AutoMix: Automatically Mixing Language Models

Anonymous ACL submission

Abstract

Large language models (LLMs) are now available from cloud API providers in various sizes and configurations. While this diversity offers a broad spectrum of choices, effectively leveraging the options to optimize computational cost and performance remains challenging. In this work, we present AutoMix, an approach that strategically routes queries to larger LMs, based on the approximate correctness of outputs from a smaller LM. Central to AutoMix is a few-shot self-verification mechanism, which estimates the reliability of its own outputs without requiring training. Given that verifications can be noisy, we employ a meta-verifier in AutoMix to refine the accuracy of these assessments. Our experiments using LLAMA2-13/GPT-4, on five context-grounded reasoning datasets demonstrate that AutoMix surpasses established baselines, improving the incremental benefit per cost by up to 86%.[1]

1 Introduction

Human problem-solving inherently follows a multi-step process: generate a solution, verify its validity, and refine it further based on verification outcomes. The emulation of this self-refinement and reflective behavior has gained attention in the recent research (Pan et al., 2023a; Madaan et al., 2023; Reid and Neubig, 2022; Schick et al., 2022; Welleck et al., 2022; Shinn et al., 2023). Current self-refine paradigms use a single model across all problem-solving stages, demonstrating effectiveness in specific scenarios (Madaan et al., 2023; Shinn et al., 2023). Yet, the intrinsic complexity and variability of tasks, from simplistic (e.g., binary classification on separable data) to complex (e.g., code generation) and potentially unsolvable (e.g., certain forms of multi-step reasoning), motivate an alternative approach of *model switching*. Model switching sequentially queries models of disparate sizes and capabilities, at each step determining whether to accept the current output or route to a more capable, but computationally intensive, model (Liu et al., 2020; Zhou et al., 2020; Madaan and Yang, 2022; Geng et al., 2021; Schuster et al., 2022).

Contemporary model-switching strategies often rely on separate routing models trained for a fixed set of tasks or require access to logits (Chen et al., 2023; Welleck et al., 2022; Reid and Neubig, 2022). However, modern LLMs are often accessible solely through black-box APIs, limiting direct model optimization. This constraint, along with the expectation that LLMs handle a broad range of tasks, creates a challenge that existing routing approaches fail to address. In response, we introduce `AutoMix`, a method that allows users to *mix* models of various sizes and capabilities and only assumes access to black-box LLM APIs. `AutoMix` consists of three steps designed within the constraints of black-box access: solution generation (using the smaller model to generate an initial answer), self-verification (using the same smaller model to assess output), and selective routing (employing larger models only when suggested by self-verification).

We formulate self-verification as an entailment problem, evaluating the consistency of generated answers with the provided context (Poliak, 2020; Dagan et al., 2022). For example, an answer discussing "desert animals" in a context focused on "aquatic life" would be flagged as inconsistent. However, recognizing that self-verification can sometimes be inconsistent or noisy (Tyen et al., 2023; Huang et al., 2023), we introduce a *meta-verifier* to evaluate the reliability of the initial verification. The meta-verifier acts as a secondary check, providing an additional layer of confidence assessment to ensure that the decision to route a task to a larger or smaller model is justifiable. Additionally, in contrast to existing model-switching approaches, which generally classify tasks as *Simple*

[1] We will release the code and data upon acceptance.

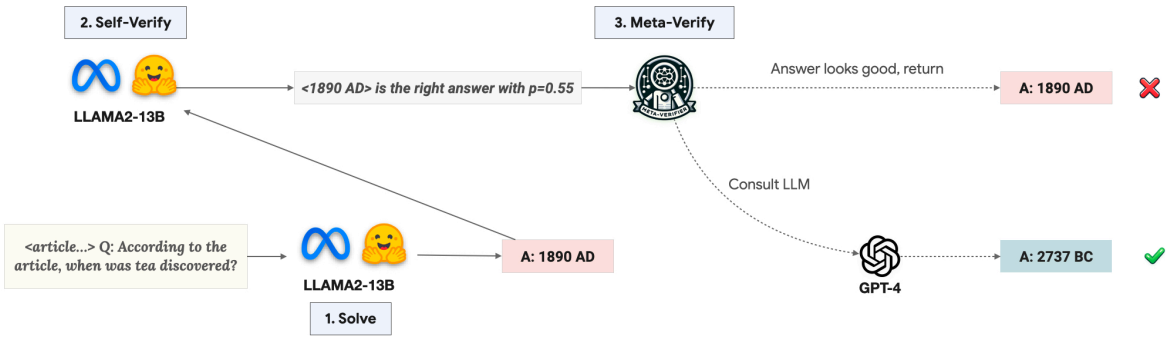

Figure 1: AutoMix: Given a context (like an article) and a question q, an initial answer (*1890 AD*) is generated with the smaller language model (SLM). The answer is self-verified by the SLM, yielding a noisy verification score. The Meta-Verifier subsequently assesses the verifier's results. Based on the meta-verifier's decision, either the initial answer (*1890 AD*) is returned, or the question is rerouted to a larger language model (LLM) to enhance accuracy.

or *Complex* for model routing (Chen et al., 2023), AutoMix includes an "unsolvable" task classification. This conserves resources by avoiding routing overly complex queries to larger models.

In summary, our contributions are: (1) We introduce AutoMix, a strategy that mixes language models (LLAMA2-13B and GPT-4) without access to internal model details. (2) We explore context-grounded entailment as a self-verification method and propose a POMDP-based meta-verifier to enhance decision reliability. (3) We introduce the *Incremental Benefit Per Unit Cost* (IBC) metric to quantify the efficacy of combined model usage. (4) Across five datasets, we demonstrate that AutoMix provides up to 86% efficiency improvement, outperforming strong baselines.

2 AutoMix

Task and setup We tackle the problem of context-grounded question answering, where given a context $\mathcal{C}$ (e.g., stories, newswire, or research article) and a question q, the model is tasked with generating an accurate and coherent answer, consistent with the provided context. Our choice of tasks is motivated by two key concerns: (1) longer queries are more computationally demanding, underscoring the need for an approach like AutoMix to navigate the cost-accuracy trade-off, and (2) the context allows for cross-checking preliminary answers with available information using self-verification (described shortly). While self-verification in reasoning tasks is challenging for LLMs (Pan et al., 2023a; Huang et al., 2023), we find that context significantly aids this process.

We deploy two distinct models: a smaller, cost-efficient model, denoted as SLM (smaller language model), and a larger, more accurate yet costly model, LLM (large language model. Our objective is to optimize performance while staying economical. An initial answer, $\mathcal{A}_s$, is generated using the smaller SLM. Further, in Appendix B.2 we extend AutoMix to 3 models by incorporating the medium language model, showing significant gains.

```
Context: {context}

Question: {question}

AI Generated Answer: {generated_answer}

Instruction: Your task is to evaluate
↪ if the AI Generated Answer is
↪ correct, based on the provided
↪ context and question. Provide the
↪ judgement and reasoning for each
↪ case. Choose between Correct or
↪ Incorrect.

Evaluation:"
```

Listing 1: **Verification Prompt.** The verification process is framed as a natural language entailment task, where the model determines the validity of the model-generated answer with respect to the context and question. We use a generic few-shot prompt for all tasks (prompt in appendix F.1).

Few-shot Verification To assess the trustworthiness of $\mathcal{A}_s$, we employ a few-shot verifier, $\mathcal{V}$, which ascertains the validity of SLM's outputs and decides if a query should be redirected to LLM. Different from existing works that perform verification by creating a new question (Weng et al., 2022; Jiang et al., 2023), we frame verification

as an entailment task (Dagan et al., 2005; Poliak, 2020; Dagan et al., 2022), aiming to determine if the answer generated by SLM aligns with the provided context. Specifically, the verifier gauges $v = \text{p}(\text{correct} = 1 \mid \mathcal{A}_s, \mathcal{C}, q)$, with correct $= 1$ indicating that $\mathcal{A}_s$ is correct. The verification prompt is outlined in Figure 1. We use the same verification prompt for all tasks. Figure 2 shows an example.

2.1 Meta-verifier

Given the potential inconsistency or noise in verifier outcomes, a secondary evaluation mechanism, which we term the *meta-verifier*, is crucial to vet the verifier's conclusions. In particular, the verifier is tasked with determining whether the SLM's answer is entailed by the context, and this decision is made without considering the inherent difficulty of the problem. Notably, routing *Unsolvable* queries to the LLM is resource-inefficient and does not enhance performance. While ascertaining the ground truth of query difficulty is non-trivial, verification probability and trends from historical data inferred using the validation set, can provide insightful guidance. Formally, we define the meta-verifier's outputs as $m(v, \mathcal{A}_s, \mathcal{C}, q) \rightarrow \{0, 1\}$, where $m = 1$ implies the verifier's output can be trusted.

Addressing the notable challenges of self-correction in large language models (Madaan et al., 2023; Huang et al., 2023), our method employs a non-LLM setup for meta-verification to avoid escalating issues like hallucination and reasoning errors (Dziri et al., 2023). The meta-verifier can adopt various learning strategies, including supervised learning, reinforcement learning, and symbolic reasoning, explored further in upcoming sections. Subsequent sections provide a deeper exploration into two particular implementations of this strategy.

Thresholding In this simplistic meta-verifier approach, the decision is made based on the probability of the verifier being correct with a threshold t, defined as $m_t(v) = 0$ for $v < t$ and $m_t(v) = 1$ for $v \geq t$. Intuitively, a high probability of verifier indicates that verifier is confident in it's decision and can be trusted. For black-box language models, the probability of correctness can be derived by sampling $k > 1$ samples at a higher sampling temperature.

Using a POMDP In the context of a meta-verifier, we observe that the queries could be categorized into three different categories: *Simple*, *Complex*, and *Unsolvable*. The simple queries are addressable by SLM itself; the complex queries are addressable by LLM but not by SLM, and Unsolvable queries are so complex that they cannot be solved by LLM or SLM. Since the ground truth state, i.e., the query category is unknown and unobserved, we formulate this decision problem as a Partially Observable Markov Decision Process (POMDP) (Monahan, 1982). POMDP presents a robust framework, offering a structured way to manage and navigate through the decision spaces where the system's state is not fully observable. A POMDP is defined by a tuple (S, A, T, R, Ω, O), where S is a set of states, A is a set of actions, T represents the state transition probabilities, R is the reward function, Ω is a set of observations, and O is the observation function.

