# STEP-TAGGING: TOWARD CONTROLLING THE GENERATION OF LANGUAGE REASONING MODELS THROUGH STEP MONITORING

**Anonymous authors**Paper under double-blind review

#### **ABSTRACT**

The field of Language Reasoning Models (LRMs) has been very active over the past few years with advances in training and inference techniques enabling LRMs to reason longer, deeper, and more accurately. However, a growing body of studies show that LRMs are still inefficient, over-generating verification and self-reflection steps. To address this challenge, we introduce the Step-Tagging framework, a lightweight sentence-classifier enabling real-time annotation of the type of reasoning steps that an LRM is generating. To cover the wide space of reasoning behaviors, we introduced ReasonType: a novel taxonomy of reasoning steps. Building on this framework, we demonstrated that careful online monitoring of the count of specific steps can produce effective interpretable early stopping criteria of LRM inferences. We evaluate the Step-tagging framework on three open-source reasoning models across two standard benchmark datasets, MATH500 and GSM8K, and achieve 30 to 40% token reduction while maintaining comparable accuracy to standard generation. This work offers a novel way to increase control over the generation of LRMs, and a new tool to study behaviors of LRMs.

#### 1 Introduction

For the past few years, the field of Language Reasoning Models (LRMs) has experienced significant growth in terms of capabilities. Initiated by the pioneering work on model prompting such as Chain-of-Thought (Wei et al., 2023) and Self-Consistency (Wang et al., 2023), Inference Time Scaling has emerged as a popular field with the goal of making models more accurate at reasoning. At the same time, fundamental work on Reinforcement Learning (RL) and Supervised Fine-Tuning (SFT) as part of Training Time Scaling has led to the release of strong reasoning models.

However, recent surveys have shown that LRMs need to generate a very large number of tokens—several thousands—in order to generate an accurate answer on challenging questions (Qu et al., 2025; Chen et al., 2025b; Sui et al., 2025a). This behavior makes reasoning models extremely inefficient - scaling in both compute resources and inference time. Although recent works have suggested solutions to this problem, most of them overlook the possibility of monitoring the output generated by a model in order to dynamically manage the reasoning of the model, leaving a significant gap in the literature. To address this challenge, this paper aims to offer a new perspective on the efficiency of LRMs by focusing on online monitoring of models. Our contributions are as follows:

- **Formalization of the Reasoning step concept:** From the literature, we have observed numerous definitions of reasoning steps. We first present a comprehensive review and formalize a more generalized definition of what constitutes a reasoning step for LRMs. We then propose *ReasonType*, the first taxonomy of reasoning steps, enabling a structured identification of reasoning behaviors.
- **Step-Tagging module:** We introduce the *Step-Tagging* module (see Figure 1), an online lightweight sentence classifier capable of identifying the nature of each step that the LRMs are generating. This novel framework offers a tool to systematically monitor the generation of LRMs.
- Early-Stopping Framework: We observed that LRMs often generate the correct answer early in the output sequence. Leveraging the *Step-Tagging* module, we found that the type of reasoning steps plays a role in determining the early-stopping condition. Based on these observations, we built an interpretable early-stopping framework that dynamically stops token generation based on

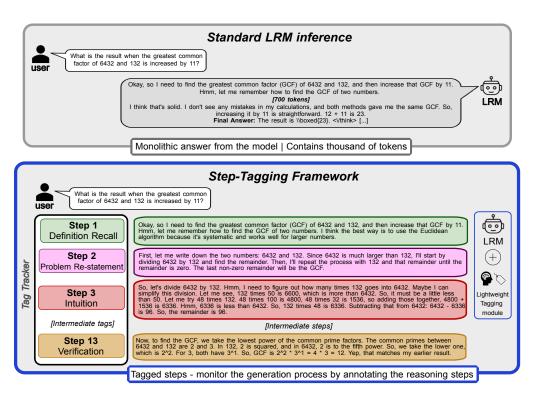


Figure 1: Step-Tagging: a framework for monitoring the generation of LRMs - example on sample 39 from MATH500 test with DS-Qwen14B, using the ReasonType taxonomy - seed 42

reasoning steps types and counts, calibrated on both models and problem complexity. Tested on three open-source LRMs across two reasoning datasets, our framework reduced token generation by 30-40% while maintaining a comparable accuracy.

The paper is organized as follows. We first review research on LRM efficiency. We then propose a definition of a reasoning step and a taxonomy of reasoning step types. Building on this, we present our Step-Tagging module that can segment and label the reasoning steps within the output of an LRM. We also propose an early stopping mechanism, based on frequency constraints defined on reasoning step types. Finally, we present a set of experiments that validate our framework.

#### 2 RELATED WORK

To render models less verbose and more efficient, Train and Test Time Scaling approaches have been explored (Qu et al., 2025; Li et al., 2025; Chen et al., 2025a). Also, recent work has explored monitoring the generation of LRMs. The *Related-Work* section in the Appendix C complements this section, defining the inefficiency problem of reasoning models and its origins.

**Efficient Reasoning through Training.** Using SFT approaches, work such as Xia et al. (2025) explored fine-tuning models on compressed reasoning traces to limit the verbosity of LRM generation. Other papers have suggested various RL algorithms designed to make models more efficient. For instance, Luo et al. (2025); Team et al. (2025); Yu et al. (2025) showed that including a length component in the reward function leads to more efficient training and inference.

**Efficient Reasoning during Inference.** Researchers have also explored Inference Time Scaling technique to increase the efficiency of models (Qu et al., 2025). *Model Switch* uses a router module to select small or large models for inference depending on the complexity of the problem (Ong et al., 2025). Similarly, *System Switch* looked at dynamically selecting inference settings based on the problem (Aytes et al., 2025). *Length Budgeting* aims to reduce the budget allocated to the generation of answers. Works such as Lee et al. (2025); Han et al. (2025); Xu et al. (2025) showed that

careful prompt engineering can lead to more efficient generation compared to standard inference. In addition, Pu et al. (2025) demonstrated that calibration experiments can be performed to estimate the optimal number of tokens to solve particular problems. However, these techniques are hardly interpretable since they rely on either prompt engineering or black box techniques.

Monitoring LRM generation. We observe an emerging theme of research on monitoring LRM generation at a step level. Specifically, Lee & Hockenmaier (2025) proposes a taxonomy of reasoning traces evaluators. However, the authors acknowledged that existing monitoring approaches are not adapted to complex reasoning traces. Moreover, Zeng et al. (2025) showed that monitoring the generation of LRMs can enhance their performance by balancing both Exploration and Exploitation. But their technique does not monitor the reasoning traces at a step-level, and requires many inference trials. As a result, existing works often overlook the question of how to dynamically monitor LRMs reasoning during single inferences. To better understand how the generation reflects on the model's reasoning, we must begin by precisely defining what constitutes a *reasoning step*.

#### 3 How to define a reasoning step?

The concept of a reasoning step is central in evaluating and improving the generation of LRMs. However, defining a reasoning step remains a non-trivial problem. As highlighted by Yao et al. (2023); Lee & Hockenmaier (2025); Cao et al. (2025), the step segmentation depends on the models, the problem, and different research goals lead to various definitions. In this section, we survey existing approaches and select the one that leads to the most robust definition.

**Token-per-token generation.** From Schuurmans et al. (2024), we can formalize the auto-regressive generation of text of LLMs, and thereby LRMs. We assume  $x_{1:s}$  is an (1,s) dimensional vector containing the tokens of the input sequence, where each token  $x_i \in V = \{v_1, ..., v_V\}, |V|$  being the size of the vocabulary. We can approximate the next-token generation as following:

$$P_{\pi_{\theta}}(y|x_{1:s}) \approx \prod_{i=1}^{n} P_{\pi_{\theta}}(y_i|x_{1:s+i-1})$$
 (1)

where  $P_{\pi_{\theta}}(y|x_{1:s})$  is the probability of generating the output sequence  $y=y_{1:n}=x_{s+1:n+s}$  in an auto-regressive manner, and  $\pi_{\theta}$  is the model parametrized by  $\theta$ .

