wahltinez, Pankil Botarda, Parker Barnes, Paul Barham, Paul Michel, Pengchong Jin, Petko Georgiev,
 614 Phil Culliton, Pradeep Kuppala, Ramona Comanescu, Ramona Merhej, Reena Jana, Reza Ardeshir Rokni,
 615 Rishabh Agarwal, Ryan Mullins, Samaneh Saadat, Sara Mc Carthy, Sarah Cogan, Sarah Perrin, Sébastien M. R.
 616 Arnold, Sebastian Krause, Shengyang Dai, Shruti Garg, Shruti Sheth, Sue Ronstrom, Susan Chan, Timothy
 617 Jordan, Ting Yu, Tom Eccles, Tom Hennigan, Tomas Kociský, Tulsee Doshi, Vihan Jain, Vikas Yadav, Vilobh
 618 Meshram, Vishal Dharmadhikari, Warren Barkley, Wei Wei, Wenming Ye, Woohyun Han, Woosuk Kwon,
 619 Xiang Xu, Zhe Shen, Zhitao Gong, Zichuan Wei, Victor Cotrutz, Phoebe Kirk, Anand Rao, Minh Giang,
 620 Ludovic Peran, Tris Warkentin, Eli Collins, Joelle Barral, Zoubin Ghahramani, Raia Hadsell, D. Sculley,
 621 Jeanine Banks, Anca Dragan, Slav Petrov, Oriol Vinyals, Jeff Dean, Demis Hassabis, Koray Kavukcuoglu,
 622 Clement Farabet, Elena Buchatskaya, Sebastian Borgeaud, Noah Fiedel, Armand Joulin, Kathleen Kenealy,
 623 Robert Dadashi, and Alek Andreev. Gemma 2: Improving open language models at a practical size, 2024. URL
 624 <https://arxiv.org/abs/2408.00118>.
- 625 Dustin Tran, Matthew D. Hoffman, Rif A. Saurous, Eugene Brevdo, Kevin Murphy, and David M. Blei. Deep
 626 probabilistic programming. In *International Conference on Learning Representations*, 2017.
- 627 Luong Trung, Xinbo Zhang, Zhanming Jie, Peng Sun, Xiaoran Jin, and Hang Li. ReFT: Reasoning with reinforced
 628 fine-tuning. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual
 629 Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 7601–7614, Bangkok,
 630 Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.410. URL
 631 <https://aclanthology.org/2024.acl-long.410/>.
- 632 Eric Veach and Leonidas J. Guibas. Optimally combining sampling techniques for monte carlo rendering. In
 633 *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH
 634 '95, pp. 419–428, New York, NY, USA, 1995. Association for Computing Machinery. ISBN 0897917014. doi:
 635 10.1145/218380.218498. URL <https://doi.org/10.1145/218380.218498>.
- 636 Greg C. G. Wei and Martin A. Tanner. A monte carlo implementation of the em algorithm and the poor man's
 637 data augmentation algorithms. *Journal of the American Statistical Association*, 85:699–704, 1990. URL
 638 <https://api.semanticscholar.org/CorpusID:123027134>.
- 639 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Huai hsin Chi, F. Xia, Quoc Le, and Denny Zhou.
 640 Chain of thought prompting elicits reasoning in large language models. *ArXiv*, abs/2201.11903, 2022. URL
 641 <https://api.semanticscholar.org/CorpusID:246411621>.
- 642 Ronald J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning.
Mach. Learn., 8(3–4):229–256, May 1992. ISSN 0885-6125. doi: 10.1007/BF00992696. URL
 643 <https://doi.org/10.1007/BF00992696>.
- 644 Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Beibin Li, Erkang Zhu, Li Jiang, Xiaoyun Zhang, Shaokun
 645 Zhang, Jiale Liu, Ahmed Hassan Awadallah, Ryen W White, Doug Burger, and Chi Wang. Autogen: Enabling
 646 next-gen LLM applications via multi-agent conversations. In *First Conference on Language Modeling*, 2024.
 647 URL <https://openreview.net/forum?id=BAaky1hNKS>.
- 648 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen
 649 Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan
 650 Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jing Zhou, Jingren Zhou,
 651 Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li,
 652 Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang,
 653 Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang
 654 Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan
 655 Qiu. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025.

658 Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D.
659 Manning. HotpotQA: A dataset for diverse, explainable multi-hop question answering. In Ellen Riloff, David
660 Chiang, Julia Hockenmaier, and Jun’ichi Tsujii (eds.), *Proceedings of the 2018 Conference on Empirical
661 Methods in Natural Language Processing*, pp. 2369–2380, Brussels, Belgium, October–November 2018.
662 Association for Computational Linguistics. doi: 10.18653/v1/D18-1259. URL <https://aclanthology.org/D18-1259/>.

664 Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. ReAct:
665 Synergizing reasoning and acting in language models. In *International Conference on Learning Representations
666 (ICLR)*, 2023.

668 Mert Yuksekgonul, Federico Bianchi, Joseph Boen, Sheng Liu, Pan Lu, Zhi Huang, Carlos Guestrin, and James
669 Zou. Optimizing generative ai by backpropagating language model feedback. *Nature*, 639:609–616, 2025.

670 Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah Goodman. Star: Bootstrapping reasoning with rea-
671 soning. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Ad-
672 vances in Neural Information Processing Systems*, volume 35, pp. 15476–15488. Curran Associates,
673 Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/639a9a172c044fbb64175b5fad42e9a5-Paper-Conference.pdf.

675 Yongchao Zhou, Andrei Ioan Muresanu, Ziwen Han, Keiran Paster, Silviu Pitis, Harris Chan, and Jimmy Ba.
676 Large language models are human-level prompt engineers. In *The Eleventh International Conference on
677 Learning Representations*, 2023. URL <https://openreview.net/forum?id=92gvk82DE->.

679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704

APPENDICES

Program View	Probabilistic View
τ -typed object	Random variable $\in \Sigma^*$ restricted to strings $\in \text{valid}(\tau)$
LM adaptor with weights θ , with output restricted to τ -typed objects	Unnormalized conditional distribution $p_{LM}(\mathbf{z}_t \mid \mathbf{z}_s; \theta) \mathbb{I}(\mathbf{z}_t \in \text{valid}(\tau))$
Deterministic algorithm $f : \tau_i \rightarrow \tau_o$	Degenerate distribution $\delta_{\text{canon}(f(\text{parse}(x, \tau_i)))}(y)$
<code>parse</code> and <code>canon</code> functions that convert typed objects to/from LM inputs/outputs	Measurable maps between object domain \mathcal{O} and string domain Σ^*
Executing a workflow to obtain $\mathbf{z}_{1\dots M}$	Sampling from joint unnormalized probability $\tilde{p}_\theta(\mathbf{z}_{1\dots M}) = \prod_k \tilde{p}_\theta(\mathbf{z}_{T_k} \mid \mathbf{z}_{S_k})$
Probability that a stochastic workflow succeeds	$\mathcal{Z}_\theta = \Pr_{p_\theta}(\text{all nodes are valid})$

Table 1: Dual semantics: how TAC concepts map between their program and probabilistic views.

A BACKGROUND AND RELATED WORK

TACs sit at the intersection of probabilistic modeling, workflow composition, and LLM adaptation. We organize the related work thematically.

Formalizing and Executing LM Workflows. We approach LLM workflows from the perspective of probabilistic programming. Probabilistic programming languages (PPLs) tailored for machine learning, such as Edward (Tran et al., 2017) and Pyro (Bingham et al., 2019), combine differentiable components with stochastic control flow to define complex distributions. TACs share this goal but address a distinct challenge inherent to typed LLM workflows: enforcing type constraints restricts the LLM’s support, rendering the distribution *unnormalized* ($\mathcal{Z}_\theta \neq 1$). Traditional PPLs typically assume normalized models. TACs draw inspiration from classical structured prediction (Lafferty et al., 2001; Belanger & McCallum, 2016), which provides tools for handling unnormalized models. Our formulation connects these threads, treating type compliance itself as the partition function, enabling a specialized, tractable optimization objective (TACSTaR).

In contrast, programmatic LM frameworks such as DSPy (Khattab et al., 2022; 2024), LMQL (Beurer-Kellner et al., 2023), and LangChain (Chase, 2022) expose LMs through pipelines with declarative constraints. These systems typically optimize the *inputs* (prompts or few-shot exemplars) to a fixed system, rather than casting the entire workflow as a single probabilistic object with learnable continuous parameters and a likelihood objective. While some proposals optimize weights within such pipelines (e.g., BetterTogether (Soylu et al., 2024)), TACs differ fundamentally in their principled probabilistic formulation, enabling theoretically justified training (§3.2) and advanced inference techniques (§3.3).

Optimizing Composed Systems. The standard approach to optimizing LM workflows involves difficult discrete optimization over the space of possible prompts, often addressed through heuristic search (Zhou et al., 2023; Pryzant et al., 2023; Yuksekgonul et al., 2025) or reinforcement learning (Jafari et al., 2024), both of which can suffer from high variance and computational cost.

TACs instead leverage gradient-based optimization. This builds upon methods that adapt LMs for reasoning, such as STaR (Zelikman et al., 2022) and ReFT (Trung et al., 2024), which were inspired by techniques like Chain-of-Thought (CoT) (Wei et al., 2022) and Self-Refine (Madaan et al., 2023). We adopt the spirit of STaR, but generalize it within a hypergraph framework to propose typed, multi-step rationalizations (§3.2). Furthermore, we introduce an amortized variant that learns to propose rationalizations, rather than relying solely on heuristics (§3.3).

752 Enforcing Structure and Compliance. To improve output reliability, various methods enforce grammar-based
 753 constraints during LLM generation (Poesia et al., 2022; Geng et al., 2023; McCarthy et al., 2023; Beurer-Kellner
 754 et al., 2024; Geng et al., 2025). These methods generally modify *local* conditional distributions over next tokens at
 755 inference time, masking out continuations incompatible with the grammar. In contrast, our objective learns
 756 parameters so that type-compliant trajectories carry increasing probability mass *globally*, improving both validity
 757 and task accuracy through training.

758 Parameter-efficient adaptation. LoRA and related PEFT methods (Houlsby et al., 2019; Hu et al., 2022; Li
 759 & Liang, 2021; Lester et al., 2021; Liu et al., 2022) enable light-weight adaptation. We use small adaptors to
 760 highlight data-efficiency and show that gains stem from *typed workflow learning* rather than sheer capacity.

762 Connection to Reinforcement Learning. The TACSTaR training procedure (§3.2) can also be viewed through
 763 the lens of policy optimization. As Zelikman et al. (2022) observed, the STaR objective closely resembles the
 764 REINFORCE algorithm (Williams, 1992). The M-step in TACSTaR can be interpreted as optimizing the workflow
 765 policy under REINFORCE with a binary reward for generating the correct output.

We adopt the MC-EM framing as it provides a principled approach for likelihood maximization in the presence of
 767 annotated output data. While advanced RL techniques (*e.g.*, PPO (Schulman et al., 2017) or actor-critic methods
 768 (Konda & Tsitsiklis, 1999)) might be applicable, they introduce complexity, such as training value functions,
 769 which are difficult to estimate over complex, typed latent spaces. Furthermore, the exploration challenge in sparse
 770 reward settings is significantly mitigated by the rationalization heuristic and the inference network in Amortized
 771 TACSTaR (§3.3), which guide sampling towards successful trajectories using known outputs.

773 B ADDITIONAL STUDIES ON WORKFLOW PATTERN DESIGN

775 In this section, we conduct additional experiments that vary the pattern structures, and evaluate how such changes
 776 affect performance. Specifically, we would like to answer the following questions:

- 778 • **(§B.2) Is adaptation with reasoning workflows effective?** The TAC framework gives practitioners great
 779 freedom in designing a workflow that reason in the process. We hypothesize that adapting with such
 780 explicit structures improves performance on tasks that require complex reasoning.
- 781 • **(§B.3) How do TAC design variations affect performance?** We evaluate how such TAC design variations
 782 for the same task affect performance.