In our scenario, the states S correspond to the three question categories: *Simple*, *Complex*, and *Unsolvable*. Actions are denoted as either reporting the SLM answer or routing to the LLM. Observations, in the form of verifier output v, enable the POMDP to ascertain its belief state, which is

Context

Which country grows the most tea? The answer is India. It grows three times as much as China. Which country drinks the most tea? It's neither China nor Japan. It's Great Britain. [...] In general, there are two kinds of tea. Black tea and green tea. [...] One customer put the bag into a pot. Then he just poured hot water over it. And the tea bag was born. **Shen Nong was the first to drink tea. (Shen was a Chinese emperor.) This was about 2737 B.C.** Shen had bad digestion. So he drank several cups of hot water daily[...] Tea became the drink of China.

Question

When did Shen Nong drink tea?

Generated Answer (by LLAMA2-13B)

He took it in 1990.

Verifier Output (by LLAMA2-13B)

The context does not mention that he took tea in 1990. The AI-generated answer is **Incorrect.**

Figure 2: **Context-Grounded Self-Verification using LLAMA2-13B in Action.** The example showcases the verifier, utilizing the *same model* as the answer generator, identifying and rejecting an inaccurate answer—*He took it in 1990*—by effectively leveraging the context.

procedure ANSWERQUERY($\mathcal{C}, q$)
▷ $\mathcal{C}$: Context, q: Question, SLM/LLM: Small/large language model
$\mathcal{A}_s \leftarrow$ SOLVE(SLM, $\mathcal{C}, q$)
$v \leftarrow$ SELF-VERIFY($\mathcal{A}_s, \mathcal{C}, q$)
if META-VERIFY($v, \mathcal{A}_s, \mathcal{C}, q$) **then**
return $\mathcal{A}_s$
else
$\mathcal{A}_l \leftarrow$ SOLVE(LLM, $\mathcal{C}, q$)
return $\mathcal{A}_l$
end if
end procedure

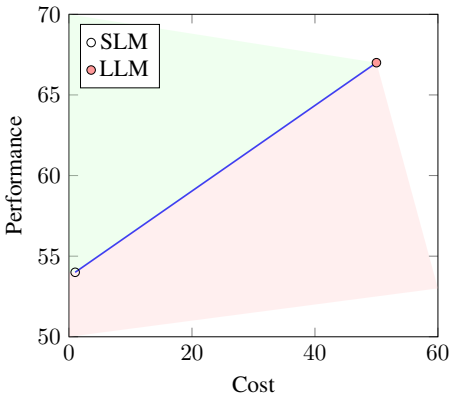

Figure 3: **Left:** `AutoMix` algorithm. **Right**: Performance vs. Cost curve. The slope between SLM and LLM provides a way to the Incremental Benefit per Cost (IBC) for methods that mix models. Methods with a steeper slope than this reference when plotted against SLM have a positive IBC (green region), whereas those below the reference have a negative IBC (red region), falling into the red region.

a probability distribution over S. For instance, a high verifier confidence in the correctness of $\mathcal{A}_s$ would increase the belief in the *Simple* state. The solution to the POMDP subsequently yields a policy that maps belief states to actions, effectively deciding whether to invoke the LLM based on a balance of expected future rewards and computational costs. See Appendix B.1 for more details. Although the POMDP framework inherently handles sequences of decisions, we confine our approach to a single-decision scenario (horizon or episode length 1) for simplicity, with the potential for extension to streaming settings for optimizing across multiple queries or a fixed time duration.

3 Cost-Performance Efficiency Analysis

In our approach to leveraging model performance, it is essential to consider not only the raw accuracy of predictions but also the associated computational or monetary costs. To that end, we introduce a metric to understand the efficiency of the models in terms of cost. We use C_M and P_M to denote the cost and performance of a method M. We also use C_{SLM} and C_{LLM}, and P_{SLM} and P_{LLM}, to denote the cost and performance of using the SLM and LLM, respectively.

Incremental Benefit Per Cost (IBC) We introduce methods, denoted by M, to optimally integrate SLM and LLM. For each method M, we associate a cost C_M and performance P_M. To quantify the utility of M over SLM, we define the metric *Incremental Benefit Per Cost* (IBC) as IBC_M (Equation (3)).

$$\text{IBC}_M = \frac{P_M - P_{\text{SLM}}}{C_M - C_{\text{SLM}}}, \tag{1}$$

$$\text{IBC}_{\text{BASE}} = \frac{P_{\text{LLM}} - P_{\text{SLM}}}{C_{\text{LLM}} - C_{\text{SLM}}}, \tag{2}$$

$$\Delta_{\text{IBC}}(M) = \frac{\text{IBC}_M - \text{IBC}_{\text{BASE}}}{\text{IBC}_{\text{BASE}}} \times 100 \tag{3}$$

The IBC metric captures the efficiency of performance enhancement relative to the additional cost. For comparative evaluation, we set a baseline IBC, IBC_{BASE}, representing the benefit of *always* using LLM over SLM. Finally, we compare methods using Δ_{IBC}, which compares the IBC of a specific method with IBC_{BASE}. A positive IBC lift suggests that M achieves performance increments more cost-effectively than a standalone LLM, whereas a negative lift indicates reduced efficiency (Figure 3) Please see Appendix B.2 for a discussion on extending IBC to multiple models.

Geometric Interpretation On a Performance vs. Cost plot, consider the line segment joining the data points of the small language model (SLM) and the large language model (LLM). This segment's slope represents a basic rate of performance increase for each additional unit of cost. The Incremental Benefit per Cost (IBC) for any method M is the slope of the line from the SLM point to the point representing M(Figure 3). A method M that lies above the SLM-LLM segment provides a steeper slope, indicating a favorable IBC (and a positive Δ_{IBC}). Conversely, if M lies below the segment, it suggests an unfavorable or negative IBC. Our primary objective is to identify or

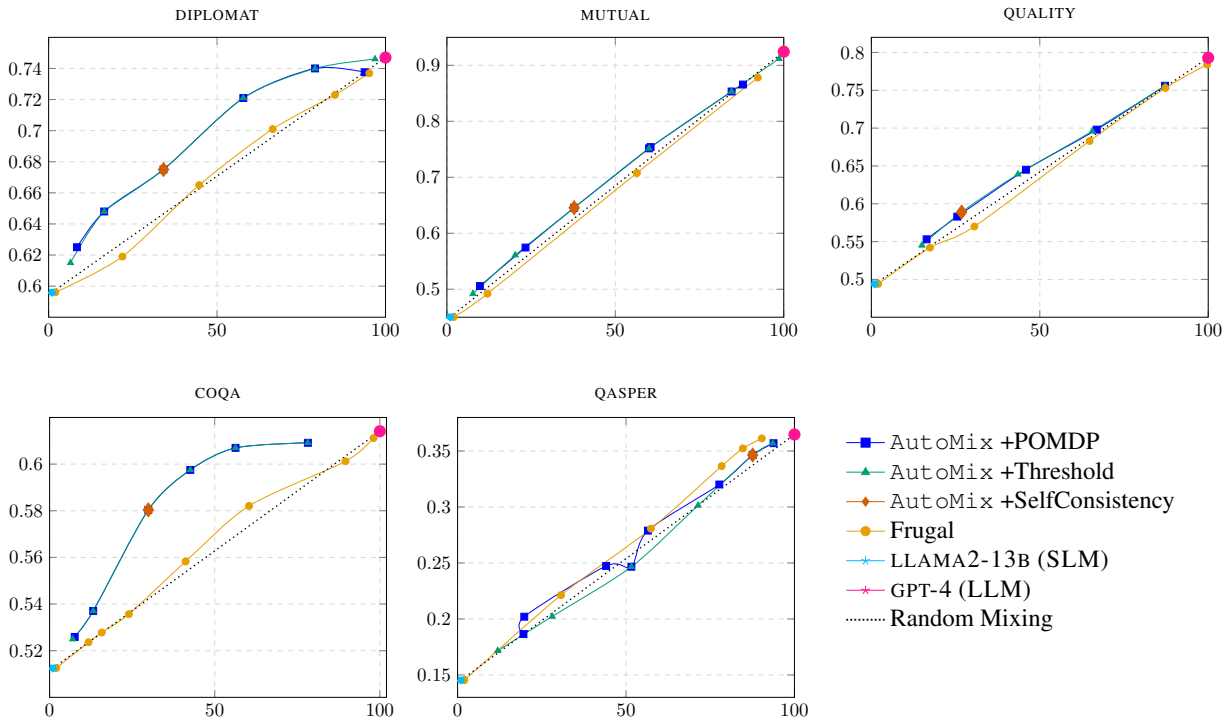

Figure 4: **Main Results:** performance **(y-axis)** vs. cost (**x-axis**) for different methods on the small and large LLAMA2-13/GPT-4. POMDP based meta-verifier is consistenly above the linear interpolation of SLM-LLM, signifying a higher incremental benefit per unit cost (IBC).

develop methods that yield a consistently positive IBC, maximizing performance enhancements for each additional unit of cost.

Cost Calculation To evaluate the efficiency of a method M that leverages both the Small Language Model (SLM) and the Large Language Model (LLM), we define a cost metric, C_M. This metric incorporates the costs of both initial answer generation and verification by the SLM, as well as potential routing to the LLM. Specifically, the total cost C_M is computed as $C_M = 2 \times C_{\text{SLM}} + w_{\text{LLM}} \times C_{\text{LLM}}$. Here, C_{SLM} and C_{LLM} represent the costs of a single query to the SLM and LLM, respectively. The factor $w_{\text{LLM}} \in [0, 1]$ denotes the proportion of times the LLM is used, with $w_{\text{LLM}} = 1$ indicating exclusive use and $w_{\text{LLM}} = 0$ denoting no usage. It's important to note that while our framework uses the SLM for verification, alternative verifiers could be incorporated, which would adjust the cost formula accordingly.

While various complexities determine the pricing of these APIs (Dehghani et al., 2021), given our emphasis on black-box utilization of large language models, we choose to represent cost simply: the monetary expense charged to the end user by the language model APIs.