#### 3.1 What is a reasoning step?

Rather than viewing the model's output as a monolithic text sequence, recent work has shifted toward decomposing generation of LRMs into discrete steps. This decomposition enables finer-grained analysis of model behavior and facilitates targeted interventions. From the literature, we identified four principal methods to segment the output from a model into distinct thoughts (Appendix E):

- Token or sentence level: Näively, thoughts can be decomposed into token (Yao et al., 2023) or sentence level (Fu et al., 2023). However, for complex reasoning problems, these definitions are not ideal since reasoning steps are composed of multiple sentences in mathematical reasoning.
- Paragraph level: LLMs and LRMs such as Deepseek-R1, QwQ, or GPT are natively generating back-to-line symbols between two thoughts (e.g. .\n\n). Since this observation is model agnostic, it has been adopted by several works (Cao et al., 2025; Park et al., 2024; Lightman et al., 2023). However, Cao et al. (2025) emphasized that this approach on its own is not enough to correctly distinguish each steps. Models tend to output these symbols frequently, and so using them as delimiters for reasoning can result in over-estimating the real number of steps. Figure 8 in the Appendix E supports this observation, where more than 10 back-to-line symbols are generated for the same type of step.
- Dynamic steps using special token: Another common approach is to prompt the model to force the generation of special tokens to split the thoughts (e.g. <next\_step>). While some works have used this strategy (Zelikman et al., 2024; Sui et al., 2025b; Paul et al., 2024), it suffers from low reliability and efficiency. Indeed, this approach artificially generates more tokens, and prompt engineering could cause mistakes since models are not pre-trained to perform this sub-task.

To clearly identify and monitor reasoning steps, most approaches are insufficient to split the reasoning traces of models. This motivated us to find an alternative approach to segmenting the traces.

#### 3.2 Model agnostic reasoning step definition

Since most models separate paragraphs and thoughts using back-to-line symbols, using this token is a useful starting point to segment reasoning steps. However, to mitigate the over-segmentation problem caused when back-to-line symbols are generated too frequently, Cao et al. (2025) set a minimal number of tokens k per reasoning step and merge reasoning steps shorter than k together. We adopt a similar definition of a reasoning step: a step is delimited by ".\n\n", and we set a minimal number of tokens k on the size of each step to avoid considering many small steps.

Step Generation. As introduced in the previous section, we decompose the output sequence of LRMs into discrete reasoning steps. Building upon Cao et al. (2025) and the definition in Equation 1, we first formalize the notion of *stepwise generation*. Let  $y=y_{1:n}\in V^n$  be the output token sequence generated by the model over the vocabulary V. We define a reasoning delimiter token  $\alpha\in V$ , such as  $\alpha=\text{``.}\n^n$ . Let  $R=\{r_0=1,\ldots,r_i,\ldots,r_{T'}=n\}$  denote the indices in y corresponding to the occurrence of  $\alpha$  in y.  $r_0$  and  $r_{T'}$  correspond to the first and last indexes of  $y_{1:n}$ . Based on these indices, we define a sequence - of length T' - of reasoning steps formed by  $y_{1:n}$  with the delimiter  $\alpha$ :

$$S^* = \{s_1^*, \dots, s_i^*, \dots, s_{T'}^*\}, \text{ such as } s_i^* = y_{r_{i-1}:r_i}$$
 (2)

where each step  $s_i^*$  corresponds to a sub-part of the full output y. However, we observe that models tends to generate back-to-line symbols frequently. To reduce redundancy and noise from short or fragmented steps - highlighted by Cao et al. (2025) - we introduce a minimum token threshold  $k \in \mathbb{N}$  such as:

$$S = \{s_1, \dots, s_T\}, \quad \text{with } |s_j| \ge k \text{ for all } j \in [1, T].$$

For any original step  $s_i^* \in S^*$  such as  $|s_i^*| < k$ , we continue the generation until the merged span reaches another delimiter and its length exceeds the threshold k. In this formulation, each new reasoning step is initiated by the generation of  $\alpha$ , offering a more consistent definition between models. Algorithm 1 in Appendix F.1 formalizes our definition.

#### 4 STEP-TAGGING MODULE

In the previous section, we formalized the method we selected to segment reasoning steps of LRMs. Building on this, we introduce *Step-Tagging*, a lightweight module capable of identifying, discriminating, and tagging reasoning steps in real-time during inference.

**Objective.** Our definition of a reasoning step enables users to segment reasoning steps within model outputs. However, this definition alone does not allow the user to annotate the segmented steps with reasoning types. This annotation would enable users to track logical transitions within model outputs. To do this, we must first define a tag dictionary  $\mathcal{T}_{tags}$  (i.e., a label space of reasoning step tags) that covers the types of reasoning steps generated by models. Essentially, given a sequence of reasoning steps  $S = \{s_1, s_2, \ldots, s_T\}$ , we wish to label each step  $s_i$  with a tag  $\tau_i \in \mathcal{T}_{tags}$ . Formally, we are looking to construct a step-tagging function  $\phi$  such as:

$$\forall i \in [1, T], \phi(s_i) = \tau_i \tag{4}$$

where  $s_i \in S$  is a reasoning step from the full output sequence y, where  $|s_i| \geq k$ ,  $\phi$  is the step-tagging function, and  $\tau_i \in \mathcal{T}_{tags}$  is the reasoning tag associated to the step  $s_i$ .

**Taxonomy of the type of steps.** To enable fine-grained monitoring of reasoning behavior, we need to know the different types of reasoning steps that are typically generated by LRMs (i.e., we need to define  $\mathcal{T}_{tags}$ ). To do so, we created a taxonomy based on the outputs of both DeepSeek-R1-Distill-Llama-8B (DeepSeek-AI et al., 2025) and QwQ-32B (Team, 2025) models.

Inspired by prior work on model behavior analysis (Galichin et al., 2025; Kuznetsov et al., 2025), we first created a prompt to identify distinct types of reasoning steps in the traces (see Appendix I.1). We then sampled 40 reasoning traces from the MATH500 train dataset (covering two samples per difficulty level for each model) and using our prompt submitted the traces to GPT-40-mini (OpenAI et al., 2024). The prompt resulted in a series of different step-types. We merged overlapping categories, to construct a taxonomy that reflects the temporal and reasoning progression of model's traces. We refer to this taxonomy as *ReasonType* (Figure 2) encompassing 13 categories, including early-stage behaviors such as *Problem Re-statement*, later reasoning stages like *Verification* and *Exploration*. To validate our taxonomy, we conducted ablation studies (see Appendix H).

Figure 2: ReasonType - A taxonomy of reasoning step types as per gpt-4o-mini

Early-stopping criteria. In the following section, we see that LRMs tend to generate the answer early in the output sequence, with step-types following an ordered pattern. Based on this observation, the central challenge that we address is to determine when to stop the generation of LRMs based on step tags, creating an interpretable stopping criterion. Assuming that our Step-Tagging framework can effectively monitor the steps (Equation 4), we can define a constraint on the frequency of a given step type. Each constraint operates online, over a running sequence of reasoning steps  $S_{\text{running}} = \{s_1, \ldots, s_j\}$ , where each step  $s_i$  is associated with a tag  $\tau_i \in \mathcal{T}$ . We define the constraint  $c_{\tau^*}$  as:

$$c_{\tau^*}(S_{\text{running}}, \delta) = \mathbf{1}[f_{\text{freq}}(S_{\text{running}}, \tau^*) \le \delta] \text{ with } f_{\text{freq}}(S_{\text{running}}, \tau^*) = \sum_{i=1}^{J} \mathbf{1}[\tau_i = \tau^*]$$
 (5)

where  $c_{\tau^*}(S_{\text{running}}, \delta)$  is the constraint on the tag type  $\tau^*$  over the step-sequence  $S_{\text{running}}$  being generated, given the threshold  $\delta$ .  $f_{\text{freq}}(S_{\text{running}}, \tau^*)$  is the occurrence of the type-step  $\tau^*$  over the running sequence  $S_{\text{running}}$ . While the constraint  $c_{\tau^*}$  is satisfied, the generation continues. If the constraint is violated, the generation stops (see Appendix F.2 for more implementation details).

To facilitate the evaluation of early-exit answers, we prompted the models right after the last step being generated, and allowed an additional budget of 100 tokens. We used the following prompt: " $\n \n I$  am confident in my answer. Here is the final answer. $\n \n **Final Answer**$ ". We borrowed this approach from Muennighoff et al. (2025), who showed that this intervention helped the model to provide explicitly its current best answer - thereby facilitating evaluation.

#### 5 EXPERIMENTAL SETTING

Our paper contains two objectives. First, our goal is to prove that lightweight classifiers can effectively monitor the generation of LRMs. Furthermore, we show that the Step-Tagging framework can be used to implement an interpretable early-stopping criterion to make the generation of LRMs more efficient. We will first motivate our choices of datasets and inference settings followed by the step-tagging pipeline and the choice of metrics to measure the performance of this pipeline.