783 B.1 END-TO-END TRAINABLE WORKFLOWS AS TACs.

785 The declarative and flexible nature of TACs enable practitioners to rapidly implement end-to-end trainable
 786 workflows. We implement some common patterns as TACs:

- 788 • **Direct adaptation** of an LM to the downstream task without any latent structure corresponds to common
 789 supervised PEFT methods surveyed in §A. The **direct** pattern (Fig. 5a) is a singleton TAC with no latent
 790 nodes.
- 791 • **Adapting with latent rationales** corresponds to patterns that learn to generate rationales for the task at
 792 hand Zelikman et al. (2022). There are several possible TAC structure designs that incorporate rationales:
 793 for example, **cot-type-structure** (Fig. 5b) maps the input to a rationale-output typed object, from which
 794 the task output is deterministically extracted. Alternatively, **cot-cascade-structure** (Fig. 1a) introduce
 795 rationales as distinct nodes in the TAC hypergraph, which transforms into the task output under an adaptor.
- 796 • **Trainable self-refinement** refers to an end-to-end trainable variant of self-refine (Madaan et al., 2023),
 797 where the model first sketches a task output, and iteratively refine it. Without TAC, a practitioner would
 798 have to resort to manually writing tedious postprocessing functions for the intermediate results. On the
 other hand, the TAC counterpart **refine-structure** (Fig. 6 in §F) is straightforward.

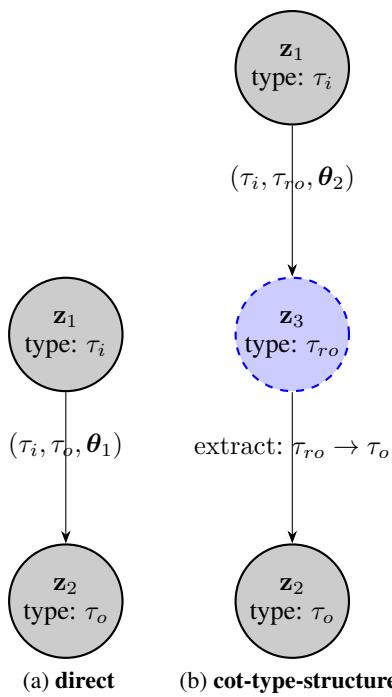


Figure 5: Workflow patterns experimented in this paper, with increasing structural complexity from left to right. In the most complicated pattern **expression-cascade-structure** we illustrate the workflow with example node values. Dashed-boundary nodes indicate variables that are not observed at training time. And solid-boundary nodes indicate nodes with training time observable values. A main message of this work is that instead of defining workflows imperatively as fixed-parameter systems, **we treat an entire typed workflow as a single probabilistic program, whose parameters are lightweight PEFT modules, allowing end-to-end training with latent variables.**

For the MGSM-SymPy task, we experiment with the **expression-cascade-structure** pattern (Fig. 1b), which additionally imposes the constraint that the output must be a rational number represented by an arithmetic expression tree. Such type constraints often reflect business logic (for example, we expect the MGSM dataset to have rational number answers), and may be necessary when the TAC forms a component in a larger system.

B.2 EFFECTIVENESS OF ADAPTATION WITH REASONING WORKFLOWS

To evaluate whether adaptation with reasoning workflows is effective, we compare **cot-cascade-structure**, and **refine-structure** TACs against **direct** on the 3 tasks MGSM, FinQA and HotPotQA, on base models gemma-2-27b-it and gemma-1.1-7b-it. Table 2 shows that both **cot-cascade-structure** significantly outperforms **direct** on MGSM and FinQA on both gemma-2-27b-it and gemma-1.1-7b-it. But **cot-cascade-structure** slightly underperforms **direct** on HotPotQA. These results largely agree with the meta study done by Sprague et al. (2024a), which also reported that tasks that require arithmetic and symbolic reasoning, such as MGSM and FinQA, benefit the most from CoT, while a huge portion of previous work saw that CoT degrades performance for multi-hop QA. However, we note that the **refine-structure** TAC (Fig. 6) consistently

846 outperform the **direct** baseline in all 3 tasks on `gemma-2-27b-it`, showcasing the effectiveness of the adaptive
 847 refinement paradigm.
 848

849 850 851 Dataset	852 853 854 gemma-2-27b-it			855 856 857 858 859 860 gemma-1.1-7b-it	
	852 853 854 direct	852 853 854 cot-cascade-structure	852 853 854 refine-structure	852 853 854 direct	852 853 854 cot-cascade-structure
MGSM	24.7		82.2	78.6	5.1
FinQA	17.3		36.0	23.7	3.0
HotPotQA	34.0		32.0	39.0	—

856 Table 2: Comparison between **direct** and reasoning workflows. For the MGSM dataset, we report per-language
 857 accuracies in Table 5. The difference between best performing runs and **direct** are statistically significant/marginally
 858 significant: for MGSM and FinQA $p < 0.05$ (both `gemma-2-27b-it` and `gemma-1.1-7b-it`), and for
 859 HotPotQA $p = 0.07$ under paired permutation tests. Per-language accuracy numbers of the MGSM dataset are in
 860 §H.

861
 862 **Task adaptation with TACSTaR is effective.** To evaluate whether the efficacy of TACs can be attributed to our
 863 proposed TACSTaR method, we also compare adapted TAC workflows against those with the same hypergraph
 864 structure, but with un-adapted weights (*i.e.*, all adaptors in the TAC use base model weights). Both TACSTaR trained
 865 and un-adapted models use the same structured LangFun prompts that are similar to examples listed in §N. The
 866 significant gap between adapted and un-adapted results in Table 3 indicate that the TACSTaR algorithm is effective.
 867 Notably, un-adapted models still outperform **direct** workflows (listed in Table 2), indicating that LangFun’s
 868 type-inducing prompts can invoke somewhat effective test-time computation over the TAC hypergraph structure.
 869

870 Task	871 Structure	872 TACSTaR	873 Un-adapted
MGSM	cot-cascade-structure	82.2	45.4
MGSM	cot-type-structure	80.4	74.7
MGSM-SymPy	expression-cascade-structure	75.9	69.5
FinQA	cot-cascade-structure	36.0	13.0
HotPotQA	refine-structure	39.0	24.0

877 Table 3: Comparison between TACSTaR-adapted and un-adapted `gemma-2-27b-it`. The differences are all
 878 statistically significant ($p < 0.05$) under paired permutation tests.
 879

880 B.3 EFFECTS OF DIFFERENT TAC DESIGNS

882 **Decoupling rationale and output modeling helps performance.** **cot-cascade-structure** (Fig. 1a) achieves a
 883 higher score than **cot-type-structure** (Fig. 5b) on the MGSM task (Table 4), suggesting that modeling the rationale
 884 and task output generation with distinct adaptors helps performance. By using distinct adaptors, the workflow
 885 allows specialization: the first adaptor focuses on reasoning, while the second specializes in synthesis, reducing the
 886 complexity burden on a single monolithic step. The positive result again highlights how the TAC formalism can
 887 help practitioners iterate and experiment with different multi-adaptor cascade designs, which would be tedious
 888 otherwise.
 889

890 **Robustness to Semantic Constraints.** Comparing performance on MGSM and the more constrained MGSM-
 891 SymPy task reveals a key advantage of the TAC framework’s robustness. As shown in Table 4, the best-performing
 892 TAC model sees a modest performance drop, from 82.2% on MGSM to 75.9% on MGSM-SymPy, when required

893 to generate a valid symbolic expression.¹³ This contrasts sharply with the prompt-optimizing baseline (Fig. 2). The
 894 best DSPy configuration experiences a much more significant degradation, plummeting from 81.9% on MGSM to
 895 just 57.1% on MGSM-SymPy. The substantially smaller performance drop for TACs underscores the brittleness of
 896 discrete prompt optimization when faced with strict structural requirements. The TAC framework’s gradient-based
 897 adaptation within a typed system proves to be significantly more resilient, making it a more reliable paradigm for
 898 tasks demanding structural compliance.

	MGSM	MGSM-SymPy
cot-type-structure	cot-cascade-structure	expression-cascade-structure
80.4	82.2	75.9

900 Table 4: Effects of different TAC designs on the MGSM dataset, demonstrating the impact of workflow structure
 901 on performance. The **cot-cascade-structure** (which decouples rationale generation from the final answer
 902 synthesis) outperforms the monolithic **cot-type-structure**. The **expression-cascade-structure** result shows strong
 903 performance on the more constrained MGSM-SymPy task.

910 C ALGORITHMS

911 C.1 FORWARD AND BACKWARD

912 Algorithm 1 (`forward`) executes the probabilistic program represented by a TAC $C = (\mathbf{Z}, \mathbf{E})$. Starting from a
 913 given input node value \mathbf{z}_1^* , the algorithm traverses the hypergraph following a topological order, and terminates
 914 when all edges $\in \mathbf{Z}$ have been visited. `forward` takes C and \mathbf{z}_1^* as input arguments. `forward` also takes the
 915 following as arguments:

- 916 • sampler configuration κ for different sampling techniques, *e.g.*, varying temperature, nucleus, and top- k
 917 sampling
- 918 • maximum number of sampling attempts

919 Algorithm 2 (`backward`) takes as input (C, \mathbf{Z}^*) , where $C = (\mathbf{Z}, \mathbf{E})$ where $\mathbf{E} = (e_1 \dots e_K)$ is a TAC, and
 920 \mathbf{Z}^* are value assignments of \mathbf{Z} . We assume the log probability $p_{LM}(\mathbf{y} \mid \mathbf{x}; \theta_k)$ is auto-differentiable with
 921 regard to all adaptor hyperedges in a TAC. Algorithm 2 returns unnormalized log joint probabilities of \mathbf{Z}^*
 922 under C : $\log \tilde{p}_\theta(\mathbf{Z}^*)$, the per-node generation log probabilities $(\log p_\theta(z_2 \mid \cdot) \dots \log p_\theta(z_M \mid \cdot))$, and also
 923 gradients of LM adaptors: $\nabla_{\theta_k} \log \tilde{p}_\theta(\mathbf{Z}^*)$ for adaptor hyperedges’ indices k . We note that `backward` is easily
 924 parallelizable: all adaptor edges can be processed at the same time.

925 C.2 TACSTAR

926 The TACStaR algorithm (Algorithm 3) takes as input $(C, \{x_i^*, y_i^* \mid i \in [1..D_{\text{train}}]\})$, where C is the TAC to train,
 927 and $\{(x_i^*, y_i^*) \mid i \in [1..D_{\text{train}}]\}$ is the training dataset. As we described in §3.2, TACStaR uses a ‘fallback TAC’
 928 heuristics in hope to obtain a sample when the forward algorithm fails.

929 **Building Fallback TAC.** Given a TAC $C = (\mathbf{Z}, \mathbf{E})$ with input node and output node typed τ_i and τ_o respectively,
 930 we build its fallback TAC $C_{\text{fallback}} = (\mathbf{Z}', \mathbf{E}')$ (denoted as the function `build_fallback` in Algorithm 3) as
 931 follows:

932 ¹³Sample expressions generated under **expression-cascade-structure** are listed in §M.

940 **Algorithm 1** TAC Forward Algorithm (forward)

941 **Input:** TAC cascade $C = (\mathbf{Z}, \mathbf{E})$ where $\mathbf{Z} = \{\mathbf{z}_1 \dots \mathbf{z}_M\}$ and $\mathbf{E} = \{\mathbf{e}_1 \dots \mathbf{e}_K\}$, input object: \mathbf{z}_1^* , sampler
 942 configuration κ, N_{\max} for maximum number of sampling attempts.
 943 **Output:** Sampled values $(\mathbf{z}_2^*, \dots, \mathbf{z}_M^*)$.