4 Experiments

Setup We use LLAMA2-13B (Touvron et al., 2023) as our smaller language model (SLM) and GPT-4 (OpenAI, 2023) as the larger language model (LLM), assigning a relative cost of 1 unit for the SLM and 100 units for the LLM. This reflects the actual price disparity between these models [2]. Furthermore, in practical setups, SLM might be deployed with on-premise hardware, and LLM might be only available through relatively expensive APIs, further skewing the cost ratio. The cost ratio between models can shift significantly based on specific deployment scenarios. For instance, for a user with access to a single A6000 GPU, running LLAMA2-13B might incur virtually no cost, while utilizing GPT-4 could prove prohibitively expensive. Please see Appendix D for more details on the experimental setup.

Datasets We experiment with a diverse set of datasets: i) QASPER (Dasigi et al., 2021): Question answering over research papers; ii) QUAL-

[2] https://openai.com/pricing, https://together.ai/

ITY (Pang et al., 2022): Multiple-choice questions (MCQ) on long articles and stories; iii) COQA (Reddy et al., 2019): Conversational comprehension requiring coreference and pragmatic reasoning; iv) MUTUAL (Cui et al., 2020): Multi-turn dialogue reasoning (next response prediction); v) DIPLOMAT (Li et al., 2023): Pragmatic identification and reasoning questions on multi-turn dialogues. We use the F1 score for QASPER and COQA, and accuracy for the remaining datasets. To manage input complexity, we retain a context subset (max 3500 tokens) retrieved using the question as a key. Retrieval is performed with `all-MiniLM-L6-v2` sentence embedding model (Reimers and Gurevych, 2019). We also experiment with additional datasets, CNLI and NARRATIVE-QA from Scrolls, and observe similar trends. Details are in the Appendix E.

We utilize the validation sets from Shaham et al. (2022) for QASPER, and QUALITY, and use the prompts from Shaham et al. (2023). For COQA, MUTUAL, and DIPLOMAT, we employ its validation split and adapt the QUALITY prompt. Regardless of the dataset, we provide identical input prompts to both SLM and LLM to ensure consistent input processing costs. The output length is fixed in multi-choice datasets like QUALITY, and the brevity of responses in other datasets allows us to assume uniform output processing costs. We use greedy decoding (temperature 0) and draw a single sample for both the SLM and LLM.

Baselines We compare against FrugalGPT (F) (Chen et al., 2023) as our baseline. FrugalGPT uses a finetuned DistillBert model (Sanh et al., 2019) as a verifier. If the verifier's confidence probability for a given question, context, and SLM answer falls below a set threshold, the query is routed to the LLM. Due to its significantly lower operational cost, we assign a cost of 0 to the verifier.

Proposed approaches We experiment with three different types of meta-verifiers: i.) **`AutoMix` + Self-Consistency**: Uses the majority decision from verifier from 8 drawn samples and performs the decision without any meta-verification. ii) **AutoMix + Thresholding**: Routes queries to the LLM if the verifier confidence is below a dataset-specific threshold (optimized on the validation set). We use a threshold for each dataset that yields the highest Δ_{IBC} on the validation set. iii) **AutoMix + POMDP**: This method optimizes routing decisions using a POMDP solver (Smith and Simmons, 2006) as a meta-verifier. The POMDP is learned on the validation set, and takes decisions based on the verifier outputs (detailed in Appendix B.1).

4.1 Main Results

Figure 4 shows performance vs. cost curves for various datasets and model-mixing methods. On 4 of the 5 datasets, `AutoMix`-POMDP and `AutoMix`-Threshold outperform `FrugalGPT`, staying above the SLM-LLM curve and yielding better performance per unit cost. Gains achieved by `AutoMix` over `FrugalGPT` are impressive because `FrugalGPT` has access to domain-specific trained routers and incurs no verification cost.

Further, `AutoMix`-POMDP shows consistent positive Δ_{IBC} across all evaluated costs. These results show that `AutoMix`, utilizing self-verification and meta-verification, can effectively mix LLAMA2-13B and GPT-4 on a wide range of tasks without access to model weights or domain-specific routing data.

	`AutoMix` + P	`AutoMix` + T	FrugalGPT
DIPLOMAT	**58.5**	50.1	-7.0
MUTUAL	**12.4**	11.8	-8.7
COQA	83.1	**86.5**	2.6
QASPER	8.5	-0.2	**9.4**
QUALITY	**10.3**	9.4	-4.7

Table 1: **Δ_{IBC} values:** `AutoMix` + T and `AutoMix` + P are variations of our proposed method with thresholding (T) and POMDP (P) based meta-verifiers, respectively. **`AutoMix` + POMDP** demonstrates a robust and consistent Δ_{IBC} across all datasets, implying a judicious utilization of computational resources. Despite having access to domain specific training and a 0-cost verifier, FrugalGPT underperforms `AutoMix` on 4/5 datasets.

Table 1 compares the average Δ_{IBC} across across five cost regions of equal size for each method and dataset. `AutoMix`-POMDP outperforms `FrugalGPT` on 4/5 datasets. On COQA and DIPLOMAT, it achieves substantial Δ_{IBC} gains of 86% and 58%, respectively. Even where these gains are less pronounced, as in QASPER, `AutoMix`-POMDP remains highly competitive with `FrugalGPT`.

Figure 5 shows the accuracy of using POMDP-based meta-verifier over Verifier-SC. We see significant improvements across all datasets, with absolute gains of up to 17%, demonstrating our proposed meta-verifier's importance in few-shot verification setups. Notably, even modest savings in computational cost can translate to significant fi-

nancial implications at the scale of LLM operations, underscoring the economic relevance of our approach.

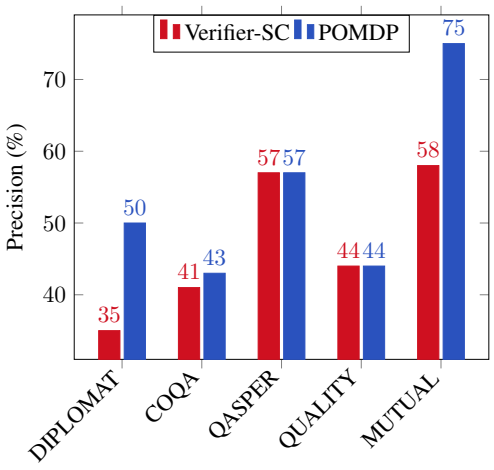

Figure 5: **Right:** The precision of the meta-verifier for both POMDP and Verifier-Self-Consistency (Verifier-SC) approaches across various datasets. Across all scenarios, the POMDP method consistently wins or ties with up to 43% relative performance gains.

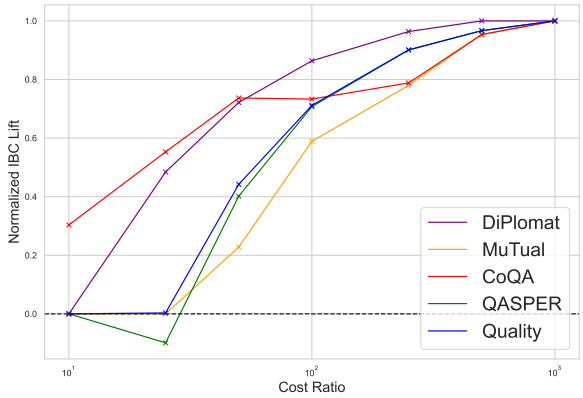

Figure 6: Normalized Δ_{IBC} for different cost regions. Increasing cost-ratio, results in better performance, highlighting an important criteria users need to consider while using this method.

5 Analysis

5.1 Key findings and takeaway

Effect on Cost Ration on `AutoMix` Our main experiments assumed a cost ratio of 1:100 between locally hosted LLAMA2-13B and GPT-4. Next, we analyze how changes in the cost ratio influence the Incremental Benefit-Cost (IBC) values across different settings. The results in Figure 6 show that for a cost ratio as low as 1:10, AutoMix starts delivering better performance per unit cost.

AutoMix is Effective in Low-Resource Scenarios Figure 9 (Appendix) demonstrates the performance dynamics of AutoMix and FrugalGPT with varying validation sizes. Notably, our method significantly outperforms FrugalGPT with limited data, despite the latter's domain-specific training and zero verifier cost. However, as training data increases, FrugalGPT narrows the performance gap by leveraging domain-specific training, albeit still trailing by 20%. This pattern indicates that AutoMix provides a particularly advantageous solution in real-world scenarios where data may be scarce.

Effectiveness of Few-shot Self-Verification In Appendix A.1, we evaluate few-shot self-verification quantitatively and qualitatively. We observe that the self-verification can effectively use context to identify errors in answers generated by SLM in many cases.

Improving Self-Verification with Task-Specific Prompt Engineering We explore the impact of task-specific prompt engineering on self-verification performance in Appendix A.2. While prompt engineering improves verifier accuracy, our meta-verifier remains robust in various settings and can beneficially leverage even a weak verifier.

5.2 Results of Automix w/ 3 Models

In this section, we evaluate the performance of AutoMix when applied to a three-model scenario, a setting we call AutoMix_3. Specifically, we employ LLAMA2-13B as the SLM, LLAMA2-70B as the MLM, and GPT-4 as the LLM. The results of this evaluation are presented in Figure 7.

AutoMix_3 consistently outperforms the baselines across the cost regions. We also compare AutoMix_3 against a baseline, *Union* AutoMix, which selects between the two-model variants $\text{AutoMix}_{SLM-MLM}$ and $\text{AutoMix}_{MLM-LLM}$, depending on the cost requirements specified by the end-user. For instance, if the desired average cost is less than that of the MLM, $\text{AutoMix}_{SLM-MLM}$ is employed, whereas $\text{AutoMix}_{MLM-LLM}$ is utilized for cost regions exceeding that of the MLM.