**Datasets.** To assess our approach, we selected two state-of-the-art reasoning datasets:

- MATH500 (Hendrycks et al., 2021): This dataset includes 12,500 mathematical questions spanning 5 different levels of complexity, allowing diversity in analysis of efficiency of reasoning behaviors. We selected the curated version from Lightman et al. (2023), containing 500 test samples and selected 1,000 training samples to form an equivalent distribution of complexity level.
- **GSM8K** (Cobbe et al., 2021): This dataset contains 8,792 mathematical questions. We selected 3,000 train instances, and the original 1,318 test samples. Overall, this dataset is more homogeneous where questions involve logical mathematical reasoning, and include a larger number of questions which is good for benchmarking models.

**Model selection.** To apply our framework a user must have access to the fine-grained reasoning traces of LRMs. However, many high-performing closed-source models (such as, o3 and Claude 3.7) do not expose raw reasoning traces. Instead, these models output summaries of thinking tokens generated, which can bias the estimation of their efficiency compared to open-source models. In contrast, open-source models like DeepSeek-R1 and QwQ consistently provide reasoning traces. For this reason, we focus our analysis exclusively on *DeepSeek-R1-Distill-Llama-8B*, *DeepSeek-R1-Distill-Qwen-14B* and *QwQ-32B*, which offer the granularity needed to monitor the reasoning process. This choice is motivated by their variety in term of size and performance, full open-source availability, and diversity in providers.

 **Inference setting.** To monitor the steps and intervene in the generation process, we suggest a new definition of the generation process of LRMs. We assume that each model generates one token at a time, and we split the steps dynamically. However, for the purposes of our experiments instead of re-designing the generation process, we performed standard inference and applied our *Step-Tagging* and *Early-Stopping* algorithms *offline*. To ensure the robustness and reproducibility of our approach, we generated five outputs per test sample using fixed random seeds (namely 40, 41, 42, 43, and 44), with deterministic decoding.

**Metrics.** To assess the model's performance on challenging reasoning tasks, the Avg@k, Pass@k, and Cons@k are common metrics (Chen et al., 2021; 2025a; Yu et al., 2025). The Pass@k measures the proportion of the samples where at least one of k attempts leads to the correct answer, while the Cons@k consider a sample correct if all k attempts are correct. Since we are interested about both performance and robustness of our approach, we selected the Avg@5, the Pass@5 and the Cons@5 as the quantitative metrics. Assessing the performance of LRMs on mathematical questions is challenging. This is due to the open nature of the question. For our experiments, we selected the  $Math-Verify^1$  library which is a common metric to assess mathematical problems. It uses text extraction and formal verification. This metric also reported strong correctness compared to other evaluation methods such as Harness (Zhibin Gou, 2024) or Qwen-Math Verifier (Huang et al., 2025).

**Baselines.** To assess the effectiveness of our early-stopping approach, we define two baselines:

- Ideal Early stopping  $\mathcal{IES}$ : We observe a growing understanding that, up to a token-budget, thinking longer may be leading to worse results. Muennighoff et al. (2025) observes that certain models achieved correct answers at the beginning, but sometimes backtracked to a wrong answer. Inspired by this work, we define the *Ideal Early Stopping*, which prunes the remaining steps after the first occurrence of the correct answer based on our metric if any. In this case, this baseline is theoretical since the ground truth label is needed for each inference (see Appendix J).
- **Prompt-guided efficiency**  $\mathcal{P}_{guided}$ : We also observe that LRMs are sensitive to the input prompt (Lee et al., 2025). In this case, we compare our framework with user-prompt and system-prompt variants, with Zero-Shot and Few-Shot prompts that aim to reduce the reasoning computation while retaining accuracy. We explicitly instructed the models to not generate verbose output, or over-verification steps. We selected 4 variants, namely: zero-shot user and system prompt, and few-shot system prompt with 1 and 3 examples:  $\mathcal{P}_{user}^{(0)}$ ,  $\mathcal{P}_{system}^{(0)}$ ,  $\mathcal{P}_{system}^{(1)}$ , respectively. The prompts used to establish these baselines are listed in Appendix J.

#### 5.1 IMPLEMENTATION OF THE STEP-TAGGER MODULE

Training data generation. Given that our reasoning step taxonomy was created using GPT-40-mini OpenAI et al. (2024) the most direct way to label a reasoning trace would be to use GPT-40-mini. However, this GPT-40-mini annotation is costly, each step requiring more than a second to be annotated (see Table 6 in Appendix I.2). Consequently, instead, we used GPT-40-mini to label a dataset of reasoning traces with the labels from the taxonomy that we use to train lighter weight reasoning step classifiers. We constructed training datasets by running each LRMs on 1,000 samples from MATH500 train and 3,000 samples from GSM8K train datasets (with a seed of 42). For each step  $s_i$  in generated outputs, we prompted GPT-40-mini to assign a tag  $\tau_i$  (Appendix I.1).

Sentence classifiers. We selected the bert-base-uncased sentence classifier (Devlin et al., 2019) to construct our Step-Tagging framework, including a single hidden layer. Given the large and fine-grained nature of our taxonomy (13 distinct step types), training a multi-class classifier is challenging due to significant class imbalance. To address this, we trained separate binary classifiers for each step-type. This approach notably improved detection accuracy across low-frequency categories, and fits our definition of early-stopping constraint: one step-type per early-stopping criteria. We used a *balanced cross-entropy* to enhance the performance of the models on low-represented classes. We implemented an early-stopping criteria, and a maximum of 5 epochs. The batch size is 16 and we used an AdamW optimizer with a learning rate of  $2.10^{-5}$ . To evaluate the performance of our classifiers, we computed the Macro-F1 and Micro-F1 on the test datasets. While the *Macro-F1* helps to identify the classifier's ability to detect rare classes, the *Micro-F1* offers a more global view on the step detector's performance across all steps.

<sup>1</sup>https://github.com/huggingface/Math-Verify

#### 5.2 Step-split settings

Minimal step size k. To apply our definition of reasoning steps, users first need to set the value of k. From the literature, this task is not straightforward since Cao et al. (2025) uses different values of k based on the problems and models (e.g.  $k \approx 100$ s). The value of k directly affects the granularity of the monitoring. A *small value* would imply very small steps, sometimes splitting the same thoughts between many steps, leading to extremely frequent monitoring. Conversely, a *large value* of k would imply significantly large steps, including multiple thoughts and hence leading to biased monitoring.

Selecting the optimal k value. To balance fine-grained monitoring and reasoning efficiency, we rely on two proxies to determine the value of k. First, the  $\mathcal{IES}$  accuracy provide us a way to assess if individual steps contains more than one thought. Second, the Step-Tagger performance informs us about the semantic meaning of steps under a given a value of k. Based on three ablation studies that we conducted in Appendix G, we set k to 60, 30 and 100 for DS-Llama8B, DS-Qwen14B and QwQ-32B, respectively.

#### 5.3 EARLY-STOPPING CONSTRAINTS

Early-Stopping calibration using a Pareto curve. To select the correct constraints (tag-type  $\tau$  and threshold  $\delta$ ) we rely on the training datasets, and on the synthetic generated tags. Figure 3 presents the number of tokens vs. accuracy of every tag-type with values of threshold ranging from 0 to 20, for the DS-Llama8B model on our train MATH500 per complexity level. We first observe that LRMs tend to generate an increasing number of tokens when the complexity of the prompt increases. Figure 3 also shows that early-stopping constraints are dependent on the complexity. For this reason, we selected one constraint per complexity level for the MATH500 dataset, while for the GSM8K dataset, we chose a unique constraint since we assume problems to carry equivalent complexity. Furthermore, we observe that constraints form a *Pareto curve* (Lee et al., 2025), and each step-types results in different trade-off between accuracy and token-count (see Appendix H.3). On the strength of this observation, we set three Early-Stopping criteria that target specific trade-offs between accuracy and efficiency: ST-ES 95%, ST-ES 90% and ST-ES 85%. For each setting, we select the tag-type and threshold that lies closest to the Pareto frontier. Selected constraints and calibrations for the other models and datasets are shown in the Appendix K.

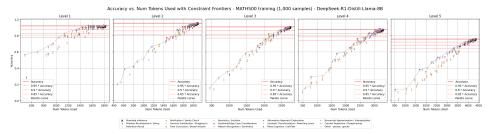


Figure 3: Early-Stopping selection using Pareto - DS-Llama8B on MATH500 train - seed 42

**LRM Router: dynamic inference.** The MATH500 dataset involves various complexity levels. For efficient inference, dynamic routing is needed. However, in real-world settings, the complexity level of a question is sometimes unknown. To address this challenge, we grouped levels  $\{1,2\}$  and  $\{3,4,5\}$  into two complexity levels, and identified common constraints for both clusters. To route the inference settings, we trained a BERT classifier on the input problem, using the full MATH500 train dataset. Performance of this classifier is reported in Table 8, Appendix L. For the MATH500 dataset, we included a fourth Early-Stopping criteria, namely *ST-ES Router*.