944 1: Determine a topological ordering of edges in \mathbf{E} . Let the sorted hyperedges be $e'_1 \dots e'_K$.
 945 2: $\mathbf{Z}_{\text{already_sampled}}^* \leftarrow \{\mathbf{z}_1^*\}$.
 946 3: **for** $k \in [1..K]$ **do**
 947 4: Assert the source nodes of e'_k is a subset of $\mathbf{z}_{\text{already_sampled}}$.
 948 5: **if** $e'_k = (\tau_i, \tau_o, \theta)$ is a type-constrained LM adaptor **then**
 949 6: # type-constrained LM adaptors have a single source node and a single target node.
 950 7: $\mathbf{x} \leftarrow$ canonicalized representation of e'_k 's source node.
 951 8: **while** number of attempts $\leq N_{\max}$ **do**
 952 9: Try draw $\mathbf{y} \sim p_{LM}(\cdot | \mathbf{x}; \theta, \kappa)$
 953 10: **if** $\text{parse}(\mathbf{y}, \tau_o) \neq \text{error}$ **then**
 954 11: $t \leftarrow$ index of e'_k 's target node.
 955 12: $\mathbf{z}_t^* \leftarrow \mathbf{y}$
 956 13: $\mathbf{Z}_{\text{already_sampled}}^* \leftarrow \mathbf{Z}_{\text{already_sampled}}^* \cup \{\mathbf{z}_t^*\}$
 957 14: break
 958 15: **end if**
 959 16: **end while**
 960 17: **else if** e'_k is a deterministic algorithm f **then**
 961 18: # In this work we assume f 's inputs and outputs are sorted by node index in C .
 962 19: $\mathbf{O}_{f\text{input}} \leftarrow$ parsed objects of e'_k 's source nodes, sorted by node index.
 963 20: $\mathbf{O}_{f\text{output}} \leftarrow f(\mathbf{O}_{f\text{input}})$
 964 21: $\mathbf{Z}_{f\text{output}}^* \leftarrow$ canonicalized representations of objects $\in \mathbf{O}_{f\text{output}}$, sorted by node index.
 965 22: $\mathbf{Z}_{\text{already_sampled}}^* \leftarrow \mathbf{Z}_{\text{already_sampled}}^* \cup \mathbf{Z}_{f\text{output}}^*$
 966 23: **end if**
 967 24: **end for**
 968 25: **return** $\mathbf{Z}_{\text{already_sampled}}^* - \{\mathbf{z}_1^*\}$.

- 969
- 970 • The input node of C_{fallback} : \mathbf{z}'_1 is of the product type $\tau_{io} = \tau_i \times \tau_o$, representing a data container that
 971 holds one object of type τ_i and another object of type τ_o .
 - 972 • All other nodes $\in \mathbf{Z}$ have their counterpart nodes in Z' (with the same types and indices).
 - 973 • We copy each hyperedge $e \in \mathbf{E}$ over to \mathbf{E}' , connecting nodes with the same indices. In the case that e is
 974 a deterministic algorithm hyperedge, and has \mathbf{z}_1 as one of its source nodes, we modify the counterpart
 975 hyperedge e' to have a deterministic algorithm that first extracts the original object $\text{parse}(\mathbf{z}_1)$ from
 976 $\text{parse}(\mathbf{z}'_1)$, and then pass $\text{parse}(\mathbf{z}_1)$ to the original algorithm as input.

977 Adaptors in C_{fallback} use no-op weights, falling back to the behavior of the base model. We denote such no-op
 978 weights as θ_0 . For example, Fig. 7 is the C_{fallback} for Fig. 5b.

981 C.3 AMORTIZED TACSTAR

982 The Amortized TACStaR algorithm (Algorithm 4) builds upon Algorithm 3 to introduce an inference network TAC.
 983 While C_{fallback} used fixed no-op weights that behave identical to the base language model, Amortized TACStaR
 984 leverages an inference network TAC C' with trainable parameters.

987 **Algorithm 2** TAC Backward Algorithm (backward)

988 **Input:** $C = (\mathbf{Z}, \mathbf{E})$ and sample $\mathbf{Z}^* = \{\mathbf{z}_1^*, \mathbf{z}_2^*, \dots, \mathbf{z}_M^*\}$

989 **Output:** $(\log \tilde{p}_\theta(\mathbf{Z}^*), (\log p_\theta(z_2 | \cdot) \dots \log p_\theta(z_M | \cdot)), \{\nabla_{\theta_k} \log \tilde{p}_\theta(\mathbf{Z}^*) | e_k \in \mathbf{E} \text{ is an adaptor hyperedge}\})$

990 1: Initialize log-probability accumulator $\mathcal{L} \leftarrow 0$.

991 2: **for** each LM adaptor hyperedge $e_k = (\tau_i, \tau_o, \theta_k)$ **do**

992 3: Let $\mathbf{z}_i^* \in \mathbf{Z}^*$, $\mathbf{z}_o^* \in \mathbf{Z}^*$ be the sample value of e_k 's input and output nodes \mathbf{z}_i (typed τ_i) and \mathbf{z}_o respectively.

993 4: $(\ell, \mathbf{g}_k) \leftarrow \text{peft_backward}(\log p_{LM}(\mathbf{z}_o^* | \text{canon}(\text{parse}(\mathbf{z}_i^*, \tau_i)); \theta))$.

994 5: $\mathcal{L} \leftarrow \mathcal{L} + \ell$

995 6: keep track of ℓ by its node index.

996 7: **end for**

997 8: **# For nodes from deterministic hyperedges, set log prob to 0 as they have no learnable parameters.**

998 9: **return** $(\mathcal{L}, (\log p_\theta(z_2 | \cdot) \dots \log p_\theta(z_M | \cdot)), \{\mathbf{g}_k | e_k \in \mathbf{E} \text{ is an adaptor hyperedge}\})$.

1001
1002
1003 **Building the inference network C' .** Given a TAC $C = (\mathbf{Z}, \mathbf{E})$ with input node and output node typed τ_i and τ_o respectively, we build the adaptive fallback TAC $C' = (\mathbf{Z}', \mathbf{E}')$ (denoted as the function `build_infer_net` in Algorithm 4). At a high level, every adaptor hyperedge that generates latent variables in C is mapped into a counterpart in C' that also depends on both observed a τ_i -typed input and a τ_o -typed output, now encoded as \mathbf{z}'_1 , typed τ_{io} . Specifically we build C' with the following procedure:

- 1008 • The input node of C' : \mathbf{z}'_1 is of the product type $\tau_{io} = \tau_i \times \tau_o$, as with `build_fallback`.
- 1009 • All nodes $\in \mathbf{Z}$ have their counterpart nodes in Z' (with the same types and indices), except for $\{\mathbf{z}_1, \mathbf{z}_2\}$.¹⁴
- 1010 • For each hyperedge $e \in \mathbf{E}$,
- 1011 – In the case that e is a deterministic algorithm hyperedge, and has \mathbf{z}_1 as one of its source nodes, we add a counterpart hyperedge e' that connect counterpart nodes in \mathbf{Z}' , with its deterministic algorithm modified to typecheck, as `build_fallback`.
- 1012 – Otherwise, e is an adaptor hyperedge. Denoting its source node as \mathbf{z}_s and target node as \mathbf{z}_t :
- 1013 * If $\mathbf{z}_t = \mathbf{z}_2$, we continue since \mathbf{z}_t has no counterpart C' .
- 1014 * If $\mathbf{z}_s = \mathbf{z}_1$ and $\mathbf{z}_t \neq \mathbf{z}_2$, we add a counterpart hyperedge $e' = (\tau_s, \tau_t, \theta_{\text{new}})$ connecting counterpart nodes \mathbf{z}'_s and \mathbf{z}'_t . θ_{new} indicates the parameter vector of a new LM adaptor.
- 1015 * Otherwise, $\mathbf{z}_s \neq \mathbf{z}_1$ and $\mathbf{z}_t \neq \mathbf{z}_2$. In this case, we create e' to be an adaptor that is conditioned on both \mathbf{z}'_s and \mathbf{z}'_1 . To achieve this goal, we introduce into C' a helper node \mathbf{z}''_s typed $\tau_{ios} = \tau_i \times \tau_o \times \tau_s$, and a helper hyperedge e'' that has source nodes $\{\mathbf{z}'_1, \mathbf{z}'_s\}$, and target node $\{\mathbf{z}''_s\}$. e'' is a deterministic edge that combines values in \mathbf{z}'_1 and \mathbf{z}'_s into the 3-object container \mathbf{z}''_s . Finally, we add e' that connects \mathbf{z}''_s to \mathbf{z}_t as the adaptor transformation $(\tau_{ios}, \tau_t, \theta_{\text{new}})$, where θ_{new} again indicates the parameter vector of a new LM adaptor.

1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1030
1031
1032
1033
Adaptors in C' are new adaptors. And we train C alternately with C' in Algorithm 4. The algorithm to train C' is listed in Algorithm 5.

C.4 UPDATING C'

We train the inference network C' to better approximate the posterior distribution defined by C alternately (§3.3). In other words, we update adaptor parameters in C' so that sampled latent variables of C' ($(\hat{\mathbf{z}}_3, \dots, \hat{\mathbf{z}}_M)$

¹⁴We arbitrarily designate a node $\in \mathbf{Z}'$ that does not have an outgoing hyperedge as the output node for syntactic conformity.

1034 **Algorithm 3** TACSTaR Training Algorithm

1035 **Input:** Training pairs $\mathcal{D}_{\text{train}} = \{(x_i^*, y_i^*) \mid i \in [1..|\mathcal{D}_{\text{train}}|]\}$, TAC C , sampler configuration κ .

1036 1: $C_{\text{fallback}} \leftarrow \text{build_fallback}(C)$

1037 2: **for** epoch in $[1.. \text{num_epochs}]$ **do**

1038 3: $S \leftarrow \{\} \# \text{Successful samples}$

1039 4: **for** training pair $(x^*, y^*) \in \mathcal{D}_{\text{train}}$ **do**

1040 5: $\mathbf{z}_1^* \leftarrow \text{canon}(x^*)$

1041 6: **# E-step (Sampling Latent Variables):**

1042 7: $(\hat{\mathbf{z}}_2 \dots \hat{\mathbf{z}}_M) \leftarrow \text{Forward}(C, \mathbf{z}_1^*)$.

1043 8: **# Filtering (Validity Check):**

1044 9: Initialize $\text{error_flag} \leftarrow \text{false}$.

1045 10: Set $\text{error_flag} \leftarrow \text{true}$ if errors in E-step or $\text{parse}(\hat{\mathbf{z}}_2) \neq y^*$.

1046 11: **# Heuristics Fallback (Addressing Forward Failure):**

1047 12: **if** error_flag is true **then**

1048 13: $\mathbf{z}'_1^* \leftarrow \text{canon}((x^*, y^*))$

1049 14: $(\hat{\mathbf{z}}'_2 \dots \hat{\mathbf{z}}'_M) \leftarrow \text{forward}(C_{\text{fallback}}, \mathbf{z}'_1^*)[0]$.

1050 15: **if** no error was raised and $\text{parse}(\hat{\mathbf{z}}'_2) = y^*$ **then**

1051 16: $(\hat{\mathbf{z}}_2 \dots \hat{\mathbf{z}}_M) \leftarrow (\hat{\mathbf{z}}'_2 \dots \hat{\mathbf{z}}'_M)$

1052 17: Set $\text{error_flag} \leftarrow \text{false}$.

1053 18: **end if**

1054 19: **end if**

1055 20: **if** error_flag is false **then**

1056 21: $S \leftarrow S \cup \{(\mathbf{z}_1^*, \hat{\mathbf{z}}_2 \dots \hat{\mathbf{z}}_M)\}$

1057 22: **end if**

1058 23: **end for**

1059 24: **# M-step (Parameter Update):**

1060 25: **for** $(\mathbf{z}_1^*, \hat{\mathbf{z}}_2 \dots \hat{\mathbf{z}}_M) \in S$ **do**

1061 26: $\mathbf{G} \leftarrow \text{backward}(C, (\mathbf{z}_1^*, \hat{\mathbf{z}}_2 \dots \hat{\mathbf{z}}_M))[2]$

1062 27: $\text{optimize}(C, \mathbf{G})$

1063 28: **end for**

1064 29: **end for**

1065 obtained using $\text{forward}(C', \text{canon}(x^*), \kappa)$ follow the normalized distributions under C (obtained using
1066 $\text{backward}(C, (\text{canon}(x^*), \text{canon}(y^*), \hat{\mathbf{z}}_3, \dots, \hat{\mathbf{z}}_M))$). To promote diversity of samples, we additionally
1067 obtain samples from C_{fallback} (§C.2). Let $\mathbf{Z} = (\mathbf{z}_3^*, \dots, \mathbf{z}_M^*)$ be a sample out of G collected samples
1068 ($\mathbf{Z}^{(1)}, \dots, \mathbf{Z}^{(G)}$) from C_{fallback} and C' . We approximate the posterior probability of \mathbf{Z} under C , conditioning on
1069 $\mathbf{z}_1^* = \text{canon}(x^*)$, $\mathbf{z}_2^* = \text{canon}(y^*)$ under the balance heuristic (Veach & Guibas, 1995) as

$$\hat{p}_{\text{posterior}}(\mathbf{Z}) \propto \frac{(N_{\text{fallback}} + N_{\text{infer}})\tilde{p}_{\text{model}}}{N_{\text{fallback}}p_{\text{fallback}} + N_{\text{infer}}p_{\text{infer}}}, \quad (5)$$

1070 where $\tilde{p}_{\text{model}} = \tilde{p}_C(\mathbf{z}_1^*, \mathbf{z}_2^*, \mathbf{z}_3^*, \dots, \mathbf{z}_M^*)$, $p_{\text{fallback}} = \prod_{m=3}^M p_{LM}(\mathbf{z}_m^* \mid \mathbf{z}_m^* \text{ source node}; \theta_0)$, and $p_{\text{infer}} =$
1071 $\prod_{m=3}^M p_{LM}(\mathbf{z}_m^* \mid \mathbf{z}_m^* \text{ source node}; \theta_{\text{new}})$. These values are all obtained using the backward algorithm.¹⁵ We
1072 denote the number of samples attempted (including errors) on $C_{\text{fallback}} = N_{\text{fallback}}$, the number of samples attempted
1073 (including errors) on $C' = N_{\text{infer}}$. $\hat{p}_{\text{posterior}}$ is normalized over the mixture so that $\sum_{g=1}^G \hat{p}_{\text{posterior}}(\mathbf{Z}^{(g)}) = 1$.