Further, we consider a baseline, Chained AutoMix, by chaining $\text{AutoMix}_{SLM-MLM}$ with $\text{AutoMix}_{MLM-LLM}$. The query first goes to the SLM, and an $\text{AutoMix}_{\text{SLM-MLM}}$ decides between reporting the SLM answer or rout-

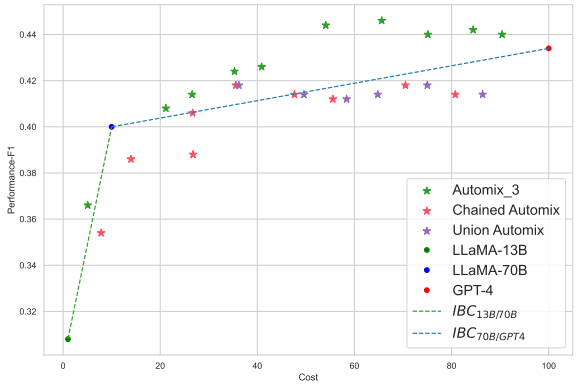

Figure 7: AutoMix with 3 models: LLAMA2-13B, GPT-4 and GPT-4. AutoMix method shows consistent IBC lifts for both SLM-MLM and MLM-LLM regions. Further, compared to chaining two AutoMix models or using the union of two AutoMixes, AutoMix_3 provide significant improvements.

ing to the MLM. In the latter's case, a second $\text{AutoMix}_{\text{MLM}-\text{LLM}}$ repeats the procedure using the MLM and LLM models. Chained AutoMix underperforms across the board, as it cannot directly route queries from the SLM to the LLM. Additionally, whenever 'Chained AutoMix' prompts the MLM, it invariably uses the costly verifier, even in cases where it might not be necessary. We refer readers to Appendix B.2, C, D.1 for more details.

6 Related Work

Self-Verification AutoMix aligns in spirit with works that aim to perform self-verification for reasoning problems, such as Weng et al. (2023); Jiang et al. (2023) (see Pan et al. (2023a) for a survey of recent self-verification and correction approaches). However, AutoMix uniquely harnesses context for verification instead of relying on LLM's knowledge (Dhuliawala et al., 2023), which can be challenging for reasoning problems (Madaan et al., 2023; Huang et al., 2023), and introduces a meta-verifier mechanism to offset the verifier's potential noise. Further, unlike Madaan et al. (2022), who utilize a corpus of past mistakes to gauge the likelihood of a model error for a new question, AutoMix uniquely utilizes context for verification. Finally, different from works that rely on external knowledge bases for verifying the outputs of language models (Peng et al., 2023; Gao et al., 2023; Pan et al., 2023b), AutoMix uses the context supplied with the question to verify the answer.

Our meta-verification approach can also be seen in the context of conformal prediction (Angelopoulos et al., 2023; Vovk et al., 2005) for a more robust self-verification. Ren et al. (2023) tie meta-verification more closely with conformal predictions for robot navigation, showing that layering predictions from a language model with a secondary mechanism help in identifying situations that do not have adequate information for action.

Mixing Models Distinct from related work optimizing LLM inference cost by model switching and external verifiers (Chen et al., 2023; Zhu et al., 2023; vSakota et al., 2023), AutoMix obviates the need for verifier training through few-shot SLM model prompting and does not require upfront access to all input queries. When needed, the meta-verifier learned with only as few as 200 samples outperforms training specialized models. Our work is thus aligned with recent work that aims at composing different models and external tools for inference time improvement of language models (Khattab et al., 2023; Press et al., 2022; Yao et al., 2022; Zhou et al., 2022).

Adaptive Computation In contrast to adaptive computation and model routing methods that preempt computation via intermediate representations (Liu et al., 2020; Zhou et al., 2020; Schuster et al., 2021; Geng et al., 2021; Schuster et al., 2022; Madaan and Yang, 2022), AutoMix necessitates no architectural modifications and assumes only black-box access to APIs. Further, unlike AdaptiveConsistency (Aggarwal et al., 2023), which optimizes inference within a single LLM model, AutoMix flexibly optimizes between two models and transcends its utility in Self-Consistency.

7 Conclusion

AutoMix integrates black-box large language model (LLM) APIs into a multi-step problem-solving framework, optimizing the computational cost and performance trade-offs. AutoMix opens avenues for several interesting research directions. First, while self-verification and correction are challenging for LLMs in general, we find promising results using context-grounded few-shot verification, indicating that similar approaches may yield gain in other scenarios. Secondly, our work interweaves Good Old-Fashioned Artificial Intelligence (GOFAI) approaches with LLMs, demonstrating that the incorporation of a POMDP can boost the accuracy of a noisy few-shot verifier.

Limitations

While our empirical evidence demonstrates effectiveness, the broader applicability of AutoMix may vary depending on the specific models and datasets used. Further, AutoMix assumes a context-grounded reasoning setup for effective self-verification, which excludes tasks like factual question-answering and commonsense reasoning. Finally, as open-source models get powerful and inference costs decrease, serving a strong model for all queries might be feasible. However, there are still likely going to be latency and availability trade-offs that might be handled using AutoMix.

References

Pranjal Aggarwal, Aman Madaan, Yiming Yang, and Mausam. 2023. Let's sample step by step: Adaptive-consistency for efficient reasoning with llms. *ArXiv*, abs/2305.11860.

Anastasios N Angelopoulos, Stephen Bates, et al. 2023. Conformal prediction: A gentle introduction. *Foundations and Trends® in Machine Learning*, 16(4):494–591.

Lingjiao Chen, Matei A. Zaharia, and James Y. Zou. 2023. Frugalgpt: How to use large language models while reducing cost and improving performance. *ArXiv*, abs/2305.05176.

Leyang Cui, Yu Wu, Shujie Liu, Yue Zhang, and Ming Zhou. 2020. Mutual: A dataset for multi-turn dialogue reasoning. *ArXiv*, abs/2004.04494.

Ido Dagan, Oren Glickman, and Bernardo Magnini. 2005. The pascal recognising textual entailment challenge. In *Machine learning challenges workshop*, pages 177–190. Springer.

Ido Dagan, Dan Roth, Fabio Zanzotto, and Mark Sammons. 2022. *Recognizing textual entailment: Models and applications*. Springer Nature.

Pradeep Dasigi, Kyle Lo, Iz Beltagy, Arman Cohan, Noah A Smith, and Matt Gardner. 2021. A dataset of information-seeking questions and answers anchored in research papers. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4599–4610.

Mostafa Dehghani, Yi Tay, Anurag Arnab, Lucas Beyer, and Ashish Vaswani. 2021. The efficiency misnomer. In *International Conference on Learning Representations*.

Shehzaad Dhuliawala, Mojtaba Komeili, Jing Xu, Roberta Raileanu, Xian Li, Asli Celikyilmaz, and Jason Weston. 2023. Chain-of-verification reduces hallucination in large language models. *arXiv preprint arXiv:2309.11495*.

Nouha Dziri, Ximing Lu, Melanie Sclar, Xiang Lorraine Li, Liwei Jian, Bill Yuchen Lin, Peter West, Chandra Bhagavatula, Ronan Le Bras, Jena D Hwang, et al. 2023. Faith and fate: Limits of transformers on compositionality. *arXiv preprint arXiv:2305.18654*.

Luyu Gao, Zhuyun Dai, Panupong Pasupat, Anthony Chen, Arun Tejasvi Chaganty, Yicheng Fan, Vincent Zhao, Ni Lao, Hongrae Lee, Da-Cheng Juan, et al. 2023. Rarr: Researching and revising what language models say, using language models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16477–16508.

Shijie Geng, Peng Gao, Zuohui Fu, and Yongfeng Zhang. 2021. Romebert: Robust training of multi-exit bert. *arXiv preprint arXiv:2101.09755*.

Jie Huang, Xinyun Chen, Swaroop Mishra, Huaixiu Steven Zheng, Adams Wei Yu, Xinying Song, and Denny Zhou. 2023. Large language models cannot self-correct reasoning yet. *arXiv preprint arXiv:2310.01798*.

Weisen Jiang, Han Shi, Longhui Yu, Zhengying Liu, Yu Zhang, Zhenguo Li, and James T Kwok. 2023. Backward reasoning in large language models for verification. *arXiv preprint arXiv:2308.07758*.

Omar Khattab, Arnav Singhvi, Paridhi Maheshwari, Zhiyuan Zhang, Keshav Santhanam, Sri Vardhamanan, Saiful Haq, Ashutosh Sharma, Thomas T Joshi, Hanna Moazam, et al. 2023. Dspy: Compiling declarative language model calls into self-improving pipelines. *arXiv preprint arXiv:2310.03714*.

Tomáš Kočiskỳ, Jonathan Schwarz, Phil Blunsom, Chris Dyer, Karl Moritz Hermann, Gábor Melis, and Edward Grefenstette. 2018. The narrativeqa reading comprehension challenge. *Transactions of the Association for Computational Linguistics*, 6:317–328.

Yuta Koreeda and Christopher D Manning. 2021. Contractnli: A dataset for document-level natural language inference for contracts. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 1907–1919.

Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*.

Teven Le Scao and Alexander M Rush. 2021. How Many Data Points is a Prompt Worth? In *NAACL*.

Hengli Li, Songchun Zhu, and Zilong Zheng. 2023. Diplomat: A dialogue dataset for situated pragmatic reasoning. *ArXiv*, abs/2306.09030.

Jiachang Liu, Dinghan Shen, Yizhe Zhang, Bill Dolan, Lawrence Carin, and Weizhu Chen. 2021a. What Makes Good In-Context Examples for GPT-3? *arXiv:2101.06804 [cs]*. ArXiv: 2101.06804.

Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2021b. Pre-train, Prompt, and Predict: A Systematic Survey of Prompting Methods in Natural Language Processing. *arXiv preprint arXiv:2107.13586*.

Weijie Liu, Peng Zhou, Zhiruo Wang, Zhe Zhao, Haotang Deng, and Qi Ju. 2020. Fastbert: a self-distilling bert with adaptive inference time. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6035–6044.

Aman Madaan, Niket Tandon, Peter Clark, and Yiming Yang. 2022. Memory-assisted prompt editing to improve GPT-3 after deployment. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2833–2861, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al. 2023. Self-refine: Iterative refinement with self-feedback. *arXiv preprint arXiv:2303.17651*.

Aman Madaan and Yiming Yang. 2022. Flowgen: Fast and slow graph generation. *arXiv preprint arXiv:2207.07656*.