#### 6 MONITORING LRMs USING STEP-TAGGER

To validate our taxonomy, we analyse the distribution and sequence of the step-tag labels generated by GPT-40-mini, then we evaluate the performance of our sentence classifiers.

**Reasoning patterns.** First, we observe that our Step-Tagging framework allows us to clearly follow the reasoning progression of the model. Figures 34 and 35 in the Appendix I.2 present an analysis (and validation) of reasoning patterns exhibited by models based on the step-types identified in the reasoning traces generated by the models.

Step frequency. Figure 4 presents the frequency of each step-type in the GPT-40-mini labels. The plot shows a high frequency of Verification, confirming our observations from the lit-We also note that the frequency seems to depend on the problem complexity and models. For DS-Llama8B and DS-Qwen14B, Formula Substitution steps are very occurrent for GSM8K ( $\approx$ 60%), while Exploration and Self-Talk steps are more frequent for QwQ-32B on both datasets.

378

379

380

381

382

383

384

385

386

387

388

389

390

391

392

393

394

395

396

397

398

399

400

401

402

403

404

405 406 407

408 409

410

411

412

413

414

415

416

417

418 419

420

421

422

423

424

425

426

427 428

429

430

431

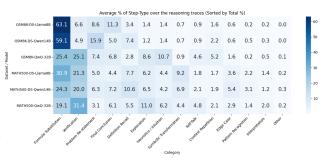


Figure 4: Step-type distribution from GPT-40-mini

**Performance of step monitoring.** Figure 5 presents the performance of the binary step-classifiers on the selected step-types constraints for the DS-Llama8B model. We observe that the Micro-F1 is generally high across most steps for all models across all datasets - ranging from 0.89 to 0.97, which demonstrates that the classifiers are good at detecting step-tags. Moreover, we also reported the macro-F1 score since the distribution of step-types is highly imbalanced (see Figure 4).

We observe lower scores, notably for Context Repetition with 0.65 (Context Repetition is a rare step type, representing 1.7% of the labels, and so we attribute this relatively low score to label imbalance). However, the scores remain relatively high, particularly for Verification and Exploration. Figure 5: Step-Tagger performance - DS-Llama8B



We interpret the strong performance of the classifiers as validating our reasoning step taxonomy in the sense that it indicates that the step types are distinct (i.e., they reflect types with separable properties). Figures 42 and 43 in the Appendix M present similar results for DS-Qwen14B and *OwO-32B*, respectively.

#### STEP-TAGGING EARLY-STOPPING (ST-ES) CRITERIA

Next, we show in this section that Step-Tagging modules can effectively be used as an early-stopping criteria. Figure 6 presents the average token count against the Avg@5 for the three LRMs on the MATH500 and GSM8K datasets. Each plot compares the performance trade-offs between the baselines and the ST-ES criteria. Table 9 in the Appendix O reports the quantitative metrics of the baselines and our approach on the three models, for the 5 seeds that we selected.

 $\mathcal{P}_{\text{guided}}$  baselines. We first notice that simple instruction on the models results in strong tokenreduction, achieving 20% to 60% saved tokens across configurations. Specifically, it seems that the baselines are giving much better results on QwQ-32B, and the system-prompt variants generally lead to more token-reduction for the Deepseek models.

Strong performance of the ST-ES. Next, we observe that our ST-ES criteria effectively leads to more efficient generation, with all ES-ST settings lying on the left side of the Efficiency line compared to the Standard inference for all models. Furthemore, the ST-ES criteria appear to outperform most  $\mathcal{P}_{guided}$  baselines for both Deepseek models.

Indeed, we observe that our ST-ES criteria is performing well on the DS-Llama8B model on both datasets since almost all ST-ES configurations lies on the Pareto front. On MATH500 (Figure 6(a)), ST-ES Router and ST-ES 85% achieved approximately the same token reduction as  $\mathcal{P}_{\text{system}}^{(0)}$  and  $\mathcal{P}_{system}^{(3)}$  (27% and 34%, respectively), while achieving higher accuracy. On GSM8K (Figure 6(d)), ST-ES 90% achieves the same token reduction as  $\mathcal{P}_{system}^{(1)}$  (around 41%) while maintaining higher Avg@5 (0.799 vs. 0.754, respectively). Furthermore, results on the DS-Qwen14B model also show good performance of criteria leading to significant token-reduction (10 to 32%), with some configurations lying on the Pareto front. However, the  $\mathcal{P}_{guided}$  settings appears to lead to more efficient inference, notably for MATH500. In addition, the criteria suffers from more accuracy loss, as the ST-ES Line are more vertical than for DS-Llama8B.

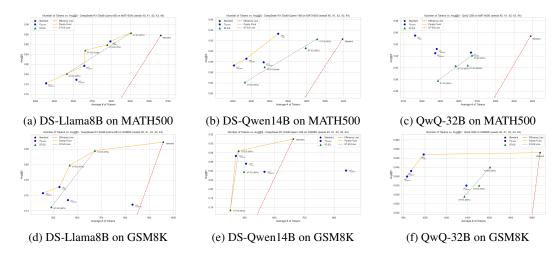


Figure 6: Number of Tokens vs. Avg@5 -  $\mathcal{P}_{guided}$  Baselines vs. ST-ES criteria - The efficiency lines in red highlight the configurations that improve the efficiency relative to the standard inference, while the Pareto frontiers in yellow show the most efficient approaches. The Step-Tagging Early-Stopping framework achieved up to 30 to 40% of token-count saving, with minimal accuracy loss.

ST-ES faces challenges on the QwQ model. In contrast, the ST-ES criteria shows nuanced results on the QwQ-32B model. Baselines are stronger, and for the same token gains as the Deepseek models, the accuracy loss seems higher. We suspect that this observation can be attributed to two factors.

First, Figure 7 presents the average per percentage of the full output sequence for the three models. DS-Llama8B and DS-Qwen14B appear to generate correct answers earlier in their output sequences, but sometimes continue reasoning, leading to the destruction of the correct current answer. In particular, this is the case on easier problems (Level 1-3 MATH500) and GSM8K, where a drop in the accuracy can be observed at around 40-50% stopping. By contrast, QwQ-32B exhibits more stable accuracy gains as the token count increases, which suggests that the model is more conservative of its current solution in the way it constructs its response.

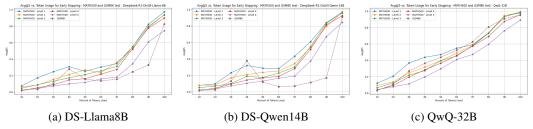


Figure 7: Early-Stopping Avg@5 per percentage of the full output sequence

Second, we suspect that larger models are better at controlling the length of their generation using specific prompts. We observe that the baselines from QwQ-32B are much more efficient than the ones from both DS-Llama8B and DS-Qwen14B, which are smaller models. We note that Lee et al. (2025) report similar findings with larger models achieving higher *Upper Bound of Token Reduction* when prompted to compress their reasoning.

#### 8 Conclusion

This work offers a novel view on both monitoring and efficiency of LRMs. We propose *ReasonType*, a novel taxonomy of reasoning steps, and demonstrated that users can effectively track the reasoning flow of the generation. We validated our taxonomy using the performance of sentence classifiers, paving the way for more work on the monitoring of reasoning steps.

Furthermore, we show that the frequency of step-type in the reasoning traces can be used as a reliable and interpretable early-stopping criterion. Through careful monitoring of certain step-types, our framework can enhance the control of the generation of RLMs enabling a significant reduction in token generation (up to 40%) while preserving performance.

#### REPRODUCIBILITY STATEMENT

We took several measures to ensure the reproducibility of our experiments, namely:

- Code availability: The source code that we developed to conduct our experiments is available in the submission ZIP folder.
- Experimental Settings: We listed in Section 5 the experimental settings. This includes the datasets used, the models (open-source available on HuggingFace), the parameters of the algorithms, the prompts of the models, the evaluation functions, and the environment setups (seeds and deterministic decoding). We also included scripts to reproduce the experiments we lead. We used one or two A100-80GB GPUs to run our experiments.