1074
1075
1076
1077
1078
1079 ¹⁵backward algorithm as presented in this work computes both gradients and probabilities. In our implementation we do
1080 not compute gradients when they are not needed; but we omit this subtlety in Algorithm 2.

1081 **Algorithm 4** Amortized TACStaR Training Algorithm

1082 **Input:** Training pairs $\mathcal{D}_{\text{train}} = \{(x_i^*, y_i^*) \mid i \in [1..|\mathcal{D}_{\text{train}}|\}\}, \text{TAC } C, \text{sampler configuration } \kappa.$

1083 1: $C' \leftarrow \text{build_infer_net}(C)$

1084 2: **for** epoch in $[1.. \text{num_epochs}]$ **do**

1085 3: $S \leftarrow \{\} \# \text{Successful samples}$

1086 4: **for** training pair $(x^*, y^*) \in \mathcal{D}_{\text{train}}$ **do**

1087 5: $\mathbf{z}_1^* \leftarrow \text{canon}(x^*)$

1088 6: **# E-step (Sampling Latent Variables):**

1089 7: $(\hat{\mathbf{z}}_2 \dots \hat{\mathbf{z}}_M) \leftarrow \text{Forward}(C, \mathbf{z}_1^*).$

1090 8: **# Filtering (Validity Check):**

1091 9: Initialize $\text{error_flag} \leftarrow \text{false}.$

1092 10: Set $\text{error_flag} \leftarrow \text{true}$ if errors in E-step or $\text{parse}(\hat{\mathbf{z}}_2) \neq y^*.$

1093 11: **# Heuristics Fallback (Addressing Forward Failure):**

1094 12: **if** error_flag is true **then**

1095 13: $\mathbf{z}'_1 \leftarrow \text{canon}((x^*, y^*))$

1096 14: $(\hat{\mathbf{z}}'_2 \dots \hat{\mathbf{z}}'_M) \leftarrow \text{forward}(C_{\text{fallback}}, \mathbf{z}'_1)[0].$

1097 15: **if** no error was raised and $\text{parse}(\hat{\mathbf{z}}'_2) = y^*$ **then**

1098 16: $(\hat{\mathbf{z}}_2 \dots \hat{\mathbf{z}}_M) \leftarrow (\hat{\mathbf{z}}'_2 \dots \hat{\mathbf{z}}'_M)$

1099 17: Set $\text{error_flag} \leftarrow \text{false}.$

1100 18: **end if**

1101 19: **end if**

1102 20: **if** error_flag is true **then**

1103 21: $(\hat{\mathbf{z}}_3 \dots \hat{\mathbf{z}}_M) \leftarrow \text{forward}(C', \mathbf{z}_1^*)[0]$

1104 22: Set $\text{error_flag} \leftarrow \text{false}$ if no errors in previous step.

1105 23: **end if**

1106 24: **if** error_flag is false **then**

1107 25: $S \leftarrow S \cup \{(\mathbf{z}_1^*, \mathbf{z}_2^*, \hat{\mathbf{z}}_3, \dots, \hat{\mathbf{z}}_M)\}$

1108 26: **end if**

1109 27: **end for**

1110 28: **# M-step (Parameter Update):**

1111 29: **for** $(\mathbf{z}_1^*, \hat{\mathbf{z}}_2 \dots \hat{\mathbf{z}}_M) \in S$ **do**

1112 30: $\mathbf{G} \leftarrow \text{backward}(C, (\mathbf{z}_1^*, \hat{\mathbf{z}}_2 \dots \hat{\mathbf{z}}_M))[2]$

1113 31: $\text{optimize}(C, \mathbf{G})$

1114 32: **end for**

1115 33: $C' \leftarrow \text{update inference network } C' (\S C.4).$

1116 34: **end for**

1117 Algorithm 5 updates adaptors in C' to bring its unnormalized distribution closer to Eq. (5). Since the self-
1118 normalized approximation of the posterior distribution is consistent but biased, we require minimum numbers of
1119 samples from C' and C_{fallback} .

1121 **D FORMAL STATEMENTS AND PROOFS REGARDING TYPE COMPLIANCE**

1122 **Well-specifiedness.** Let $C = (\mathbf{Z}, \mathbf{E})$. We define well-specifiedness for TAC: we say $\boldsymbol{\theta} = \{\boldsymbol{\theta}_1 \dots \boldsymbol{\theta}_K\}$ is well-specified if for every LM adaptor $e_k = (\tau_i, \tau_o, \boldsymbol{\theta}_k) \in \mathbf{E}$ and for every valid canonicalized string x of type τ_i , the LM distribution p_{LM} only has support over valid outputs of type τ_o . Formally, $\forall \text{ valid } x, \sum_{\mathbf{y} \in \mathcal{D}_{\text{valid}}(\tau_o)} p_{LM}(\mathbf{y} \mid x; \boldsymbol{\theta}_k) = 1$ iff $\boldsymbol{\theta}$ is well-specified.

1128 **Algorithm 5** update_infer_net

1130 **Input:** Training pair (x^*, y^*) , model TAC C , sampler configuration κ , inference network C' , non-adaptive
 1131 fallback C_{fallback} , number of samples from C_{fallback} : G_{fallback} , number of samples from C' : G_{infer} .
 1132 1: $\mathbf{z}'_1 \leftarrow \text{canon}((x^*, y^*))$, $\mathbf{z}_1^* \leftarrow \text{canon}(x^*)$, $\mathbf{z}_2^* \leftarrow \text{canon}(y^*)$.
 1133 2: $\mathbf{Z}_{\text{collected}} \leftarrow []$
 1134 3: **# In our implementation we give up and raise an error after 30 unsuccessful attempts.**
 1135 4: **while** number of successful samples from $C_{\text{fallback}} < G_{\text{fallback}}$ **do**
 1136 5: Try $(\hat{\mathbf{z}}_2, \hat{\mathbf{z}}_3, \dots, \hat{\mathbf{z}}_M) \leftarrow \text{forward}(C_{\text{fallback}}, \mathbf{z}'_1, \kappa, 1)$
 1137 6: **if** previous step succeeded **then**
 1138 **# We discard** $\hat{\mathbf{z}}_2$ **from** C_{fallback} .
 1139 Append $(\hat{\mathbf{z}}_3, \dots, \hat{\mathbf{z}}_M)$ to $\mathbf{Z}_{\text{collected}}$.
 1140 **end if**
 1141 7: **end while**
 1142 11: $N_{\text{fallback}} \leftarrow$ numbers of attempts on C_{fallback}
 1143 12: **while** number of successful samples from $C' < G_{\text{infer}}$ **do**
 1144 13: Try $(\hat{\mathbf{z}}_3, \dots, \hat{\mathbf{z}}_M) \leftarrow \text{forward}(C', \mathbf{z}'_1, \kappa, 1)$
 1145 14: **if** previous step succeeded **then**
 1146 Append $(\hat{\mathbf{z}}_3, \dots, \hat{\mathbf{z}}_M)$ to $\mathbf{Z}_{\text{collected}}$.
 1147 **end if**
 1148 17: **end while**
 1149 18: $N_{\text{infer}} \leftarrow$ numbers of attempts on C'
 1150 19: $G \leftarrow G_{\text{fallback}} + G_{\text{infer}}$
 1151 20: **Assert** $G = |\mathbf{Z}_{\text{collected}}|$
 1152 21: Compute $[\hat{p}_{\text{posterior}}(\mathbf{Z}^{(1)} \dots \hat{p}_{\text{posterior}}(\mathbf{Z}^{(G)})]$ using Eq. (5).
 1153 22: Sample $g \in [1..G]$ with probability proportional to $\hat{p}_{\text{posterior}}(\mathbf{Z}^{(g)})$.
 1154 23: $\mathbf{G} \leftarrow \text{backward}(C', \mathbf{Z}^{(g)})[2]$.
 1155 24: $\text{optimize}(C', \mathbf{G})$

1156 We first prove that hyperedges are locally normalized (*i.e.*, the partition function is 1) when θ is well-specified:

1157 **Lemma 1.** *If θ is well-specified, then for any hyperedge $e_k \in \mathbf{E}$ and any valid assignment \mathbf{x} to its source nodes,*
 1158 *the local partition function $Z_k = 1$.*

1160 *Proof.* e_k is either an LM adaptor or a deterministic algorithm:

- 1162 • If e_k is an LM adaptor, $Z_k = \sum_{\mathbf{y}} \tilde{p}_{\theta}(\mathbf{y} \mid \mathbf{x}; e_k) = \sum_{\mathbf{y} \in \text{valid}(\tau_o)} p_{LM}(\mathbf{y} \mid \mathbf{x}; \theta_k) = 1$.
 1163 • If e_k is a deterministic algorithm, by Eq. (2) $Z_k = \sum_{\mathbf{y}} \tilde{p}(\mathbf{y} \mid \mathbf{x}; e_k) = \tilde{p}(\text{canon}(f(\text{parse}(\mathbf{x}, \tau_i)))) + 0 = 1 + 0 = 1$.
 1166

1167 \square

1168 We then use induction based on the TAC C 's topological structure.

1169 **Lemma 2.** *Let θ be a well-specified parameter vector for TAC $C = (\mathbf{Z}, \mathbf{E})$. The conditional partition function*
 $\mathcal{Z}_{\theta}(\mathbf{z}_1) = 1$.

1170 *Proof.* We use induction on the number of nodes k , following the topological sort $\mathbf{z}_1, \dots, \mathbf{z}_M$. For clarity, here we
 1171 abuse the subscript notation for topological order, and therefore \mathbf{z}_M (instead of \mathbf{z}_2) is the output.

1175 Let C_k be the sub-TAC induced by $\{\mathbf{z}_1, \dots, \mathbf{z}_k\}$. Its partition function is $\mathcal{Z}_k(\mathbf{z}_1) = \sum_{\mathbf{z}_2 \dots \mathbf{z}_k} \prod_{m=2}^k \tilde{p}_{\theta}(\mathbf{z}_m \mid S_m)$, where S_m denotes the source nodes of \mathbf{z}_m under its corresponding hyperedge.

1178 **Base Case.** $k = 1$. C_1 has only \mathbf{z}_1 . $\mathcal{Z}_1(\mathbf{z}_1) = 1$ since the product is empty.

1180 **Inductive Step.** We assume $\mathcal{Z}_{k-1}(\mathbf{z}_1) = 1$. First we rewrite $\mathcal{Z}_k(\mathbf{z}_1)$ by explicitly summing over \mathbf{z}_k . Since
1181 $\mathbf{z}_1, \dots, \mathbf{z}_k$ is a topological order, the source nodes of \mathbf{z}_k : S_k is a subset of $\{\mathbf{z}_1, \dots, \mathbf{z}_{k-1}\}$. We thus rewrite $\mathcal{Z}_k(\mathbf{z}_1)$
1182 as

$$1183 \mathcal{Z}_k(\mathbf{z}_1) = \sum_{\mathbf{z}_2 \dots \mathbf{z}_k} \left(\prod_{m=2}^{k-1} \tilde{p}_{\theta}(\mathbf{z}_m \mid S_m) \right) \cdot \left(\sum_{\mathbf{z}_k} \tilde{p}_{\theta}(\mathbf{z}_k \mid S_k) \right). \quad (6)$$

1187 We discuss the summands by the validity of $\mathbf{z}_2 \dots \mathbf{z}_{k-1}$:

- 1189 • If $\mathbf{z}_2 \dots \mathbf{z}_{k-1}$ is valid: by Lemma 1 the term $\sum_{\mathbf{z}_k} \tilde{p}_{\theta}(\mathbf{z}_k \mid S_k) = 1$. This summand is therefore
1190 $\prod_{m=2}^{k-1} \tilde{p}_{\theta}(\mathbf{z}_m \mid S_m)$.
- 1191 • If $\mathbf{z}_2 \dots \mathbf{z}_{k-1}$ is not valid: by Eqs 1 and 2 this summand is 0.