Swaroop Mishra, Daniel Khashabi, Chitta Baral, Yejin Choi, and Hannaneh Hajishirzi. 2021. Reframing Instructional Prompts to GPTk's Language. *arXiv preprint arXiv:2109.07830*.

George E. Monahan. 1982. State of the art—a survey of partially observable markov decision processes: Theory, models, and algorithms. *Management Science*, 28:1–16.

OpenAI. 2023. Gpt-4 technical report.

Liangming Pan, Michael Saxon, Wenda Xu, Deepak Nathani, Xinyi Wang, and William Yang Wang. 2023a. Automatically correcting large language models: Surveying the landscape of diverse self-correction strategies. *arXiv preprint arXiv:2308.03188*.

Liangming Pan, Xiaobao Wu, Xinyuan Lu, Anh Tuan Luu, William Yang Wang, Min-Yen Kan, and Preslav Nakov. 2023b. Fact-checking complex claims with program-guided reasoning. *arXiv preprint arXiv:2305.12744*.

Richard Yuanzhe Pang, Alicia Parrish, Nitish Joshi, Nikita Nangia, Jason Phang, Angelica Chen, Vishakh Padmakumar, Johnny Ma, Jana Thompson, He He, et al. 2022. Quality: Question answering with long input texts, yes! In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5336–5358.

Baolin Peng, Michel Galley, Pengcheng He, Hao Cheng, Yujia Xie, Yu Hu, Qiuyuan Huang, Lars Liden, Zhou Yu, Weizhu Chen, et al. 2023. Check your facts and try again: Improving large language models with external knowledge and automated feedback. *arXiv preprint arXiv:2302.12813*.

Adam Poliak. 2020. A survey on recognizing textual entailment as an nlp evaluation. *arXiv preprint arXiv:2010.03061*.

Ofir Press, Muru Zhang, Sewon Min, Ludwig Schmidt, Noah A Smith, and Mike Lewis. 2022. Measuring and narrowing the compositionality gap in language models. *arXiv preprint arXiv:2210.03350*.

Siva Reddy, Danqi Chen, and Christopher D Manning. 2019. Coqa: A conversational question answering challenge. *Transactions of the Association for Computational Linguistics*, 7:249–266.

Machel Reid and Graham Neubig. 2022. Learning to model editing processes. *arXiv preprint arXiv:2205.12374*.

Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.

Allen Z Ren, Anushri Dixit, Alexandra Bodrova, Sumeet Singh, Stephen Tu, Noah Brown, Peng Xu, Leila Takayama, Fei Xia, Jake Varley, et al. 2023. Robots that ask for help: Uncertainty alignment for large language model planners. *arXiv preprint arXiv:2307.01928*.

Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.

Timo Schick, Jane Dwivedi-Yu, Zhengbao Jiang, Fabio Petroni, Patrick Lewis, Gautier Izacard, Qingfei You, Christoforos Nalmpantis, Edouard Grave, and Sebastian Riedel. 2022. Peer: A collaborative language model.

Tal Schuster, Adam Fisch, Jai Gupta, Mostafa Dehghani, Dara Bahri, Vinh Q Tran, Yi Tay, and Donald Metzler. 2022. Confident adaptive language modeling. *arXiv preprint arXiv:2207.07061*.

Tal Schuster, Adam Fisch, Tommi Jaakkola, and Regina Barzilay. 2021. Consistent accelerated inference via confident adaptive transformers. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4962–4979.

Uri Shaham, Maor Ivgi, Avia Efrat, Jonathan Berant, and Omer Levy. 2023. Zeroscrolls: A zero-shot benchmark for long text understanding.

Uri Shaham, Elad Segal, Maor Ivgi, Avia Efrat, Ori Yoran, Adi Haviv, Ankit Gupta, Wenhan Xiong, Mor Geva, Jonathan Berant, and Omer Levy. 2022. SCROLLS: Standardized CompaRison over long language sequences. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 12007–12021, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Noah Shinn, Beck Labash, and Ashwin Gopinath. 2023. Reflexion: an autonomous agent with dynamic memory and self-reflection.

Trey Smith and Reid Simmons. 2006. Focused real-time dynamic programming for mdps: Squeezing more out of a heuristic. In *AAAI*, pages 1227–1232.

Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.

Gladys Tyen, Hassan Mansoor, Peter Chen, Tony Mak, and Victor Cărbune. 2023. Llms cannot find reasoning errors, but can correct them! *arXiv preprint arXiv:2311.08516*.

Vladimir Vovk, Alexander Gammerman, and Glenn Shafer. 2005. *Algorithmic learning in a random world*, volume 29. Springer.

Marija vSakota, Maxime Peyrard, and Robert West. 2023. Fly-swat or cannon? cost-effective language model choice via meta-modeling. *ArXiv*, abs/2308.06077.

Sean Welleck, Ximing Lu, Peter West, Faeze Brahman, Tianxiao Shen, Daniel Khashabi, and Yejin Choi. 2022. Generating sequences by learning to self-correct. *arXiv preprint arXiv:2211.00053*.

Yixuan Weng, Minjun Zhu, Shizhu He, Kang Liu, and Jun Zhao. 2022. Large language models are reasoners with self-verification. *arXiv preprint arXiv:2212.09561*.

Yixuan Weng, Minjun Zhu, Fei Xia, Bin Li, Shizhu He, Kang Liu, and Jun Zhao. 2023. Large language models are better reasoners with self-verification. *CoRR, abs/2212.09561*.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R Narasimhan, and Yuan Cao. 2022. React: Synergizing reasoning and acting in language models. In *The Eleventh International Conference on Learning Representations*.

Shuyan Zhou, Uri Alon, Frank F Xu, Zhengbao Jiang, and Graham Neubig. 2022. Docprompting: Generating code by retrieving the docs. In *The Eleventh International Conference on Learning Representations*.

Wangchunshu Zhou, Canwen Xu, Tao Ge, Julian McAuley, Ke Xu, and Furu Wei. 2020. Bert loses patience: Fast and robust inference with early exit. *Advances in Neural Information Processing Systems*, 33:18330–18341.

Banghua Zhu, Ying Sheng, Lianmin Zheng, Clark W. Barrett, Michael I. Jordan, and Jiantao Jiao. 2023. On optimal caching and model multiplexing for large model inference. *ArXiv*, abs/2306.02003.

Attribute	Content
Context	Which country grows the most tea? ... Shen sipped it. He liked it. He drank it all. Shen was proud of his new drink. *Truncated for brevity*
Question	When did he take it?
Generated Answer	He took it in 1990.
Gold Answer	about 2737 B.C
Verifier Output	The context does not mention that he took tea in 1990. Verification Decision: The AI generated answer is Incorrect.

Table 2: An example where the generated answer is verifiably incorrect, citing an unsupported date (1990), and the verifier successfully catches this discrepancy.

A Verifier Qualitative Analysis

A.1 How effective is few-shot self-verification?

One notable contribution of this work is the concept of few-shot self-verification of outputs. Self-Verification, especially for reasoning problems, poses its own set of challenges; however, our setup has a unique advantage: the capacity to utilize context to validate answers. For instance, the model can identify factual inaccuracies in the answer or discern apparent contradictions that might not have been evident during the initial response. But does this advantage translate to effective self-verification in practice? As depicted in Figure 8, aside from the CNLI dataset, few-shot self-verification succeeds in accurately identifying correct examples by assigning them higher probabilities across all other datasets.

Qualitative Analysis Representative Examples from our qualitative analysis are shown in Tables 2, 3, and 4.

FrugalGPT vs. `AutoMix` at different levels of data availability Figure 9 demonstrates the performance dynamics of `AutoMix` and FrugalGPT with varying validation data sizes. Notably, our method significantly outperforms FrugalGPT with limited data, despite the latter's domain-specific training and zero verifier cost. However, as training data increases, FrugalGPT narrows the performance gap by leveraging its domain-specific training. This pattern indicates that `AutoMix` provides a particularly advantageous solution in real-world scenarios where data may be scarce.

A.2 Domain-specific vs. Domain independent verifier

We used a single verifier with the LLAMA2-13B model to help steer the model. To avoid excessive prompt engineering, we used a generic prompt for all datasets. However, task-specific prompts generally help (Le Scao and Rush, 2021; Liu et al., 2021b; Mishra et al., 2021; Liu et al., 2021a). To investigate this, we create task-specific prompts for CNLI by giving examples from the legal domain in the prompt. We use LLAMA2-70B instead of GPT-4 for this experiment, due to cost constraints.

Figure 10 underscores the efficacy of employing task-specific verification prompts, ensuring a heightened probability allocation for accurate examples during the verification process. Interestingly, the enhanced verifier accuracy does not always directly translate to proportionate improvements in our proposed method, `AutoMix`, as evidenced in Table 5. This phenomenon higlights the role of meta-verifiers, adeptly negotiating through the outputs of potentially unreliable verifiers.

B Methodology

B.1 POMDP

The Partially Observable Markov Decision Process (POMDP) presents a robust framework for handling decision-making problems under uncertainty, offering a structured way to manage and navigate through the decision spaces where the system's state is not fully observable (Monahan, 1982). A POMDP is defined by a tuple (S, A, T, R, Ω, O), where S is a set of states, A is a set of actions, T represents the state transition probabilities, R is the reward function, Ω is a set of observations, and O

Attribute	Content
Context	Fort Lauderdale, Florida (CNN) – Just taking a sip of water or walking to the bathroom is excruciatingly painful for 15-year-old Michael Brewer, ... He suffered second- and third-degree burns over about two-thirds of his body, according to the hospital's associate director, Dr. Carl Schulman. *Truncated for brevity*
Question	How long before he recovers from his injury?
Generated Answer	he will never recover from his injury.
Gold Answer	He has a lifelong recovery.
Verifier Output	The context clearly states that the teen is recovering from his injuries. Verification Decision: The AI generated answer is Incorrect.

Table 3: The example demonstrates a challenging case for the verifier. The generated answer is a rephrasing of what's mentioned in the context, however the verifier fails to catch it.

Attribute	Content
Context	Pruning keeps the plant only three or four feet tall. This is an easy height for tea picking. Green tea is steamed right after the leaves are picked. Green tea _ its healthy qualities. For example, it may prevent heart disease. How did we get tea bag? The answer: by accident. *Truncated*
Question	What did they do to green tea after picking it?
SLM Answer	steamed.
Gold Answer	prune it
Verifier Output	The context mentions that green tea is steamed after being picked. Verification Decision: The AI generated answer is Correct.