#### REFERENCES

- Ibrahim Abdelaziz, Kinjal Basu, Mayank Agarwal, Sadhana Kumaravel, Matthew Stallone, Rameswar Panda, Yara Rizk, GP Bhargav, Maxwell Crouse, Chulaka Gunasekara, Shajith Ikbal, Sachin Joshi, Hima Karanam, Vineet Kumar, Asim Munawar, Sumit Neelam, Dinesh Raghu, Udit Sharma, Adriana Meza Soria, Dheeraj Sreedhar, Praveen Venkateswaran, Merve Unuvar, David Cox, Salim Roukos, Luis Lastras, and Pavan Kapanipathi. Granite-function calling model: Introducing function calling abilities via multi-task learning of granular tasks, 2024. URL https://arxiv.org/abs/2407.00121.
- Simon A. Aytes, Jinheon Baek, and Sung Ju Hwang. Sketch-of-thought: Efficient llm reasoning with adaptive cognitive-inspired sketching, 2025. URL https://arxiv.org/abs/2503.05179.
- Lang Cao, Chao Peng, Renhong Chen, Wu Ning, Yingtian Zou, and Yitong Li. Step guided reasoning: Improving mathematical reasoning using guidance generation and step reasoning, 2025. URL https://arxiv.org/abs/2410.19817.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. Evaluating large language models trained on code, 2021. URL https://arxiv.org/abs/2107.03374.
- Qiguang Chen, Libo Qin, Jinhao Liu, Dengyun Peng, Jiannan Guan, Peng Wang, Mengkang Hu, Yuhang Zhou, Te Gao, and Wanxiang Che. Towards reasoning era: A survey of long chain-of-thought for reasoning large language models, 2025a. URL https://arxiv.org/abs/2503.09567.
- Xingyu Chen, Jiahao Xu, Tian Liang, Zhiwei He, Jianhui Pang, Dian Yu, Linfeng Song, Qiuzhi Liu, Mengfei Zhou, Zhuosheng Zhang, Rui Wang, Zhaopeng Tu, Haitao Mi, and Dong Yu. Do not think that much for 2+3=? on the overthinking of o1-like llms, 2025b. URL https://arxiv.org/abs/2412.21187.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems, 2021. URL https://arxiv.org/abs/2110.14168.
- DeepSeek-AI, Daya Guo, Dejian Yang, and Haowei Zhang et al. Deepseek-r1: Incentivizing reasoning capability in Ilms via reinforcement learning, 2025. URL https://arxiv.org/abs/2501.12948.

- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding, 2019. URL https://arxiv.org/abs/1810.04805.
  - Yao Fu, Hao Peng, Ashish Sabharwal, Peter Clark, and Tushar Khot. Complexity-based prompting for multi-step reasoning, 2023. URL https://arxiv.org/abs/2210.00720.
  - Andrey Galichin, Alexey Dontsov, Polina Druzhinina, Anton Razzhigaev, Oleg Y. Rogov, Elena Tutubalina, and Ivan Oseledets. I have covered all the bases here: Interpreting reasoning features in large language models via sparse autoencoders, 2025. URL https://arxiv.org/abs/2503.18878.
  - Tingxu Han, Zhenting Wang, Chunrong Fang, Shiyu Zhao, Shiqing Ma, and Zhenyu Chen. Tokenbudget-aware llm reasoning, 2025. URL https://arxiv.org/abs/2412.18547.
  - Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset, 2021. URL https://arxiv.org/abs/2103.03874.
  - Arian Hosseini, Alessandro Sordoni, Daniel Toyama, Aaron Courville, and Rishabh Agarwal. Not all llm reasoners are created equal, 2024. URL https://arxiv.org/abs/2410.01748.
  - Yuzhen Huang, Weihao Zeng, Xingshan Zeng, Qi Zhu, and Junxian He. Pitfalls of rule- and model-based verifiers a case study on mathematical reasoning, 2025. URL https://arxiv.org/abs/2505.22203.
  - Kristian Kuznetsov, Laida Kushnareva, Polina Druzhinina, Anton Razzhigaev, Anastasia Voznyuk, Irina Piontkovskaya, Evgeny Burnaev, and Serguei Barannikov. Feature-level insights into artificial text detection with sparse autoencoders, 2025. URL https://arxiv.org/abs/2503.03601.
  - Ayeong Lee, Ethan Che, and Tianyi Peng. How well do llms compress their own chain-of-thought? a token complexity approach, 2025. URL https://arxiv.org/abs/2503.01141.
  - Jinu Lee and Julia Hockenmaier. Evaluating step-by-step reasoning traces: A survey, 2025. URL https://arxiv.org/abs/2502.12289v1.
  - Zhong-Zhi Li, Duzhen Zhang, Ming-Liang Zhang, Jiaxin Zhang, Zengyan Liu, Yuxuan Yao, Haotian Xu, Junhao Zheng, Pei-Jie Wang, Xiuyi Chen, Yingying Zhang, Fei Yin, Jiahua Dong, Zhiwei Li, Bao-Long Bi, Ling-Rui Mei, Junfeng Fang, Xiao Liang, Zhijiang Guo, Le Song, and Cheng-Lin Liu. From system 1 to system 2: A survey of reasoning large language models, 2025. URL https://arxiv.org/abs/2502.17419.
  - Jonathan Light, Wei Cheng, Benjamin Riviere, Wu Yue, Masafumi Oyamada, Mengdi Wang, Yisong Yue, Santiago Paternain, and Haifeng Chen. Disc: Disc: Dynamic decomposition improves llm inference scaling, 2025. URL https://arxiv.org/abs/2502.16706.
  - Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step, 2023. URL https://arxiv.org/abs/2305.20050.
  - Haotian Luo, Li Shen, Haiying He, Yibo Wang, Shiwei Liu, Wei Li, Naiqiang Tan, Xiaochun Cao, and Dacheng Tao. O1-pruner: Length-harmonizing fine-tuning for o1-like reasoning pruning, 2025. URL https://arxiv.org/abs/2501.12570.
  - Leena Mathur, Marian Qian, Paul Pu Liang, and Louis-Philippe Morency. Social genome: Grounded social reasoning abilities of multimodal models, 2025. URL https://arxiv.org/abs/2502.15109.
  - Mayank Mishra, Matt Stallone, Gaoyuan Zhang, Yikang Shen, Aditya Prasad, Adriana Meza Soria, Michele Merler, Parameswaran Selvam, Saptha Surendran, Shivdeep Singh, Manish Sethi, Xuan-Hong Dang, Pengyuan Li, Kun-Lung Wu, Syed Zawad, Andrew Coleman, Matthew White, Mark Lewis, Raju Pavuluri, Yan Koyfman, Boris Lublinsky, Maximilien de Bayser, Ibrahim

Abdelaziz, Kinjal Basu, Mayank Agarwal, Yi Zhou, Chris Johnson, Aanchal Goyal, Hima Patel, Yousaf Shah, Petros Zerfos, Heiko Ludwig, Asim Munawar, Maxwell Crouse, Pavan Kapanipathi, Shweta Salaria, Bob Calio, Sophia Wen, Seetharami Seelam, Brian Belgodere, Carlos Fonseca, Amith Singhee, Nirmit Desai, David D. Cox, Ruchir Puri, and Rameswar Panda. Granite code models: A family of open foundation models for code intelligence, 2024. URL https://arxiv.org/abs/2405.04324.

- Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke Zettlemoyer, Percy Liang, Emmanuel Candès, and Tatsunori Hashimoto. s1: Simple test-time scaling, 2025. URL https://arxiv.org/abs/2501.19393.
- Tergel Munkhbat, Namgyu Ho, Seo Hyun Kim, Yongjin Yang, Yujin Kim, and Se-Young Yun. Self-training elicits concise reasoning in large language models, 2025. URL https://arxiv.org/abs/2502.20122.
- Isaac Ong, Amjad Almahairi, Vincent Wu, Wei-Lin Chiang, Tianhao Wu, Joseph E. Gonzalez, M Waleed Kadous, and Ion Stoica. Routellm: Learning to route llms with preference data, 2025. URL https://arxiv.org/abs/2406.18665.
- OpenAI, Josh Achiam, Steven Adler, and Sandhini Agarwal et al. Gpt-4 technical report, 2024. URL https://arxiv.org/abs/2303.08774.
- Yoonjeong Park, Hyunjin Kim, Chanyeol Choi, Junseong Kim, and Jy-Yong Sohn. Can separators improve chain-of-thought prompting? In 2024 2nd International Conference on Foundation and Large Language Models (FLLM), pp. 493–500. IEEE, November 2024. doi: 10.1109/fllm63129. 2024.10852507. URL http://dx.doi.org/10.1109/fllm63129.2024.10852507.
- Shishir G. Patil, Tianjun Zhang, Xin Wang, and Joseph E. Gonzalez. Gorilla: Large language model connected with massive apis, 2023. URL https://arxiv.org/abs/2305.15334.
- Debjit Paul, Mete Ismayilzada, Maxime Peyrard, Beatriz Borges, Antoine Bosselut, Robert West, and Boi Faltings. Refiner: Reasoning feedback on intermediate representations, 2024. URL https://arxiv.org/abs/2304.01904.
- Xiao Pu, Michael Saxon, Wenyue Hua, and William Yang Wang. Thoughtterminator: Benchmarking, calibrating, and mitigating overthinking in reasoning models, 2025. URL https://arxiv.org/abs/2504.13367.
- Xiaoye Qu, Yafu Li, Zhaochen Su, Weigao Sun, Jianhao Yan, Dongrui Liu, Ganqu Cui, Daizong Liu, Shuxian Liang, Junxian He, Peng Li, Wei Wei, Jing Shao, Chaochao Lu, Yue Zhang, Xian-Sheng Hua, Bowen Zhou, and Yu Cheng. A survey of efficient reasoning for large reasoning models: Language, multimodality, and beyond, 2025. URL https://arxiv.org/abs/2503.21614.
- Sebastian Raschka. Understanding reasoning llms: Methods and strategies for building and refining reasoning models, 2025. URL https://magazine.sebastianraschka.com/p/understanding-reasoning-llms.
- Baptiste Rozière, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Romain Sauvestre, Tal Remez, Jérémy Rapin, Artyom Kozhevnikov, Ivan Evtimov, Joanna Bitton, Manish Bhatt, Cristian Canton Ferrer, Aaron Grattafiori, Wenhan Xiong, Alexandre Défossez, Jade Copet, Faisal Azhar, Hugo Touvron, Louis Martin, Nicolas Usunier, Thomas Scialom, and Gabriel Synnaeve. Code llama: Open foundation models for code, 2024. URL https://arxiv.org/abs/2308.12950.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms, 2017. URL https://arxiv.org/abs/1707.06347.
- Dale Schuurmans, Hanjun Dai, and Francesco Zanini. Autoregressive large language models are computationally universal, 2024. URL https://arxiv.org/abs/2410.03170.
- Kathrin Seßler, Yao Rong, Emek Gözlüklü, and Enkelejda Kasneci. Benchmarking large language models for math reasoning tasks, 2024. URL https://arxiv.org/abs/2408.10839.

- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. Deepseekmath: Pushing the limits of mathematical reasoning in open language models, 2024. URL https://arxiv.org/abs/2402.03300.
  - Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling Ilm test-time compute optimally can be more effective than scaling model parameters, 2024. URL https://arxiv.org/abs/2408.03314.
  - DiJia Su, Hanlin Zhu, Yingchen Xu, Jiantao Jiao, Yuandong Tian, and Qinqing Zheng. Token assorted: Mixing latent and text tokens for improved language model reasoning, 2025. URL https://arxiv.org/abs/2502.03275.
  - Yang Sui, Yu-Neng Chuang, Guanchu Wang, Jiamu Zhang, Tianyi Zhang, Jiayi Yuan, Hongyi Liu, Andrew Wen, Shaochen Zhong, Hanjie Chen, and Xia Hu. Stop overthinking: A survey on efficient reasoning for large language models, 2025a. URL https://arxiv.org/abs/2503.16419.
  - Yuan Sui, Yufei He, Tri Cao, Simeng Han, Yulin Chen, and Bryan Hooi. Meta-reasoner: Dynamic guidance for optimized inference-time reasoning in large language models, 2025b. URL https://arxiv.org/abs/2502.19918.
  - Kimi Team, Angang Du, Bofei Gao, and Bowei Xing et al. Kimi k1.5: Scaling reinforcement learning with llms, 2025. URL https://arxiv.org/abs/2501.12599.
  - Qwen Team. Qwq-32b: Embracing the power of reinforcement learning, March 2025. URL https://qwenlm.github.io/blog/qwq-32b/.
  - Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation language models, 2023. URL https://arxiv.org/abs/2302.13971.
  - Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models, 2023. URL https://arxiv.org/abs/2203.11171.
  - Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models, 2023. URL https://arxiv.org/abs/2201.11903.
  - Sean Welleck, Ximing Lu, Peter West, Faeze Brahman, Tianxiao Shen, Daniel Khashabi, and Yejin Choi. Generating sequences by learning to self-correct, 2022. URL https://arxiv.org/abs/2211.00053.
  - Sean Williams and James Huckle. Easy problems that Ilms get wrong, 2024. URL https://arxiv.org/abs/2405.19616.
  - Heming Xia, Chak Tou Leong, Wenjie Wang, Yongqi Li, and Wenjie Li. Tokenskip: Controllable chain-of-thought compression in llms, 2025. URL https://arxiv.org/abs/2502.12067.
  - Silei Xu, Wenhao Xie, Lingxiao Zhao, and Pengcheng He. Chain of draft: Thinking faster by writing less, 2025. URL https://arxiv.org/abs/2502.18600.
  - Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik Narasimhan. Tree of thoughts: Deliberate problem solving with large language models, 2023. URL https://arxiv.org/abs/2305.10601.
    - Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, Xin Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Guangming Sheng, Yuxuan Tong, Chi Zhang, Mofan Zhang, Wang Zhang, Hang Zhu, Jinhua Zhu, Jiaze Chen, Jiangjie Chen, Chengyi Wang, Hongli Yu, Yuxuan Song, Xiangpeng Wei, Hao

Zhou, Jingjing Liu, Wei-Ying Ma, Ya-Qin Zhang, Lin Yan, Mu Qiao, Yonghui Wu, and Mingx-uan Wang. Dapo: An open-source llm reinforcement learning system at scale, 2025. URL https://arxiv.org/abs/2503.14476. Eric Zelikman, Georges Harik, Yijia Shao, Varuna Jayasiri, Nick Haber, and Noah D. Goodman. Quiet-star: Language models can teach themselves to think before speaking, 2024. URL https: //arxiv.org/abs/2403.09629. Weihao Zeng, Yuzhen Huang, Lulu Zhao, Yijun Wang, Zifei Shan, and Junxian He. B-star: Monitoring and balancing exploration and exploitation in self-taught reasoners, 2025. URL https://arxiv.org/abs/2412.17256. Yue Zhang Zhibin Gou. Llm math evaluation harness, 04 2024. URL https://github.com/ ZubinGou/math-evaluation-harness. 

# **Appendix**

## **Table of Contents**

~~-	01 0011001100		
A	LLM Usage	16	
В	Limitations and Future Work	16	
C	Related-Work	16	
	C.1 Language Reasoning Models (LRMs)	16	
	C.2 LRMs are Inefficient	16	
	C.3 Open-source reasoning path	17	
D	Mathematical reasoning datasets	17	
E	Definition of reasoning steps	18	
F	Algorithms	19	
	F.1 Step-wise Generation	19	
	F.2 Early Stopping algorithm	20	
G	Selecting the minimal number of token $k$	21	
	G.1 Ideal Early-Stopping as a proxy for the step-size	21	
	G.2 Semantic properties implied by the parameter $k \dots \dots \dots \dots$	23	
	G.3 Influence of the parameter $k$ on the Step-Tagging Early-Stopping criteria	25	
	G.4 Summary of takeaways	27	
Н	Validation of the ReasonType taxonomy	28	
	H.1 Reason-Type taxonomy for identifying reasoning behaviors	28	
	H.2 Tracking step-types to design interpretable early-stopping criteria	29	
	H.3 Certain Step-Type are better early-stopping criteria	31	
	H.4 Comparison to alternative taxonomy	33	
	H.5 Summary of takeaways	34	
I	Analysis of reasoning step types	35	
	I.1 Prompts	35	
	I.2 Statistics on the reasoning steps count and types	35	
J	Baselines	39	
K	Calibration of the ST-ES criteria	40	
L	LLM-Router: Prompt complexity classification	41	
M	Step-Tagging performance	41	
N	Ideal-Early-Stopping	42	
o			

#### A LLM USAGE

We acknowledge the use of Large Language Models for the purpose of our experimentation in our paper. Specifically, as stated in Section 4, we relied on GPT-40-mini to set our *ReasonType* taxonomy. This approach is borrowed from work on behavior analysis of LLMs, such as Galichin et al. (2025); Kuznetsov et al. (2025).

#### B LIMITATIONS AND FUTURE WORK

Our definition of reasoning step is taken from previous work, and relies on empirical evaluation. We believe that our step taxonomy can enhance the step definition. Future work should look at leveraging the performance of step-classification to better define reasoning steps.

To train accurate Step-Tagger modules, we suspect that significantly increasing the number of traces could lead to better results. Also, down-sampling could render our training more effective, and increase the Macro-F1. In addition, a better definition of a step could lead to more effective monitoring. For instance, it would be interesting to explore dynamic values of minimal number of tokens k, making our approach even more agnostic.