1193 We can thus rewrite Eq. (6) as

$$1195 \mathcal{Z}_k(\mathbf{z}_1) = \sum_{\mathbf{z}_2, \dots, \mathbf{z}_{k-1} \mid \text{valid assignments}} \prod_{m=2}^{k-1} \tilde{p}_{\theta}(\mathbf{z}_m \mid S_m). \quad (7)$$

1199 Equation (7) can be further rewritten to sum over both valid and invalid $\mathbf{z}_2, \dots, \mathbf{z}_{k-1}$ assignments (since again by
1200 Eqs. (1) and (2), the summand is 0 for invalid assignments):

$$1201 \mathcal{Z}_k(\mathbf{z}_1) = \sum_{\mathbf{z}_2, \dots, \mathbf{z}_{k-1}} \prod_{m=2}^{k-1} \tilde{p}_{\theta}(\mathbf{z}_m \mid S_m) = \mathcal{Z}_{k-1}(\mathbf{z}_1). \quad (8)$$

1205 Since by assumption $\mathcal{Z}_{k-1}(\mathbf{z}_1) = 1$, we thus prove by induction $\mathcal{Z}_M(\mathbf{z}_1) = \mathcal{Z}_{\theta}(\mathbf{z}_1) = 1$. \square

1207 Finally, we show that Lemma 2 implies the equivalence of maximizing the normalized and unnormalized
1208 likelihoods when the true parameters are well-specified.

1209 **Theorem 1.** Let Θ be the entire parameter space and let $\Theta' \subseteq \Theta$ be the subset of well-specified parameters.
1210 Assume θ^* uniquely maximizes the normalized likelihood $p_{\theta}(\mathbf{z}_{2..M} \mid \mathbf{z}_1)$ and resides in Θ' . Then,
1211 $\hat{\theta} = \arg \max_{\theta \in \Theta} \tilde{p}_{\theta}(\mathbf{z}_{2..M} \mid \mathbf{z}_1) \implies \hat{\theta} = \theta^*$.

1213 *Proof.* First we note $\forall \theta \in \Theta, \mathcal{Z}_{\theta}(\mathbf{z}_1) \leq 1$, since for any adaptor $\sum_{\mathbf{y}} \tilde{p}_{\theta}(\mathbf{y} \mid \mathbf{x}) \leq 1$. By Eqs. (1) and (2) the
1214 global partition function must also be ≤ 1 .

1215 We rewrite the unnormalized likelihood as a product of normalized likelihood and the partition function:

$$1217 \tilde{p}_{\theta}(\mathbf{z}_{2..M} \mid \mathbf{z}_1) = p_{\theta}(\mathbf{z}_{2..M} \mid \mathbf{z}_1) \cdot \mathcal{Z}_{\theta}(\mathbf{z}_1) \quad (9)$$

1218 Since $\mathcal{Z}_{\theta}(\mathbf{z}_1) \leq 1, \forall \theta \in \Theta, \tilde{p}_{\theta}(\mathbf{z}_{2..M} \mid \mathbf{z}_1) \leq p_{\theta}(\mathbf{z}_{2..M} \mid \mathbf{z}_1)$.

1220 At the well-specified true parameters θ^* we have $\mathcal{Z}_{\theta}(\mathbf{z}_1) = 1$ by Lemma 2. Therefore $\tilde{p}_{\theta^*}(\mathbf{z}_{2..M} \mid \mathbf{z}_1) =
1221 p_{\theta^*}(\mathbf{z}_{2..M} \mid \mathbf{z}_1)$.

1222 By our assumption that θ^* maximizes normalized likelihood, $\forall \theta \in \Theta, p_{\theta^*}(\mathbf{z}_{2\dots M} \mid \mathbf{z}_1) \geq p_{\theta}(\mathbf{z}_{2\dots M} \mid \mathbf{z}_1)$.

1223 Combining everything together:

$$\begin{aligned} 1225 \quad \tilde{p}_{\theta^*}(\mathbf{z}_{2\dots M} \mid \mathbf{z}_1) &= p_{\theta^*}(\mathbf{z}_{2\dots M} \mid \mathbf{z}_1) \\ 1226 &\geq p_{\theta}(\mathbf{z}_{2\dots M} \mid \mathbf{z}_1) \\ 1227 &\geq \tilde{p}_{\theta}(\mathbf{z}_{2\dots M} \mid \mathbf{z}_1) \end{aligned}$$

1228 for all $\theta \in \Theta$. Under the assumption θ^* is unique, $\theta^* = \arg \max_{\theta \in \Theta} \tilde{p}_{\theta}(\mathbf{z}_{2\dots M} \mid \mathbf{z}_1) = \hat{\theta}$. \square

1229 **Theorem 2.** Let $\theta = \{\theta_1 \dots \theta_K\}$ be the union of a K -adaptor TAC's LM adaptor parameters. If $\forall \mathbf{z}_{k,1} \in$
1230 $\Sigma^*, \mathbf{z}_{k,2} \in \Sigma^*, \|\nabla_{\theta} (\sum \log p_{LM}(\mathbf{z}_{k,2} \mid \mathbf{z}_{k,1}; \theta))\|_{\infty} \leq G$, then $\nabla_{\theta} \log \mathcal{Z}_{\theta} \leq 2G(1 - \mathcal{Z}_{\theta})$.

1231 *Proof.* Here we fix $\mathbf{z}_1 = x$. We denote $\mathbf{z}_{2\dots M} = y$. Let $p_{LM}^{(k)}(y)$ be the k -th LM adaptor's unmasked node
1232 probability, given (x, y) as TAC input and output. We then denote $p_{\theta}(y) = \prod_k p_{LM}^{(k)}$ as a TAC's *normalized*
1233 distribution over node assignments (without masking invalid ones). The partition function $\mathcal{Z}_{\theta} = \sum_y p_{\theta}(y \mid$
1234 $x) \mathbb{I}(y \in V) = \Pr_{p_{\theta}}(V)$ where V is the set of valid node assignments.

1235 We first rewrite $\nabla_{\theta} \log \mathcal{Z}_{\theta}$ as an expectation under p_{θ} :

$$1236 \quad \nabla_{\theta} \log \mathcal{Z}_{\theta} = \mathbb{E}_{y \sim p_{\theta}(\cdot \mid V)} [\nabla_{\theta} \log p_{\theta}(y)]. \quad (10)$$

1237 Using the identity $\sum_y p_{\theta}(y) \nabla_{\theta} \log p_{\theta}(y) = 0$, we rewrite Eq. (10) as

$$1238 \quad \nabla_{\theta} \log \mathcal{Z}_{\theta} = \mathbb{E}_{y \sim p_{\theta}(\cdot \mid V)} [\nabla_{\theta} \log p_{\theta}(y)] - \mathbb{E}_{y \sim p_{\theta}} [\nabla_{\theta} \log p_{\theta}(y)]. \quad (11)$$

1239 Let $f = \nabla_{\theta} \log p_{\theta}(y)$. We can now rewrite $\|\nabla_{\theta} \log \mathcal{Z}_{\theta}\|_{\infty}$ as

$$\begin{aligned} 1240 \quad \|\nabla_{\theta} \log \mathcal{Z}_{\theta}\|_{\infty} &= \|\mathbb{E}_{p_{\cdot \mid V}} [f] - \mathbb{E}_{p_{\theta}} [f]\|_{\infty} \\ 1241 &= \|\sum_y f \cdot (p_{\theta}(y \mid V) - p_{\theta}(y))\|_{\infty} \\ 1242 &\leq \sum_y \|f\|_{\infty} \cdot |p_{\theta}(y \mid V) - p_{\theta}(y)| \\ 1243 &\leq \sum_y G \cdot |p_{\theta}(y \mid V) - p_{\theta}(y)|. \end{aligned} \quad (12)$$

1244 Noting that $\sum_y |p_{\theta}(y \mid V) - p_{\theta}(y)|$ is twice the total variation between p_{θ} and $p_{\theta}(\cdot \mid V)$, and that the total
1245 variation between p_{θ} and $p_{\theta}(\cdot \mid V)$ is $(1 - \mathcal{Z}_{\theta})$ —the sum of invalid assignments' probabilities under p_{θ} —we
1246 can rewrite Eq. (12) as $\|\nabla_{\theta} \log \mathcal{Z}_{\theta}\|_{\infty} \leq 2G(1 - \mathcal{Z}_{\theta})$. \square

E IMPLEMENTATION CONSIDERATIONS

1261 In this section we discuss practical implementation considerations. In particular, we distinguish between *one-time*
1262 and *per-use* efforts.

E.1 ONE-TIME EFFORTS

1263 **Parsing and canonicalization.** There exist multiple libraries that can readily be used to implement `parse` and
1264 `canon` for typed data-holding objects in Python. One example is LangFun which we use extensively in the paper.
1265 Another popular library is Pydantic, which is used in DSPy.

1269 **Type validation logic.** As we briefly discussed in Footnote 5, the `parse` function can be used to implement
 1270 complex business logic. Such logic can usually be implemented cleanly as part of type definition (e.g., as
 1271 `__init__` and `__post_init__` methods in Python).
 1272

1273 **Algorithms.** The core TAC algorithms for execution and training (Algorithms listed in §C) are general and need
 1274 only be implemented once. The main computational bottlenecks in these algorithms are:
 1275

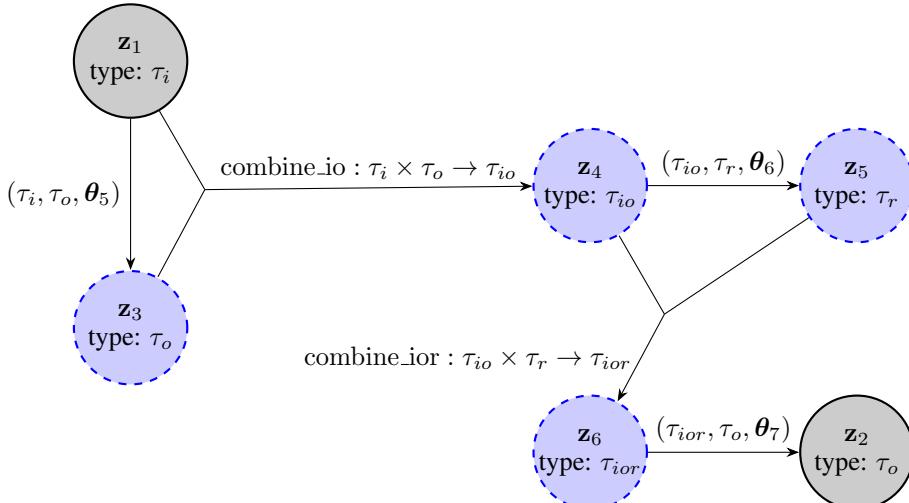
- Sampling from an LM adaptor $p_{LM}(\cdot; \theta)$.
- Evaluating the conditional probability of y given x under an LM adaptor: $p_{LM}(y | x; \theta)$.
- Computing gradients of (x, y) with regard to parameters θ : $\nabla_{\theta} \log p_{LM}(y | x; \theta)$.

1280 A practical implementation can abstract these bottlenecks away, by offloading these intensive parts to dedicated
 1281 inference servers (e.g., vLLM). The core TAC logic remains a lightweight, accelerator-agnostic program.
 1282 Furthermore, since TACs use parameter-efficient fine-tuning (PEFT), the adaptor weights and gradients are small
 1283 enough to be processed quickly, often without needing dedicated accelerators for the logic itself. This design
 1284 significantly reduces the low-level engineering burden.
 1285

1286 E.2 PER-USE EFFORTS

1287 Once the core engine is in place, a practitioner’s effort is focused on defining a TAC hypergraph for their specific
 1288 task. Since the TAC hypergraph is essentially a data flow graph, it can be represented in a way that is directly
 1289 analogous to network architecture definitions in popular neural network frameworks such as PyTorch, where the
 1290 Module’s represent hyperedges, and their `forward` methods connect the typed data nodes.
 1291

1292 F ADDITIONAL TAC DIAGRAMS OF TRAINABLE WORKFLOWS



1311 Figure 6: **refine-structure**: refinement through cascade topology engineering. This cascade models a refinement
 1312 process where an initial output sketch is iteratively refined based on generated rationales.
 1313

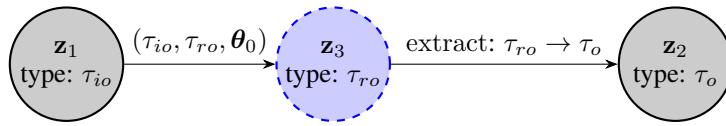


Figure 7: C_{fallback} for **cat-type-structure**. Notice that the adaptor $(\tau_{io}, \tau_{ro}, \theta_0)$ uses ‘fallback’ weights θ_0 that represent no-op weights. Since we conduct experiment on LoRA adaptors in this work, we use the zero-init vectors as θ_0 .