Table 4: An instance where the verifier deems the answer correct, and the gold label was incorrect. The verifier is able to correctly infer that the context mentions tea being steamed after picking.

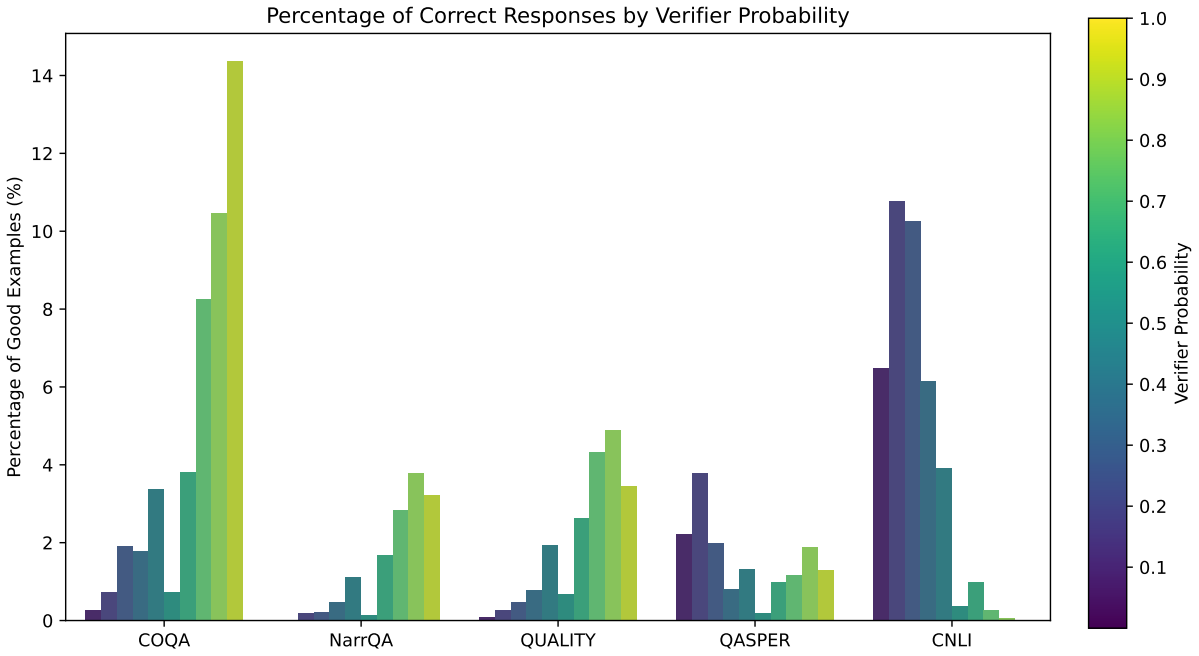

Figure 8: **Verifier Probability and Correctness:** Percentage of correct responses across distinct verifier probability bins, representing $P(\mathcal{C} = 1 \,|\, A_{\text{SLM}}, \mathcal{C}, q)$, where A_{SLM} is the answer from the Small Language Model, $\mathcal{C}$ is the context, and q is the query. Each bin represents a range of verifier probabilities and the corresponding accuracy of the responses within that probability range across various datasets. Notably, for all datasets, excluding CNLI and QASPER, a higher verification score generally corresponds to a larger proportion of correct examples, indicating that the verifier is, to an extent, capable of discerning the reliability of responses generated by itself. We use a meta-verifier to get around these noisy predictions.

	CNLI			CNLI-CV		
Method	Cost	Perf.	IBC_Lift	Cost	Perf.	IBC_Lift
SLM	1	40.1	-	1	40.1	-
FrugalGPT	37.4	59.2	**66.1**	37.4	59.2	**66.1**
Self-Consistency	47.5	52.3	-17.0	40.5	50.6	-15.5
`AutoMix`-Threshold	51.9	55.6	-3.5	28.1	46.9	-49.1
`AutoMix`-POMDP	6.7	43.5	88.7	15.8	45.2	12.4
LLM	50	55.5	-	50	55.5	-

Table 5: Despite the boost in verifier accuracy with task-specific prompts (Figure 10), `AutoMix` may not always benefit, highlighting the utility of even weak verifiers when supported by meta-verifiers.

is the observation function.

In the context of meta-verifier, the *unobservable* states (S) represent the potential correctness of the verifier's predictions, categorized as *Simple*, *Complex*, and *Unsolvable*. Note that incase of non-binary evaluation (e.g., F1-Score), *Unsolvable* indicates both the SLM and LLM have similar low performance on the input problem. Actions (A) are binary: trust the verifier or invoke the LLM. The reward function (R) quantifies the cost or gain of making a particular action in a given state, steering the decision policy towards cost-effective actions. Observations (Ω) in our model are the verifier's probability outputs, discretized into bins. Specifically, we generate k=8 samples from the verifier, discretizing our observation space in intervals of size 0.125 ranging from 0 to 1.

The observation function (O) depicts the likelihood of observing an observation given an action was taken and the system transitioned to a particular state. Using an appropriate observation function is crucial for POMDP to work. Specifically, we define observations probabilities in three ways:

- **1. Functional Form:** For each of the states s, the observation function O is defined as $O(s, v) = \frac{1}{K} \cdot v^{\gamma_s}$, where v is the verifier probability and $\gamma_s \in [0, \infty]$ is a hyperparameter for every state and K is normalizing factor. Intutively, a value of γ close to 1 indicates ideal calibration, with verifier probability v indicating true probability of being in a particular state. The values of γ_s's for the three states are determined based on the respective POMDP's performance on validation set based on the IBC-Lift.
- **2. Discrete Form:** An alternate option is to directly learn observation function O from the statistics of validation set. Since in valida-

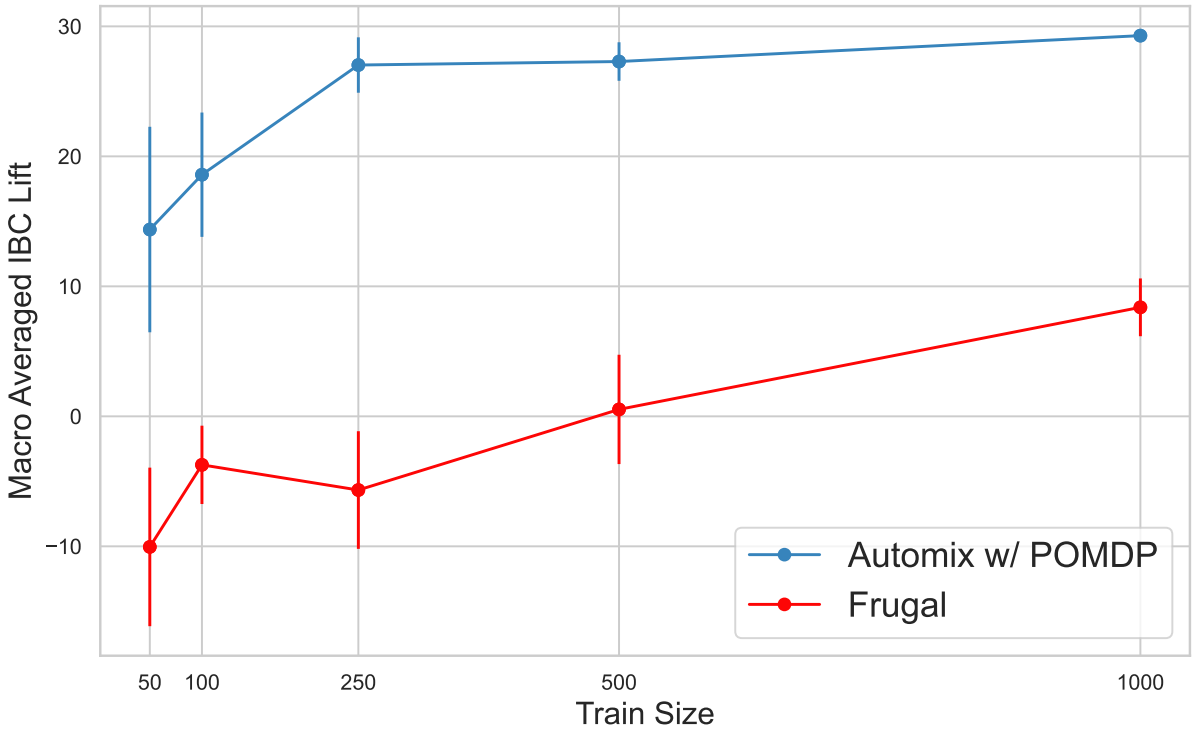

Figure 9: Comparison of AutoMix with FrugalGPT over varying Training Dataset Size. Despite zero-cost verifier and domain-specific training, FrugalGPT underperforms AutoMix. AutoMix is especially useful for limited data settings, with higher gains visible when dataset size is less than 1000.

tion set, we have access to the true state along with verifier probabilites of individual data instances, we can model observation function as $O(s, v) = \frac{\Sigma_{i=0}^{N} \mathbf{1}\{s_i=s \text{ and } v_i=v\}}{\Sigma_{i=0}^{N} \mathbf{1}\{s_i=s\}}$. The method has the advantage of being hyperparameter free and provides more accurate representation by computing the true observation probabilities on validation set. However, it performs worse than functional form, when either certain values of v or s are not well represented in validation set or in cases of high distribution shift between validation and test set.

- **3. Continous Form:** The continous form of POMDP follows the same formulation as in Discrete Form, except the fact the state space is represented by a tuple of SLM & LLM performance. Specifically, state space is represented by $\mathcal{S} = \{(P_{SLM}, P_{LLM}) | P_{SLM}, PLLM \in [0, 1]\}$, where P represents the performance of corresponding model on particular question. Since the performance (eg: F1 score) can be continous values, while we have discrete data (performance on individual scores), we apply gaussian smoothing (with standard deviation 1) followed by linear interpolation, to get observation probabilities for this continous state space.

Since both these methods have their strengths, and are independent of each other, we choose the best performing method on validation set.