Further, our ST-ES criteria is näive, since it is simply based on frequency monitoring. Some works have started to explore confidence-based methods on the uncertainty of the logits. We believe that integrating this knowledge to our criteria could further enhance our early-stopping criteria.

#### C RELATED-WORK

#### C.1 LANGUAGE REASONING MODELS (LRMS)

The field of reasoning models has been very active over the recent years. The literature tends to divide existing models into two distinct categories, namely *System-1* and *System-2* (Li et al., 2025; Qu et al., 2025). System-1 models refer to intuitive and fast LLMs. These correspond to standard instructed LLMs. In contrast, System-2 models are slower and deeper thinking, designed to perform explicit multi-step thought, and are referred to as LRMs.

**Building System-2 models.** Driven by substantial research efforts, LLMs now excel at standard capabilities such as Natural Language Processing (NLP) (Touvron et al., 2023; OpenAI et al., 2024), code generation (Mishra et al., 2024; Rozière et al., 2024) and Function-calling (Abdelaziz et al., 2024; Patil et al., 2023). However, traditional LLMs perform poorly on reasoning task as shown by work such as Williams & Huckle (2024); Seßler et al. (2024); Hosseini et al. (2024). Research on enhancing LLM reasoning capabilities can be decomposed into two categories: Training, and Inference Time Scaling (Raschka, 2025).

**Training Time Scaling.** To enhance the performance of LLMs on reasoning tasks, a substantial body of work has explored different training techniques. A promising path has been to fine-tune LLMs on reasoning traces using RL and SFT techniques. For example, Deepseek (Shao et al., 2024) introduced the Group Relative Policy Optimization (GRPO). Unlike the classic Proximal Policy Optimization (PPO) (Schulman et al., 2017), GRPO estimates the reward using group scores, which naturally enhances model's capability to generate reasoning traces. Nevertheless, this process renders the models to be much more verbose.

**Inference Time Scaling.** By contrast, Inference Time Scaling aims to enhance model performance on reasoning tasks after training. To do so, researchers looked for reasoning instances at sampling (Wang et al., 2023; Snell et al., 2024), or prompt engineering techniques (Wei et al., 2023; Muennighoff et al., 2025). Since the generation of LLMs is often highly variable, multiple inferences could lead to a wide diversity of answers. These techniques help the reliability and correctness of models. However, they comes at the cost of both compute resources and runtime.

#### C.2 LRMs are Inefficient

From the literature, we can observe a strong push for more efficient LRMs. This section will aim to define the efficiency problem of LRMs, and the underlying gaps in the literature.

Patterns of inefficiency. It is only recently that a few papers highlighted that LRMs tend to generate thousands of tokens to solve reasoning problems. Chen et al. (2025b) was the first to formalize this issue, known as *model overthinking*, where LRMs tend to generate a disproportionate number of tokens for fairly simple problems. In this emerging field, work such as Qu et al. (2025) began to look for evidence of inefficiency and showcased examples where reasoning models exhibit patterns of inefficiency mainly due to redundancy such as verbose problem reformulation, or over-verification. Munkhbat et al. (2025) also support this claim and gives a general overview of findings in the field. Su et al. (2025) claims that majority of tokens generated by LRMs ensure textual coherence rather than core reasoning. Experimentation presented by Luo et al. (2025) supports the argument that longer answers from models does not necessarily lead to more accurate answers, and in some cases can even lead to worse answers. However, to the best of our knowledge, the literature lacks tools to systematically identify patterns of inefficiency such as redundant generation.

**Origins of inefficiency.** First, Sui et al. (2025a) observed from the Deepseek paper DeepSeek-AI et al. (2025) that GRPO training leads to a positive correlation between the accuracy of the model and the average number of tokens generated. Furthermore, Galichin et al. (2025) highlighted that LRMs such as Deepseek-R1 tend to generate tokens that are responsible for the generation of certain types of reasoning steps. Referred to as *reasoning tokens* (e.g. Wait, Hum, Let me verify, ...), the authors demonstrated that these tokens are responsible for guiding the generation of intermediate stages of reasoning, such as pausing the thought, re-evaluating the current answer, or exploring novel solutions. However, their empirical evaluation on activation steering showed that the over-reliance on such tokens tends to increase the verbosity of model output. Muennighoff et al. (2025) also supports this claim. Indeed, they demonstrated that prompting reasoning tokens during the generation forces the model to reason more.

#### C.3 OPEN-SOURCE REASONING PATH

Table 1 shows the difference between close-source and open-source models. Close-source models tends to hide the raw reasoning traces generated by models.

LRMs	Open Weights	Reasoning Traces
DeepSeek-R1	Yes	Yes
QwQ	Yes	Yes
03 / 04	No	Partial
Claude 3.7	No	Partial
Gemini 2.5 Pro	No	Partial

Table 1: Comparison of LRMs and reasoning traces - *Partial* stands for models that do give access to full reasoning traces

#### D MATHEMATICAL REASONING DATASETS

Table 2 presents the selected reasoning datasets, including their references, and the number of samples per training and testing folds. Due to computational resources constraints, we limit the size of the training datasets to smaller subsets of their full versions - specifically, 1,000 samples for MATH500 Lightman et al. (2023) and 3,000 samples for GSM8K Cobbe et al. (2021). These values were selected to have approximately twice the size of the test datasets of training samples. We used a seed of 42 to infer training datasets.

Dataset	# Train	# Train used	# Test	
MATH500	12,000	1,000	500	
GSM8K	7,474	3,000	1,318	

Table 2: Description of selected mathematical dataset

#### E DEFINITION OF REASONING STEPS

Table 3 gives an overview of the different definitions of reasoning steps in the literature. Notably, we can define 4 categories, namely: token, sentence, paragraph, and dynamic step delimitation levels.

Step Type	Ref.	Method	M.A.
Token	Light et al. (2025)	_	X
TORCH	Yao et al. (2023)	=	X
Sentence	Fu et al. (2023)	[.,;]	X
Sentence	Yao et al. (2023)	_	X
	Cao et al. (2025)	.\n\n	X
Paragraph	Lightman et al. (2023)	\n	
Taragraph	Fu et al. (2023)	\n	X
	Park et al. (2024)	n n, ###, ***	X
	Yao et al. (2023)	i	
	Light et al. (2025)	Sampling	X
	Zelikman et al. (2024)	<  startofthought  >	
Dynamia	Sui et al. (2025b)	<next_step></next_step>	
Dynamic	Fu et al. (2023)	"step i"	
	Paul et al. (2024)	<equation></equation>	
	Welleck et al. (2022)	special tokens	

Table 3: Overview of types of step-delimiters - M.A. for Model Agnostic

Figure 8 presents an example of reasoning step from the Deepseek-R1-Distill-Llama8B model. We observe that the delimiter .  $\n \n$  is very often being generated, specifically when the model executes a series of operations.

#### **Example of Reasoning step**

[...] \n\n Wait, let me double-check my calculations to make sure I didn't make any mistakes.\n\n First, f(-2): \n\n 3\*(-2) is -6, minus 2 is -8. Denominator: -2 -2 is -4. -8 divided by -4 is 2. Correct.\n\n f(-1): \n\n 3\*(-1) is -3, minus 2 is -5. Denominator: -1 -2 is -3. -5 divided by -3 is 53. Correct.\n\n f(0): \n\n 3\*0 is 0, minus 2 is -2. Denominator: 0 -2 is -2. -2 divided by -2 is 1. Correct.\n\n Adding them up:  $2 + 53 + 1.\n\n$  Converting to thirds: 2 is 63, 1 is 33, so 63 + 53 + 33 = 143. Yep, that seems right.\n\n So, I think I did everything correctly. I don't see any mistakes in my calculations or reasoning.

Figure 8: Extract of the reasoning traces of DS-Llama8B - from inference of question 2 of MATH500 test dataset. The extract of the reasoning trace showcases a verification step of the model. Back to line symbols are highlighted in bold.

#### F **ALGORITHMS**

972

973 974

975 976

977

978

979 980

981 982

983

984

985

986

987

988

989

990

991

992

993

994

995

996

997

998

999

1005

#### STEP-WISE GENERATION F.1

To generate the reasoning traces of models step-by-step, we need to modify the model.generate function from Hugging Face. However, this process comes at the cost of latency in model generation since we need to interrupt the generation process at each step. The algorithm is presented in Algorithm 1.