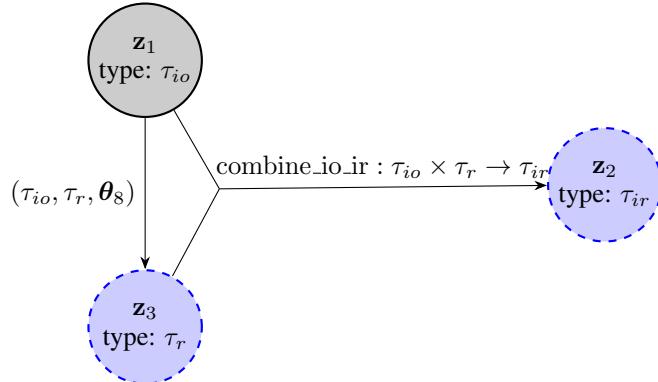


Figure 8: Inference network TAC' for **cat-type-structure**.

G FURTHER DETAILS OF EXPERIMENT SETUP

Data splits. We focus on the low-data regime of task adaptation in this work. For MGSM and MGSM-SymPy, each language has 100/30/120 training/validation/test examples respectively. The splits are 100/30/100 and 100/30/300 for HotPotQA and FinQA respectively. For HotPotQA and FinQA, we use the first entries from the original dataset files as our training and evaluation subsets. For MGSM experiments, we train and evaluate on each language separately. For MuSR tasks, the splits are 100/30/120 and 100/30/126 respectively.

Evaluation. We look at exact match accuracy scores of the answers for all 5 tasks. For MGSM-SymPy experiments, we convert answers from the dataset to integers; as for the model predictions, we evaluate the expressions as rational numbers under SymPy¹⁶, and cast the results as integer numbers. We do not make use of additional clues from the datasets (e.g., the rationales provided for the 8 examples in MGSM datasets).

G.1 TAC SETUP

Training procedure. We train all workflows that have latent variables with our TACSTaR and Amortized TACSTaR algorithms, except for the original (untyped) STaR experiments. Since **direct** experiments do not have latent variables, we train those models using the ordinary cross entropy loss. In all experiments we use a batch size of 8. The Adam optimizer (Kingma & Ba, 2014) is used throughout all experiments, with a learning rate of $5e-5$. We early-stop if no higher validation score is achieved for 4 consecutive epochs. The sampler configuration κ is set to use a combination of top-K and nucleus sampling (Holtzman et al., 2020), where we first choose the top 40 candidates, and cut off accumulated probability mass at 0.95. To train the inference TACs, we accumulate 32 samples from C_{infer} and 16 samples from the fallback model (that is, $G = 48$ at the end of Algorithm 5).

¹⁶<https://www.sympy.org/en/index.html>

1363 **Decoding procedure for generation tasks.** Here we denote the answer type as τ_o . For
 1364 each test input instance, we obtain 32 samples $\hat{\mathbf{Z}}^{(1)} \dots \hat{\mathbf{Z}}^{(32)}$ using `forward`, bucket their out-
 1365 put node values $\text{parse}(\hat{\mathbf{z}}_2^{(1)}, \tau_o) \dots \text{parse}(\hat{\mathbf{z}}_2^{(32)}, \tau_o)$ into B bins, identified by the parsed output
 1366 $y_1 \dots y_B$. We output the answer with maximum accumulated unnormalized probability mass, namely
 1367 $\arg \max_b \sum_{s \in [1..32], \text{parse}(\hat{\mathbf{z}}_2, \tau_o) = y_b} \tilde{p}_{\theta}(\hat{\mathbf{Z}}^{(s)})$.

1368
 1369 **Decoding procedure for classification tasks.** We estimate each label c 's normalized marginal probability using
 1370 Eq. (14), with $N = 32$. We output the label with largest normalized marginal probability as prediction.
 1371

1372 **Object representation of data.** We represent input τ_i and output τ_o as Python types. The objects are encoded as
 1373 string representations under `LangFun`. We design the input and output types separately to reflect the original
 1374 dataset schemata (Listings 1 to 3). As for the rationales (represented by τ_r in **cot-type-structure** and **cot-
 1375 cascade-structure**) we represent them as lists of strings (Listing 4). Product types are represented as new Python
 1376 classes (e.g., the product of type `Question` and `Answer`, represented as τ_{io} in Figs. 7 and 8, is a new class
 1377 `QuestionAnswer`). The object representation can be arbitrarily complex, with `LangFun` handling all `canon`
 1378 and `parse` logic (for example, Listing 6 has `Answer` objects embedded in multiple types; and Listing 7 has
 1379 self-referential definitions).

```
1380 1 class Question:
1381 2     question: str
1382 3
1383 4
1384 5 class Answer:
1385 6     answer: str
```

Listing 1: Input and output type definitions for MGSM

1386
 1387

```
1388 1 class Paragraph:
1389 2     title: str
1390 3     sentences: list[str]
1391 4
1392 5
1393 6 class Context:
1394 7     paragraphs: list[Paragraph]
1395 8
1396 9
1397 10 class Answer:
1398 11     answer: str
1399 12
1400 13
1401 14 class Question:
1402 15     id: str
1403 16     question: str
1404 17     context: Context
```

Listing 2: Input and output type definitions for HotPotQA

1405
 1406

```
1407 1 class Question:
1408 2     question: str
1409 3     pre_text: list[str]
1410 4     table: list[list[str]]
```

```

1410 5     post_text: list[str]
1411 6
1412 7
1413 8     class Step:
1414 9         op: str
1415 10        arg1: str
1416 11        arg2: str
1417 12        res: str
1418 13
1419 14     class Answer:
1420 15         answer: str
1421 16
1422 17
1423 18
1424 19     class QuestionAnswer:
1425 20         question: Question
1426 21         answer: Answer
1427 22
1428 23
1429 24     class Answer:
1430 25         answer: str

```

Listing 3: Input and output type definitions for FinQA

```

1431 1     class Rationale:
1432 2         steps: list[str]

```

Listing 4: Rationale type definition

```

1435 1     class QuestionAnswer:
1436 2         question: Question
1437 3         answer: Answer

```

Listing 5: QuestionAnswer type definition

```

1440
1441 1     class ThinkingSteps:
1442 2         steps: list[str]
1443 3
1444 4
1445 5     class Paragraph:
1446 6         title: str
1447 7         sentences: list[str]
1448 8
1449 9
1450 10    class Context:
1451 11        paragraphs: list[Paragraph]
1452 12
1453 13
1454 14    class SupportingFact:
1455 15        title: str
1456 16        sentence: str
1457 17
1458 18
1459 19    class RelevantContext:

```

```

145720     sentences: list[str]
145821
145922
146023     class Answer:
146124         answer: str
146225
146326     class Question:
146427         id: str
146528         question: str
146629         context: Context
146730
146831     class QuestionAnswer:
146932         question: Question
147033         answer: Answer
147134
147235     class AnswerFirstAttemptThinkingStepsAnswer:
147336         answer_first_attempt: Answer
147437         thinking_steps: ThinkingSteps
147538         answer: Answer
147639
147740     class QuestionAnswerFirstAttempt:
147841         question: Question
147942         answer_first_attempt: Answer
148043
148144     class QuestionAnswerFirstAttemptThinkingSteps:
148245         question: Question
148346         answer_first_attempt: Answer
148447         thinking_steps: ThinkingSteps

```

Listing 6: Type definitions for **refine-structure** on HotPotQA

```

1     class Expression:
2         operator: Literal['+', '-', '*', '/']
3         left: Union[int, 'Expression']
4         right: Union[int, 'Expression']
5
6     class Answer:
7         answer: Expression

```

Listing 7: Expression type definitions in MGSM **expression-cascade-structure** experiments

G.2 DSPY SETUP

We conduct most of the DSPy experiments under v 3.0.1, but report results from DSPy v 2.6.19 for gemini-1.1-7b-it experiments since both BFSWRS and MIPROv2 struggle to generate valid outputs under DSPy v 3.0.1. Moreover, the non-optimized MGSM average accuracy is much lower under v 3.0.1 (for Native CoT it is 0.7% under v 2.6.19, and 0.2% under v 3.0.1). For all other experiments, we report results from DSPy v 3.0.1 which sets up JSON schema-based constrained decoding correctly out-of-the-box. As we noted in §4.2, constrained decoding significantly improves performance for tasks with structured output.

1504 We serve base models on vLLM v 0.10.0.
 1505

1506 **Input and output object definitions.** For structured input and output tasks, we subclass `dspy.Signature`
 1507 as `QASignature` to represent examples. The property names and types in a `QASignature` class are identical
 1508 to counterparts in TAC experiments. FinQA and MGSM-SymPy signatures are listed in Listing 8 and Listing 9
 1509 respectively.

```
1510
1511 1 class QASignature(dspy.Signature):
1512 2     pre_text: list[str] = dspy.InputField()
1513 3     table: list[list[str]] = dspy.InputField()
1514 4     post_text: list[str] = dspy.InputField()
1515 5     question: str = dspy.InputField()
1516 6     answer: str = dspy.OutputField()
```

1517 Listing 8: DSPy object signature for FinQA. Property names and types are identical to their TAC counterparts in
 1518 Listing 3

```
1519 1 class Expression(pydantic.BaseModel):
1520 2     operator: Literal['+', '-', '*', '/']
1521 3     left: Union[int, float, 'Expression']
1522 4     right: Union[int, float, 'Expression']
1523 5
1524 6 class QASignature(dspy.Signature):
1525 7     question: str = dspy.InputField()
1526 8     answer: Expression = dspy.OutputField()
```

1527 Listing 9: DSPy object signature for MGSM-SymPy. Property names and types are identical to their
 1528 TAC counterparts in Listing 7

1529 **DSPy models.** We conduct reasoning experiments on both the native `dspy.ChainOfThought` module, and
 1530 an explicitly two-step composite module that resembles TAC **cot-cascade-structure** patterns. Two-step modules
 1531 for FinQA and MuSR are listed in Listings 10 and 11 as examples.

```
1533 1 class QuestionRationale(dspy.Signature):
1534 2     question: str = dspy.InputField()
1535 3     pre_text: list[str] = dspy.InputField()
1536 4     table: list[list[str]] = dspy.InputField()
1537 5     post_text: list[str] = dspy.InputField()
1538 6     question: str = dspy.InputField()
1539 7     rationale: list[str] = dspy.OutputField()
1540 8
1541 9 class RationaleAnswer(dspy.Signature):
1542 10     rationale: list[str] = dspy.InputField()
1543 11     answer: str = dspy.OutputField()
1544 12
1545 13 class TwoStepPredictor(dspy.Module):
1546 14     def __init__(self):
1547 15         self.question_to_rationale = dspy.Predict(QuestionRationale)
1548 16         self.rationale_to_answer = dspy.Predict(RationaleAnswer)
1549 17
1550 18     def forward(self, pre_text: list[str], table: list[list[str]], post_text:
1551 19             list[str], question: str):
1552 20         r = self.question_to_rationale(question=question, pre_text=pre_text, table=
1553 21             table, post_text=post_text).rationale
```

```

1551 20     return dspy.Prediction(answer=self.rationale_to_answer(rationale=r).answer)
1552

```

Listing 10: DSPy two-step reasoning model definition for FinQA

1554

```

1555 1 class QuestionRationale(dspy.Signature):
1556 2     context: str = dspy.InputField()
1557 3     question: str = dspy.InputField()
1558 4     choices: list[str] = dspy.InputField()
1559 5     rationale: list[str] = dspy.OutputField()
1560 6
1561 7 class RationaleAnswer(dspy.Signature):
1562 8     rationale: list[str] = dspy.InputField()
1563 9     choices: list[str] = dspy.InputField()
1564 10    answer: str = dspy.OutputField()
1565 11
1566 12 class TwoStepPredictor(dspy.Module):
1567 13     def __init__(self):
1568 14         self.question_to_rationale = dspy.Predict(QuestionRationale)
1569 15         self.rationale_to_answer = dspy.Predict(RationaleAnswer)
1570 16
1571 17     def forward(self, context: str, question: str, choices: list[str]):
1572 18         r = self.question_to_rationale(question=question, context=context, choices=
1573 19             choices).rationale
1574 20         return dspy.Prediction(answer=self.rationale_to_answer(rationale=r, choices
1575 21             =choices).answer)

```

Listing 11: DSPy two-step reasoning model definition for MuSR

1573

1574

1575

Prompt optimization under DSPy. We experiment with optimizers `dspy.MIPROv2` and `dspy.BootstrapFewShotWithRandomSearch` (listed as `BFSWRS` below). For MGSM-SymPy and FinQA experiments we do not report `BFSWRS` results, as they consistently need more context length than the model maximum (8192). Moreover, for FinQA experiments we resort to `MIPROv2` 0-shot due to similar context length problems.