This POMDP mechanism allows for optimal decision-making under uncertainty, balancing the cost and reliability of invoking the LLM. Through employing standard POMDP solving algorithms such as Focused Real-Time Dynamic Programming[3] (Smith and Simmons, 2006), we derive a policy that maps belief states (probability distributions over S) to actions. During inference, the learned policy effectively decides whether to trust the verifier's output or to invoke the LLM based on a combination of expected future rewards and computational costs.

Another advantage of the POMDP-based meta-verifier is its interpretability and customizability via reward assignment. For instance, in a "Needy" state, assigning a reward of +50 for invoking the LLM indicates a preference for accurate solutions

[3]We use zmdp package https://github.com/trey0/zmdp for solving POMDP

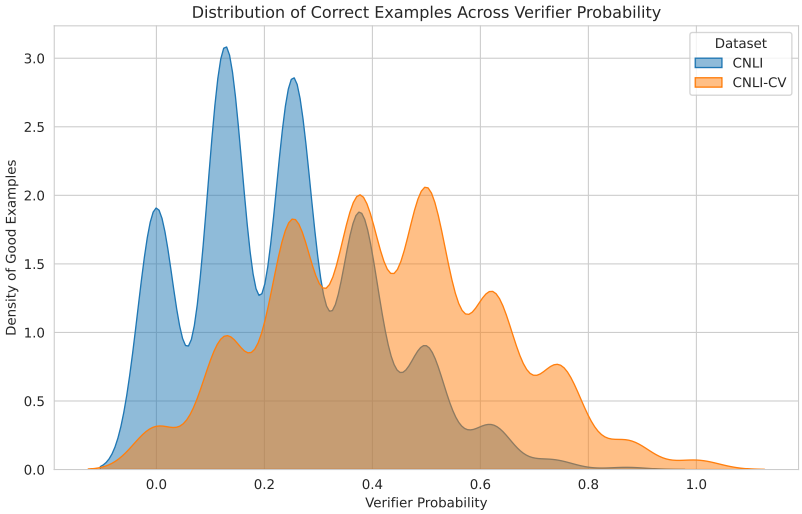

Figure 10: Enhancement of verifier accuracy using task-specific verification prompts, which allocate higher verification probabilities to more correct examples.

over computational cost. Conversely, in a "Good" state, designating a reward of -10 for trusting the SLM encourages computational savings. This enables users to strategically balance solution quality against computational expenses, aligning with specific application needs.

B.2 Integrating Three Models with AutoMix

While the fundamental approach remains consistent, the three-model scenario diverges from its two-model counterpart in two key aspects: 1) the definition of observation probabilities, and 2) the evaluation methodology.

We employ a formulation akin to the continuous form of POMDP, as described in the previous section. However, in contrast to the two-model scenario, the observations can now fall into two categories: a) SLM verifier outputs on SLM answer, and b) SLM verifier outputs on SLM answer combined with MLM verifier outputs on MLM answer. The second category allows us to model more nuanced cues regarding the impact of verifiers on the final performance improvement. For instance, Figure 12 illustrates that when both verification probabilities are available, high $\delta_{MLM-LLM}F1$ regions can be detected, which is not feasible with a single verifier. This implies that the POMDP can make more informed decisions, an advantage that is empirically demonstrated in Results 5.2.

In terms of evaluation, we consider two separate cases: 1) when the SLM-MLM-LLM curve is convex, and 2) when the curve is concave. In the convex case (as observed in the COQA dataset), it is advantageous to choose between the MLM and SLM in low-cost regions, while it is beneficial to choose between the MLM and LLM in high-cost regions. The suitable IBC curve is selected for evaluation accordingly. However, in the second case, when the IBC curves are concave, it would be more favorable to choose between the SLM and LLM, and completely ignore the MLM, as in terms of incremental performance per cost, it consistently presents a disadvantage. Thus, the $\text{IBC}_{SLM-LLM}$ is chosen for evaluation throughout. Although the evaluation presents two distinct cases, our $\texttt{AutoMix}_3$ framework is sufficiently general to identify instances where direct routing to LLM is needed even in the convex case, and also pinpoint cases where routing to MLM is beneficial in the concave scenario. This flexibility results in significantly superior performance.

C Expanding AutoMix to Three-Models

The preceding discussion focused on a two-model scenario involving the SLM and LLM. This section extends this framework to incorporate a third model, the MLM.

Our decision flow commences with the SLM generating an answer, which is then self-verified by the SLM. The verifier probability serves as an observation, guiding one of the following actions: 1) Reporting the SLM answer, 2) Running inference on the MLM or LLM and reporting the answer, or 3) Running inference on the MLM and verifying the answer. If action 3 is chosen, AutoMix has access to verification probabilities from both

```
# Meta-verifier POMDP File for narrative_qa

discount: 0.99
values: reward

# We have 6 states: 3 corresponding to the initial state before verifier is
    called, and 3 corresponding to the state after verifier is called
states: START_S START_C START_U SIMPLE COMPLEX UNSOLVABLE

# Effectively, we have 3 actions: 1.) The initial State where we run verifier
    2.) Report SLM's Answer 3.) Invoke LLM and Report its Answer
actions: Init Trust_SLM Invoke_LLM

# Observations lies in one of verifier probability bins. Eg: bin_correct_high
    represents Verifier outputs SLM answer as correct with high confidence
observations: bin_incorrect_low bin_incorrect_high bin_correct_low
    bin_correct_high

# Transition Model for Init action

T: Init
# Format: start_state : end_state : Transition_Probability

# Transition Model for Trust_SLM action
T: Trust_SLM
identity

# Transition Model for Invoke_LLM action
T: Invoke_LLM
identity

# Observation Model after "Init" action for narrative_qa
# Format: O : action : state : observation : probability

# Eaxmple: In SIMPLE cases, it is likely, SLM is correct and Verifier is
    Confident, while in UNSOLVABLE, SLM is incorrect (Lower Obs. Probability)
O : * : SIMPLE : bin_correct_high 0.8
O : * : COMPLEX : bin_correct_high 0.4
O : * : UNSOLVABLE : bin_correct_high 0.1

# Reward Model:
# Format: R: action : init_state : end_state : observation : probability

# Example: For COMPLEX state, Trusting SLM results in negative score, while
    invoking LLM results in a high +50 score.
R: Trust_SLM : COMPLEX : * : * -10
R: Invoke_LLM : COMPLEX : * : * +50
```

Figure 11: A sample POMDP specification file. POMDP requires defining states, actions, observations and relevant Transition, Observation Probabilities and Reward Values.

the SLM and MLM, which are used to decide whether to report the MLM's answer or switch to the LLM. Access to both the verifier probabilities provides AutoMix's meta-verifier with a richer observation signal. For instance, a neutral SLM verification signal combined with a neutral MLM verification signal will likely route the queries to the MLM. In comparison, an uncertain SLM verification signal and a neutral MLM verification signal will more likely be routed to LLM. In Section 5.2, we compare different variants of AutoMix, highlighting the individual importance of each state in AutoMix's formulation. Further details are provided in Appendix B.2.

Meta-Verifier in the Three-Model Case We employ a similar POMDP formulation as in the two-model scenario but with a broader range of actions due to the inclusion of the third model. The states are now represented as a tuple of performance metrics for each of the three models. Formally, the state space is denoted as $\mathcal{S} = \{(P_{SLM}, P_{MLM}, P_{LLM}) | P_{SLM}, P_{MLM}, PLLM \in [0, 1]\}$, where P denotes the performance of the respective model. For instance, if only the LLM

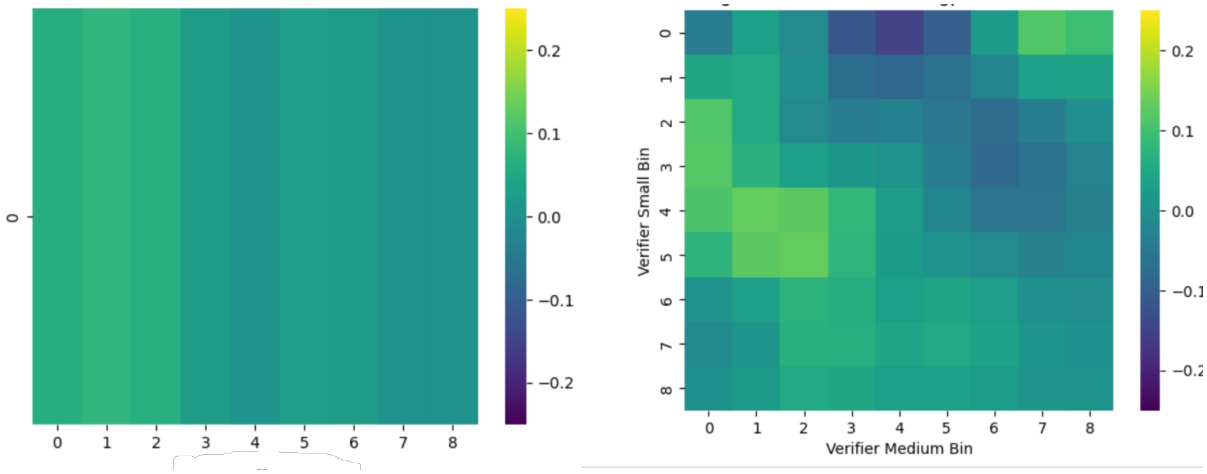

Figure 12: In the figure we compare delta improvement in F1 score from GPT-4 to GPT-4 on COQA dataset, for different verifier probabilities. The graphs are smoothened using gaussian smoothing with standard deviation=1. On left, we vary only the MLM verifier, and on right we vary both SLM and MLM verifiers. The latter case provides much richer, thus showing importance of incorporating both of them in our $\texttt{AutoMix}_3$ formulation.

can correctly solve the problem, the state will be represented as (0,0,1). AutoMix maintains a belief over all possible states and updates this belief based on the verifier probabilities, which serve as observations. The model can observe either the SLM verifier probability or the SLM and MLM verifier probabilities. The observation probabilities are learned from the validation set as in the previous section.