We set `max_errors=2` for all optimizers. For `MiPROv2` we set `auto='medium'`. For `MiPROv2` with 0-shot settings we additionally set `max_bootstrapped_demos=0`, `max_labed_demos=0`.

1576

1577

1578

1579

1580

1581

1582

1583

1584

1585

1586

1587

1588

1589

1590

1591

1592

1593

1594

1595

1596

1597

1598 **H PER-LANGUAGE TAC AND ORIGINAL STAR MGSM AND MGSM-SYMPY RESULTS**1599
1600 Per-language TAC and original STaR experimental results on tasks MGSM and MGSM-SymPy are listed in
1601 Tables 5 and 6.

Pattern	Adaptation Method	es	en	de	fr	zh	ru	ja	te	th	Average
direct	TACSTaR	27.5	27.5	25.0	25.0	23.3	25.8	23.3	18.3	26.7	24.7
cot-type-structure	TACSTaR	80.0	84.2	76.7	83.3	80.0	85.0	71.7	79.2	83.3	80.4
cot-cascade-structure	TACSTaR	87.5	87.5	83.3	85.8	80.0	87.5	74.2	73.3	80.8	82.2
refine-structure	TACSTaR	86.7	90.0	76.7	77.5	73.3	78.3	69.2	72.5	83.3	78.6
expression-cascade-structure	TACSTaR	83.3	82.5	83.3	75.8	70.0	79.2	65.8	75.0	75.8	75.9
cot-cascade-structure	un-adapted	42.5	47.5	46.7	42.5	45.0	53.3	31.7	45.0	54.2	45.4
cot-type-structure	un-adapted	77.5	79.2	80.8	76.7	68.3	79.2	68.3	69.2	73.3	74.7
expression-cascade-structure	un-adapted	76.7	71.7	69.2	70.8	68.3	68.3	63.3	70.8	73.3	69.5
cot-cascade-structure	amortized TACSTaR	84.2	91.7	86.7	83.3	82.5	81.7	70.8	77.5	83.3	82.4
N/A	original STaR	74.2	79.2	75.8	75.8	70.0	88.3	74.2	75.8	75.8	76.9

1612 Table 5: gemma-2-27b-it MGSM and MGSM-SymPy per-language accuracies (TAC and original STaR
1613 experiments).

1620 Table 6: gemma-1.1-7b-it MGSM per-language accuracies (TAC and original STaR experiments).

1621 **I PER-TASK TAC MUSR RESULTS**

1622 Per-task TAC experimental results on task MuSR are listed in Tables 7 and 8.

Decoding Method	Murder Mystery	Object Placements	Team Allocation	Average
Generation	61.7	51.6	41.7	51.6
Classification	65.0	50.0	80.0	65.0

1632 Table 7: gemma-2-27b-it MuSR per-task accuracies (TAC experiments).

Decoding Method	Murder Mystery	Object Placements	Team Allocation	Average
Generation	60.0	43.7	82.5	62.1
Classification	59.2	42.9	85.8	62.6

1639 Table 8: gemma-1.1-7b-it MuSR per-task accuracies (TAC experiments).

1641 **J PER-TASK DSPY MUSR RESULTS**

1643 Per-task DSPy experimental results on task MuSR are listed in Tables 9 and 10.

Model	Optimizer	Murder Mystery	Object Placements	Team Allocation	Average
Native CoT	None	20.8	0.0	0.0	6.9
Native CoT	MIPRO 0-shot	40.8	$7.9 \cdot 10^{-1}$	0.0	13.9
Native CoT	MIPRO	51.7	50.8	49.2	50.5
Two-step	None	52.5	14.3	22.5	29.8
Two-step	MIPRO 0-shot	55.0	27.8	19.2	34.0
Two-step	MIPRO	59.2	44.4	50.8	51.5

Table 9: gemma-2-27b-it MuSR per-task accuracies (DSPy experiments).

Model	Optimizer	Murder Mystery	Object Placements	Team Allocation	Average
Native CoT	None	10.0	3.2	3.3	5.5
Native CoT	MIPRO 0-shot	6.7	3.2	2.5	4.1
Native CoT	MIPRO	34.2	25.4	50.0	36.5
Two-step	None	33.3	5.6	16.7	18.5
Two-step	MIPRO 0-shot	35.8	1.6	15.0	17.5
Two-step	MIPRO	44.2	32.5	26.7	34.5

Table 10: gemma-1.1-7b-it MuSR per-task accuracies (DSPy experiments).

Model	Optimizer	Murder Mystery	Object Placements	Team Allocation	Average
Native CoT	None	0.0	0.0	0.0	0.0
Native CoT	MIPRO 0-shot	0.0	0.0	0.0	0.0
Native CoT	MIPRO	55.8	50.8	47.5	51.4
Two-step	None	4.2	$7.9 \cdot 10^{-1}$	0.0	1.7
Two-step	MIPRO 0-shot	3.3	1.6	0.0	1.6
Two-step	MIPRO	65.0	59.5	60.0	61.5

Table 11: Qwen3-8B MuSR per-task accuracies (DSPy experiments).

K PER-LANGUAGE DSPY MGSM AND MGSM-SYMPY RESULTS

Per-language DSPy experimental results on tasks MGSM and MGSM-SymPy are listed in Tables 12 to 14.

Model	Optimizer	es	en	de	fr	zh	ru	ja	te	th	Average
Native CoT	None	55.0	57.5	52.5	51.7	54.2	59.2	45.0	39.2	40.0	50.5
Native CoT	BFSWRS	84.2	89.2	87.5	81.7	75.0	87.5	75.0	77.5	79.2	81.9
Native CoT	MIPROv2	82.5	86.7	81.7	76.7	77.5	84.2	70.0	74.2	75.8	78.8
Two-step	None	1.7	5.8	2.5	1.7	3.3	1.7	1.7	3.3	5.0	3.0
Two-step	MIPROv2	76.7	83.3	76.7	78.3	73.3	79.2	70.0	67.5	71.7	75.2
Two-step	BFSWRS	80.8	84.2	76.7	81.7	70.0	81.7	67.5	64.2	72.5	75.5

Table 12: gemma-2-27b-it MGSM per-language accuracies (DSPy experiments).

Model	Optimizer	es	en	de	fr	zh	ru	ja	te	th	Average
Native CoT	None	$8.0 \cdot 10^{-1}$	$8.0 \cdot 10^{-1}$	$8.0 \cdot 10^{-1}$	0.0	0.0	2.5	1.7	0.0	0.0	$7.0 \cdot 10^{-1}$
Native CoT	BFSWRS	0.0	$8.0 \cdot 10^{-1}$	1.7	5.0	$8.0 \cdot 10^{-1}$	1.7	1.7	2.5	0.0	1.6
Native CoT	MIPROv2	$8.0 \cdot 10^{-1}$	1.7	2.5	2.5	1.7	0.0	1.7	$8.0 \cdot 10^{-1}$	$8.0 \cdot 10^{-1}$	1.4
Two-step	None	0.0	0.0	$8.0 \cdot 10^{-1}$	0.0	0.0	0.0	1.7	0.0	0.0	$3.0 \cdot 10^{-1}$
Two-step	MIPROv2	0.0	0.0	0.0	0.0	0.0	$8.0 \cdot 10^{-1}$	0.0	0.0	$8.0 \cdot 10^{-1}$	$2.0 \cdot 10^{-1}$
Two-step	BFSWRS	0.0	0.0	0.0	0.0	0.0	$8.0 \cdot 10^{-1}$	0.0	0.0	$8.0 \cdot 10^{-1}$	$2.0 \cdot 10^{-1}$

Table 13: gemma-1.1-7b-bit MGSM per-language accuracies (DSPy experiments).

Model	Optimizer	es	en	de	fr	zh	ru	ja	te	th	Average
Native CoT	None	56.7	66.7	55.0	45.8	47.5	59.2	45.0	49.2	45.8	52.3
Native CoT	MIPROv2	66.7	64.2	58.3	60.8	56.7	62.5	50.8	42.5	51.7	57.1
Two-step	None	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Two-step	MIPROv2	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

Table 14: gemma-2-27b-bit MGSM-SymPy per-language accuracies (DSPy experiments).

L DSPY FINQA RESULTS

DSPy experimental results on the FinQA task are listed in Table 15 and Table 16.

Model	Optimizer	Accuracy
Native CoT	None	11.7
Native CoT	MIPROv2 0-shot	12.7
Two-step	None	5.7
Two-step	MIPROv2 0-shot	10.7

Table 15: gemma-2-27b-bit FinQA accuracy (DSPy experiments).

Model	Optimizer	Accuracy
Native CoT	None	0.0
Native CoT	MIPROv2 0-shot	$6.7 \cdot 10^{-1}$
Two-step	None	0.0
Two-step	MIPROv2 0-shot	$3.3 \cdot 10^{-1}$

Table 16: gemma-1.1-7b-bit FinQA accuracy (DSPy experiments).

Model	Optimizer	Accuracy
Native CoT	None	4.3
Native CoT	MIPROv2 0-shot	5.3
Two-step	None	1.0
Two-step	MIPROv2 0-shot	12.0

Table 17: Qwen3-8B FinQA accuracy (DSPy experiments).

1739 **M EXAMPLE EXPRESSIONS FROM `expression-cascade-structure` UNDER THE**
 1740 **MGSM-SYMPY TASK**

1742 See Table 18.

1744 Question	1745 Answer	1746 Expression
1746 Nissa hires 60 seasonal workers to play elves in her department store's Santa village. A third of the elves quit after children vomit on them, then 10 of the remaining elves quit after kids kick their shins. How many elves are left?	20	$(60 - (60/3)) - 10$
1749 The expenditure of Joseph in May was \$500. In June, his expenditure was \$60 less. How much was his total expenditure for those two months?	940	$500 + 440$
1751 Tom gets 4 car washes a month. If each car wash costs \$15 how much does he pay in a year?	720	$(15 \times 4) \times 12$

1753 Table 18: Example arithmetic expressions generated for MGSM questions by `expression-cascade-structure`.

1755 **N EXAMPLE INSTRUCTION PROMPT GENERATED BY LANGFUN**

1758 The LangFun library translates requests that transformed a typed object into another typed object into natural
 1759 language instructions for LLMs, to facilitate its `parse` operations. For example, Listing 12 is a prompt generated
 1760 by LangFun for the request that transforms a `Question` object into an `Answer` object.

```

1762 1 Please respond to the last INPUT_OBJECT with OUTPUT_OBJECT according to
1763 2   OUTPUT_TYPE.
1764 3
1765 4   INPUT_OBJECT:
1766 5     1 + 1 =
1767 6
1768 7   OUTPUT_TYPE:
1769 8     Answer
1770 9
1771 10    ```python
1772 11      class Answer:
1773 12        final_answer: int
1774 13
1775 14    OUTPUT_OBJECT:
1776 15      ```python
1777 16      Answer(
1778 17        final_answer=2
1779 18      )
1780 19
1781 20
1782 21    INPUT_OBJECT:
1783 22      ```python
1784 23      Question(
1785 24        question='How are you?'
1786 25      )
1787 26
1788 27
1789 28  OUTPUT_TYPE:
```

```

178629     Answer
178730
178831     ```python
178932     class Answer:
179033         answer: str
179134     ...
179235
179336     OUTPUT_OBJECT:

```

Listing 12: Example instruction prompt generated by LangFun

O PERFORMANCE-COMPUTATIONAL COST ANALYSIS

We conduct additional experiments on FinQA (Gemma 2 27B) to analyze the cost-benefit trade-off, using token counts (*i.e.*, tokens used during training, and tokens used for evaluation on the entire test set) as a proxy for cost. On this task, the TACSTaR run processed $\sim 27M$ training tokens. The baseline DSPy optimization (MIPROv2, using the default `auto='medium'` configuration) processed $\sim 1.6M$ tokens. Decoding the 300 test examples takes $\sim 29M$ tokens for TAC (using 32 samples for robust estimation as described in §G.1) and $\sim 0.5M$ tokens for DSPy.

We evaluated whether increasing DSPy’s compute budget closes the performance gap:

- **Scaling Training:** We increased the `num_trials` hyperparameter under MIPRO v2 from 12 (under the `auto='medium'` default setting) to 32 ($\sim 3.7M$ tokens) and 300 ($\sim 33M$ tokens).
- **Scaling Inference:** We used majority voting (with the same `T=1.0`, `top_p=0.95`, `top_k=40` settings) with ensemble sizes 100 ($\sim 55M$ tokens) and 500 ($\sim 260M$ tokens).

Results listed in Table 19 show that DSPy performance plateaus quickly. Even when significantly increasing DSPy’s training and test budget (up to 9x TAC’s inference cost), the accuracy (14.3) remains far below TACs (34.0).

Training budget	Test budget	Training token count			Test token count			Accuracy
		Encoded	Decoded	Total	Encoded	Decoded	Total	
12	1	1577489	107297	1684786	489646	35302	524948	12.7
12	100	1362700	98030	1460730	50714928	4695123	55410051	12.7
12	500	1408479	95815	1504294	243050795	18271389	261322184	12.3
32	100	3246101	226967	3473068	50074600	4747957	54822557	12.0
32	500	3429654	240376	3670030	247409788	20255592	267665380	14.3
300	100	31112328	2146258	33258586	48754600	3655251	52409851	12.3
300	500	31216543	2164159	33380702	246926388	23154727	270081115	12.7
TAC		26007546	1349250	27356796	27517698	1476911	28994609	34.0

Table 19: Cost analysis. Training budget corresponds to the `num_trials` hyperparameter under MIPRO v2, and test budget corresponds to ensemble size. Accuracy numbers from the first and last rows are copied from Table 15 and Fig. 2a respectively.

1833 **P ASSESSING SNIS QUALITY UNDER LEARNED PROPOSAL DISTRIBUTIONS**

1834

1835 To quantify the effectiveness of self-normalized importance sampling (SNIS) under the trained inference networks
1836 (Amortized TACSTaR), we conducted additional experiments with the `cot-cascade-structure` pattern to
1837 estimate reverse KL divergence, $\text{KL}[\tilde{q}_\phi || p_\theta]$, which measures how well self-normalized importance samples
1838 from adapted proposal $q_\phi(\cdot | z_1^*, z_2^*)$ —as the mixture \tilde{q}_ϕ —approximates the true posterior $p_\theta(\cdot | z_1^*, z_2^*)$:
1839 $\text{KL}[\tilde{q}_\phi || p_\theta] = 0$ when $\tilde{q}_\phi = p_\theta$. Since the partition function of $\sum_z p_\theta(z_3 = z, z_2 = z_2^* | z_1^*)$ is intractable,
1840 we estimate the difference between KL divergences, comparing against SNIS distributions under the unamortized
1841 TACSTaR fallback.

1842 Specifically, we rewrite $\text{KL}[\tilde{q}_\phi || p_\theta] - \text{KL}[\tilde{p}_{\text{fallback}} || p_\theta]$ as:

1843
$$\begin{aligned}
1844 \text{KL}[\tilde{q}_\phi || p_\theta] - \text{KL}[\tilde{p}_{\text{fallback}} || p_\theta] &= \mathbb{E}_{z \sim \tilde{q}_\phi(\cdot | z_1^*, z_2^*)} [\log \tilde{q}_\phi(z | z_1^*, z_2^*) - \log p_\theta(z | z_1^*, z_2^*)] \\
1845 &\quad - \mathbb{E}_{z \sim \tilde{p}_{\text{fallback}}(\cdot | z_1^*, z_2^*)} [\log \tilde{p}_{\text{fallback}}(z | z_1^*, z_2^*) - \log p_\theta(z | z_1^*, z_2^*)] \\
1846 &= \underbrace{\mathbb{E}_{z \sim \tilde{q}_\phi(\cdot | z_1^*, z_2^*)} [\log \tilde{q}_\phi(z | z_1^*, z_2^*) - \log p_\theta(z, z_2 = z_2^* | z_1^*)]}_{\text{call this } K_{\tilde{q}_\phi}} \\
1847 &\quad - \underbrace{\mathbb{E}_{z \sim \tilde{p}_{\text{fallback}}(\cdot | z_1^*, z_2^*)} [\log \tilde{p}_{\text{fallback}}(z | z_1^*, z_2^*) - \log p_\theta(z, z_2 = z_2^* | z_1^*)]}_{\text{call this } K_{\tilde{p}_{\text{fallback}}}}
\end{aligned} \tag{13}$$

1848 where we exploit the identity

1849
$$\begin{aligned}
1850 \mathbb{E}_{q_1} [\log p_\theta(z | z_1^*, z_2^*)] - \mathbb{E}_{q_2} [\log p_\theta(z | z_1^*, z_2^*)] \\
1851 &= \mathbb{E}_{q_1} [\log p_\theta(z, z_2 = z_2^* | z_1^*) - \log \underbrace{\left(\sum_{z'} p_\theta(z^*, z_2 = z_2^* | z_1^*) \right)}_{=C}] \\
1852 &\quad - \mathbb{E}_{q_2} [\log p_\theta(z | z_1^*, z_2^*) - \log \left(\sum_{z'} p_\theta(z^*, z_2 = z_2^* | z_1^*) \right)] \\
1853 &= \mathbb{E}_{q_1} [\log p_\theta(z, z_2 = z_2^* | z_1^*)] - \mathbb{E}_{q_2} [\log p_\theta(z, z_2 = z_2^* | z_1^*)] - \mathbb{E}_{q_1}[C] + \mathbb{E}_{q_2}[C] \\
1854 &= \mathbb{E}_{q_1} [\log p_\theta(z, z_2 = z_2^* | z_1^*)] - \mathbb{E}_{q_2} [\log p_\theta(z, z_2 = z_2^* | z_1^*)].
\end{aligned}$$

1855

1856 To compute $\text{KL}[\tilde{q}_\phi || p_\theta] - \text{KL}[\tilde{p}_{\text{fallback}} || p_\theta]$, we compute the two expectations $K_{\tilde{q}_\phi}$ and $K_{\tilde{p}_{\text{fallback}}}$ in Eq. (13),
1857 constructing the self-normalized distribution $\tilde{p}_{\text{fallback}}(z | z_1^*, z_2^*) \propto \frac{p_\theta(z, z_2 = z_2^* | z_1^*)}{p_{\text{fallback}}(z | z_1^*, z_2^*)}$ with 16 samples drawn
1858 from p_{fallback} . We also construct $\tilde{q}_\phi(z | z_1^*, z_2^*)$ this way with 16 samples. We report the average difference over
1859 validation datasets on the latest MGSM checkpoints in Table 20. On average, SNIS distribution under the adapted
1860 \tilde{q}_ϕ is closer to the true posterior distribution $p_\theta(\cdot | z_1^*, z_2^*)$ than the counterpart under the default p_{fallback} .

1861 **Q ASSESSING SAMPLE DIVERSITY**

1862

1863 We conducted additional experiments on MGSM (Gemma 7B) to analyze whether diversity collapses as type
1864 compliance increases during training (§4.4). We use Effective Sample Size (ESS) normalized between 0 and 1 as
1865 the evaluation metric. We measure ESS (using 16 samples for each validation example) at the end of each training
1866 epoch on the validation datasets. Results are listed in Fig. 9. We find the average ESS fluctuates between 0.4 to
1867 0.5, implying that the samples did not collapse to the mode (in which case normalized ESS would be around
1868 $1/16 = 0.0625$).

Language	$K_{\tilde{q}_\phi}$	$K_{\tilde{p}_{\text{fallback}}}$	$\text{KL}[\tilde{q}_\phi p_\theta] - \text{KL}[\tilde{p}_{\text{fallback}} p_\theta]$
es	28.63845406	29.34799745	-0.7095433899
en	44.39640467	48.11942965	-3.723024984
de	48.86386756	49.65604005	-0.7921724921
fr	39.33989167	40.95757624	-1.617684578
zh	38.38841294	35.32150839	3.066904549
ja	40.36140304	40.1701377	0.1912653429
ru	30.32379052	31.74457529	-1.420784771
te	49.4455041	52.05616812	-2.610664018
th	39.60001892	39.65173477	-0.05171585152
Average	39.92863861	40.78057419	-0.8519355771

Table 20: Average KL divergence differences from the true posterior, between SNIS distributions under adapted and un-adapted fallback distributions.

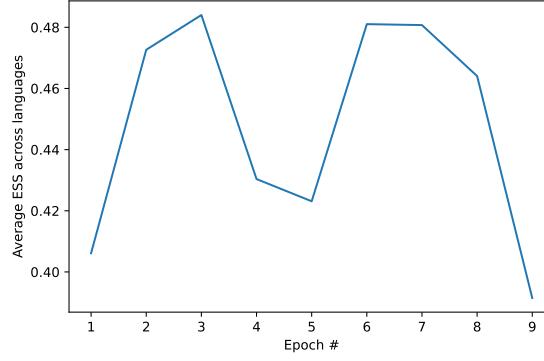


Figure 9: Average ESS at end of epochs. Sample size = 16. ESS value normalized $\in [\frac{1}{16}, 1]$ in plot.

R ESTIMATING MARGINAL PROBABILITIES VIA IMPORTANCE SAMPLING

A key advantage of the probabilistic framing of TACs is the ability to estimate marginal probabilities, even when the distribution is unnormalized. LM adaptors in a TAC can be used as proposal distributions for importance sampling.

Let \mathbf{z}_m be a node coming out of an LM adaptor. An N -sample estimate of the unnormalized probability that \mathbf{z}_m equals c , conditioned on the input \mathbf{z}_1 , $\tilde{p}(\mathbf{z}_m = c | \mathbf{z}_1; \theta)$, is:

$$\hat{p}_{|\mathbf{z}_1}(m, c, N) = \sum_{n=1}^N \left[\frac{p_{LM}(\mathbf{z}_m = c; \theta)}{N \cdot p_{LM}(\mathbf{z}_m = \mathbf{z}_m^{(n)}; \theta)} \right] \quad (14)$$

where $\mathbf{z}_m^{(n)}$ is the n -th sample of \mathbf{z}_m (drawn using Algorithm 1). Equation (14) is an unbiased importance sampling estimate of the unnormalized probability $\tilde{p}(\mathbf{z}_m = c | \mathbf{z}_1; \theta)$ (since $\text{supp}(\tilde{p}) \subseteq \text{supp}(p_{LM})$).

1927 In the special case that \mathbf{z}_m has finite support (e.g., a classification task), we can estimate the *normalized* marginal
 1928 probability using Self-Normalized Importance Sampling (SNIS) by renormalizing over all possible values c' :
 1929

$$1930 \quad 1931 \quad 1932 \quad 1933 \quad 1934 \quad 1935 \quad 1936 \quad 1937 \quad 1938 \quad 1939 \quad 1940 \quad 1941 \quad 1942 \quad 1943 \quad 1944 \quad 1945 \quad 1946 \quad 1947 \quad 1948 \quad 1949 \quad 1950 \quad 1951 \quad 1952 \quad 1953 \quad 1954 \quad 1955 \quad 1956 \quad 1957 \quad 1958 \quad 1959 \quad 1960 \quad 1961 \quad 1962 \quad 1963 \quad 1964 \quad 1965 \quad 1966 \quad 1967 \quad 1968 \quad 1969 \quad 1970 \quad 1971 \quad 1972 \quad 1973$$

$$\hat{p}(\mathbf{z}_m = c \mid \mathbf{z}_1; \boldsymbol{\theta}) = \frac{\hat{p}_{|\mathbf{z}_1}(m, c, N)}{\sum_{c'} \hat{p}_{|\mathbf{z}_1}(m, c', N)}. \quad (15)$$