D Additional Details on the Experimental Setup

For evaluation, we utilize the validation sets from Shaham et al. (2022) for QASPER, and QUALITY, and use the prompts from Shaham et al. (2023). For COQA, MUTUAL, and DIPLOMAT, we employ its validation split and adapt the QUALITY prompt. For consistency, 1000 instances are sampled from the validation set of each dataset. Regardless of the dataset, identical input prompts are dispatched to both SLM and potentially LLM, ensuring consistent input processing costs. The output length is fixed in multichoice datasets like CNLI and QUALITY, and the brevity of responses in other datasets allows us to assume uniform output processing costs. We use greedy decoding (temperature 0) and draw a single sample for both the SLM and LLM. For verification, we generate eight samples per question (temperature = 1), which has negligible cost owing to the large context. In Figure 6, we normalize Δ_{IBC} by a scaling factor such that for all datasets, the maximum is set to 1.

For running our experiments, we use LLAMA2-13B and GPT-4 models from huggingface[4]. We use vllm (Kwon et al., 2023) for hosting models for inference.

Cost Ratio: We have considered a cost ratio of 1:100 between GPT-4 and GPT-4, reflecting the API price disparity between the models, which stands at \$0.225 for LLAMA2-13B vs \$30 for GPT-4 per 1M tokens at the time of writing. Additionally, for self-verification purposes, we generate 8 samples. It is important to note, however, that the cost of generating 8 samples is negligible compared to the cost of a single sample, primarily because the major cost driver is the length of the context (e.g., generation is 60 times and 50 times smaller for QASPER and QUALITY, respectively than base context). Therefore, invoking the verifier 8 times is considered equivalent in cost to calling it once. Furthermore, in Section 6, we explore different cost ratios and observe that, even with a ratio as low as 1:25, AutoMix begins to yield non-trivial gains across most datasets.

D.1 Results of Automix w/ 3 Models

In this section, we evaluate the performance of AutoMix when applied to a three-model scenario, as described in Section C. Specifically, we employ LLAMA2-13B as the SLM, GPT-4 as the MLM, and GPT-4 as the LLM. Due to cost constraints, our evaluation is conducted on a subset of 1000

[4]Models available at: https://huggingface.co/meta-llama/Llama-2-13b-hf and https://huggingface.co/meta-llama/Llama-2-70b-hf

examples from the COQA dataset. The results of this evaluation are presented in Figure 13.

Our findings reveal that AutoMix$_3$ consistently outperforms the IBC curve for both the SLM-MLM and MLM-LLM cost regions. We also compare AutoMix$_3$ against a baseline, *Union* AutoMix, which chooses between the two-model variants AutoMix$_{SLM-MLM}$ and AutoMix$_{MLM-LLM}$, depending on the cost requirements specified by the end-user. For instance, if the desired average cost is less than that of the MLM, AutoMix$_{SLM-MLM}$ is employed, whereas AutoMix$_{MLM-LLM}$ is utilized for cost regions exceeding that of the MLM. AutoMix$_3$ outperforms the baseline consistently on all cost regions. This better performance can be attributed to the fact that AutoMix$_3$ has access to verifier probabilities from both LLAMA2-13B and GPT-4, which provides a richer signal to POMDP, resulting in taking more informed actions. Further, we consider a baseline by chaining AutoMix$_{SLM-MLM}$ with AutoMix$_{MLM-LLM}$. The query first goes to the SLM, and an AutoMix$_{\text{SLM-MLM}}$ decides between reporting the SLM answer or routing to the MLM. In the latter's case, a second AutoMix$_{\text{MLM-LLM}}$ repeats the procedure using the MLM and LLM models. We call this method 'Chained AutoMix,' and it underperforms across the board. This is primarily because it cannot directly route queries from the SLM to the LLM. Additionally, whenever 'Chained AutoMix' prompts the MLM, it invariably uses the costly verifier, even in cases where it might not be necessary. This inefficient use of resources contributes to its subpar performance.

E Additional Datasets

In this section, we evaluate AutoMix on two additional datasets: CNLI (Koreeda and Manning, 2021) and NARRATIVE-QA (Kočiskỳ et al., 2018). CNLI involves performing natural language inference on non-disclosure documents as context, while NARRATIVE-QA focuses on question answering over full-length books and movie scripts.

The results on these datasets are presented in Figure 14. On NARRATIVE-QA, AutoMix significantly outperforms all other methods, including GPT-4, for higher cost values. However, it is noteworthy that the delta between the F1-Score of LLAMA2-13B and GPT-4 is only 2%, despite their considerable differences in capabilities. This discrepancy is attributed to the very long contexts in NARRATIVE-QA, and our qualitative analysis indicates that our employed retriever, all-MiniLM-L6-v2, cannot always retrieve relevant context. As a result, GPT-4 often fails to answer questions.

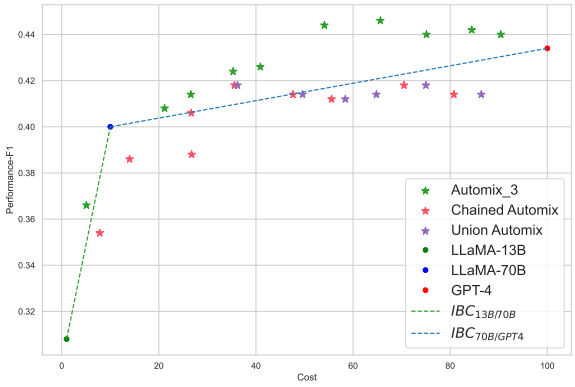

Figure 13: AutoMix with 3 models: LLAMA2-13B, GPT-4 and GPT-4. AutoMix method shows consistent IBC lifts for both SLM-MLM and MLM-LLM regions. Further, compared to chaining two AutoMix models or using the union of two AutoMixes, AutoMix$_3$ provide significant improvements.

In the case of CNLI, AutoMix-POMDP maintains a non-negative Δ_{IBC} throughout. However, FrugalGPT demonstrates superior performance. This better performance is because, despite the large number of contexts in CNLI, there are only 17 standard questions. These questions are semantically similar and often have a signature corresponding to the answer. For example, 12 out of 17 questions have less than 5% instances of Contradiction as the ground truth. Consequently, a verifier fine-tuned on such a dataset can easily learn to exploit these patterns. Our experiments show that a verifier trained solely on questions can perform comparably to FrugalGPT (23.4% average Δ_{IBC} compared to 27.2% of FrugalGPT). Since AutoMix is broadly applicable and is not dependent on the input question, FrugalGPT performs better on CNLI in particular.

F Few-Shot Prompts

F.1 Verifier Prompts

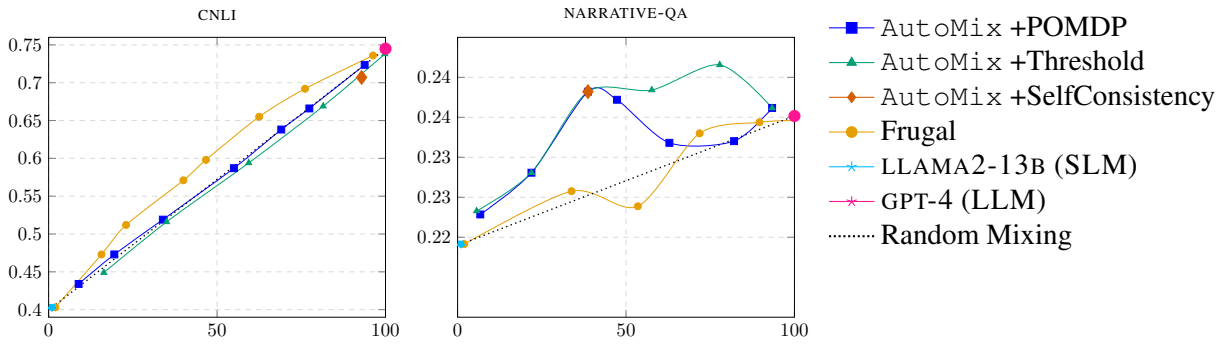

Figure 14: Aggregated performance **(y-axis)** vs. cost (**x-axis**) for different methods on the small and large LLAMA2-13/GPT-4. Results for additional datasets: NARRATIVE-QA and CNLI.

```
Story:
{relevant parts of the story}

{instruction}

Question: {question}

Answer:
```

Listing 2: **Task Prompt.** We experiment with long-context reasoning tasks, which require answering questions from stories, legal contracts, research papers, and novels.

```
Context: {context}

Question: {question}

AI Generated Answer: {generated_answer}

Instruction: Your task is to evaluate
↪  if the AI Generated Answer is
↪  correct, based on the provided
↪  context and question. Provide the
↪  judgement and reasoning for each
↪  case. Choose between Correct or
↪  Incorrect.

Evaluation:"'
```

Listing 3: **Verification Prompt.** The verification process is framed as a natural language entailment task, where the model determines the validity of the model-generated answer with respect to the context and question.

```
Context: The manuscript, discovered in 1980 in a dusty attic, turned out to be a lost work of Shakespeare.\n
Question: Whose lost work was discovered in a dusty attic in 1980?\n
AI Generated Answer: Shakespeare\n
Instruction: Your task is to evaluate if the AI Generated Answer is correct, based on the provided context and question. Provide the judgement and reasoning for each case. Choose between Correct or Incorrect.\n
Evaluation: The context specifically mentions that a lost work of Shakespeare was discovered in 1980 in a dusty attic.

Verification Decision: The AI generated answer is Correct.

---

Context: The celestial event, known as the Pink Moon, is unique to the month of April and has cultural significance in many indigenous tribes.\n
Question: In which month does the celestial event, the Pink Moon, occur?\n
AI Generated Answer: July\n
Instruction: Your task is to evaluate if the AI Generated Answer is correct, based on the provided context and question. Provide the judgement and reasoning for each case. Choose between Correct or Incorrect.\n
Evaluation: The context clearly states that the Pink Moon is unique to the month of April.

Verification Decision: The AI generated answer is Incorrect.

---

{truncated examples}

Context: {context}\n
Question: {question}\n
AI Generated Answer: {generated_answer}

Instruction: Your task is to evaluate if the AI Generated Answer is correct, based on the provided context and question. Provide the judgement and reasoning for each case. Choose between Correct or Incorrect.

Evaluation:
```

Listing 4: **Few-Shot Verifier Prompts:** 3-shot verifier prompt for evaluating the correctness of SLM's answer. The same prompt is used for all datasets.