#### **Algorithm 1** Step-wise Generation

```
Require: Prompt x; reasoning delimiter \alpha \in V; minimal step size k \in \mathbb{N}; max steps T_{\text{max}}; language model
                 \mathcal{M}; tokenizer \mathcal{T}; EOS token
             1: y \leftarrow \mathcal{T}(x)

    ▶ Tokenized input

             2: S \leftarrow []; \beta \leftarrow \emptyset

    ▷ Initialize output and buffer

             3: s \leftarrow 0
             4: while s < T_{\rm max} do
                     t \leftarrow \mathcal{M}(y)
             5:

⊳ Generate next token

                     y \leftarrow y + t
             6:
                     \beta \leftarrow \beta + \mathcal{T}^{-1}(t)
             7:

    Add decoded token to buffer

                     if EOS in y then
             8:
                                                                                                          > Stop inference if EOS generated
             9:
                          Append \beta to S
            10:
                          break
            11:
                      end if
                     if \beta ends with \alpha then
            12:
                                                                                                                    ▷ Complete and valid step
            13:
                          if \beta > k then
            14:
                               Append \beta to S
                               \beta \leftarrow"
            15:
                                                                                                                             \triangleright Increase S by one step
            16:
                               s \leftarrow s + 1
            17:
                          else
                               Continue
                                                                                             \triangleright Continue until next \alpha or EOS is generated
            18:
            19:
                          end if
            20:
                      else
1000
            21:
                          Continue
1001
            22:
                      end if
1002
            23: end while
            24: return S
1003
```

#### F.2 EARLY STOPPING ALGORITHM

Algorithm 2 lists the Step-Tagging Early-Stopping criteria. The user needs to define a constraint  $\{\tau*,\delta\}$ , and input a Binary Step-Tagger  $\phi_{\tau^*}$ , which returns 1 if the step tag is  $\tau^*$  and 0 otherwise. If the constraint breaks, the algorithm stops the generation, and prompts the model with  $\mathcal{P}_{\text{exit}}$  to give the current best answer.

#### Algorithm 2 Step-Tagger Early-Stopping

```
Require: Prompt x; reasoning delimiter \alpha \in V; minimal step size k \in \mathbb{N}; max steps T_{\max}; Reasoning
     Language Model \mathcal{M}; tokenizer \mathcal{T}; EOS token \gamma; Constraint \{\tau^*, \delta\}; Binary Step-Tagger \phi_{\tau^*}; Early-Exit
     Prompt \mathcal{P}_{exit}
 1: y \leftarrow \mathcal{T}(x)

    ▶ Tokenize the input

2: S_{running} \leftarrow [];
                                                                                                                          ▷ Initialize output
 3: t \leftarrow 0
                                                                                                      \triangleright Initialize frequency track of \tau^*
4: f_{\tau^*} \leftarrow 0
                                                                                                     ▶ Generate until constraint breaks
5: while c_{\tau*}(S_{running}, \delta) do
          Generate step s_i using \mathcal{M}, \alpha, where |s_i| > k
 7:
 8:
          if \phi_{\tau^*}(s_i) then f_{\tau^*} \leftarrow f_{\tau^*} + 1
                                                                                                                     9:
10:
               Continue the generation
11:
          end if
12:
          t \leftarrow t + 1
13: end while
                                                                                                \triangleright Infer \mathcal{M} with the early exit prompt
14: y \leftarrow \mathcal{M}(y + \mathcal{P}_{\text{exit}})
15: return y
```

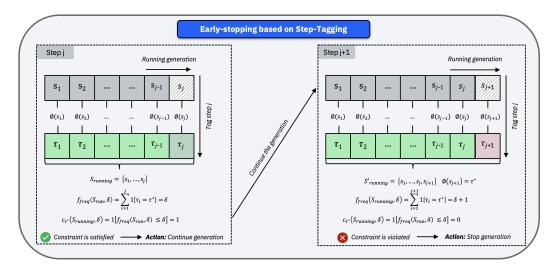


Figure 9: Illustration of early-stopping based on Step-Tagging

#### G SELECTING THE MINIMAL NUMBER OF TOKEN k

#### G.1 IDEAL EARLY-STOPPING AS A PROXY FOR THE STEP-SIZE

Grounding our approach in the literature. One crucial component of our work is the formalization of the reasoning step. From the literature, we selected the definition that seems to be the most agnostic to the model, and use-cases Cao et al. (2025) (see Section 3.2). Prior research in social sciences, which we believe can be compared to mathematical reasoning, further support that a reasoning step should be self-contained: "each reasoning step  $e_i$  represents a single piece of evidence contributing toward the social inference to select an answer  $A_a$  from A" (Mathur et al., 2025)[p.3].

**Objective and Motivations.** To support this claim, we are looking to produce a segmentation such that each unit clearly reflect a contribution toward the final answer. In addition, we note that this claim supports our problem setting: monitoring becomes more informative when the segmentation of the reasoning is well conducted. In our step definition, the step segmentation is controlled by a minimal number of token per step k. To apply this claim, we should first find a way to see if steps contain one or multiple thoughts given a value of k. This selection of k is important since we base the rest of our analysis and work on this step definition.

**Methodology.** To select optimal values of k for the three models studied, we rely on the *Ideal Early-Stopping* ( $\mathcal{IES}$ ) baseline as a signal for the quality of the step segmentation (see Section 5). First, we consider the accuracy of  $\mathcal{IES}$ , which reflects the point at which a correct answer first appears in the reasoning trace. When k is set too large, reasoning steps are likely to contain multiple distinct thoughts. In such cases, correct intermediate answers may be overwritten by later steps, which can potentially reduce  $\mathcal{IES}$  accuracy, and thereby contradicting our claim (more than one thought per step). Second, we analyze the mean number of tokens per sample given by the  $\mathcal{IES}$  baseline. If k is too large, we expect the average step length to increase, potentially erasing all efficiency gained that the  $\mathcal{IES}$  baseline is designed to provide. We applied our methodology on reasoning traces obtained on the MATH500 train dataset, for its diversity in problem complexity.

**Evaluation.** Figure 10 showcases the accuracy of the  $\mathcal{IES}$  baseline (red), and its average number of tokens per sample (blue), for the three LRMs on the MATH500 train dataset, using values of k ranging from 1 to 1,000. As expected, we observe that the accuracy generally drops when k increases. In the meantime, the minimal number of tokens of the ideal early-stopping criteria increases when k increases.

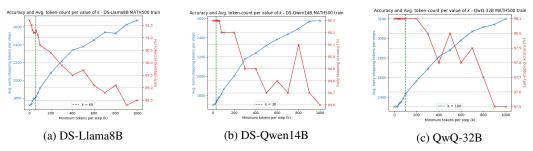


Figure 10: Selecting optimal k - Efficiency of  $\mathcal{IES}$ 

We interpret the accuracy drops when k increases as a signal that some steps are including multiple thoughts. Indeed, our metric is by definition taking the latest solution contained in each step (if any, e.g. located at the last sentence of the step). If the steps are really small, the steps are containing a minimal number of thought (as per the model's step segmentation when generating  $.\n$ ). This explains why k=1 obtained the highest accuracy for every models. However, k needs to be set higher than 1 to avoid over-segmentation (see Appendix E). Now, if a step contained a correct answer with a lower value of k, increasing k results in adding additional reasoning text - potentially incomplete or erroneous. Then our metric would assign the outcome to the latest solution within the same step, potentially overriding the correct sub-conclusion. Figures 11 and 12 illustrate this.

**Takeaway.** On the strength of these observations, we identified sweet-spots for each models, aiming to balance monitoring efficiency and supports our claims on our step-definitions. We set k=60, k=30, and k=100 for DS-Llama8B, DS-Qwen14B, and QwQ-32B, respectively.

```
1134
         Example. Figures 11 and 12 show an example of \mathcal{IES} traces for the sample 658 from MATH500
1135
         train using the DS-Qwen14B with a reasoning step-split parameter k of 30 and 300, respectively.
1136
         We selected this example because the model initially drafted a correct answer, but overwrote it
1137
         with a wrong reasoning. While the IES algorithm returned the correct answer parsed at step 8 for
1138
         the parameter k=30, the same algorithm applied for a parameter k=300 return an incorrect
         answer, and failed to early-stop. Indeed, with k=300, the same step correspond now to the step
1139
         5, and include an additional self-corrected sentence leading to a different answer. This observation
1140
         confirms that the model sometimes over-write/overthink its correct initial reasoning, and suggests
1141
         that lower values of k tends to lead to better segmentation of the reasoning traces. Overall, this
1142
         observation comforts our observation from the previous section and our selection of k.
1143
1144
         Step 1 - Okay, so I've got this geometry problem here. It says that line AB is parallel to line CD, and the measure of angle AXF is 118 degrees.
1145
         I need to find the measure of angle FYD. Hmm, let me try to visualize this. There's an Asymptote diagram provided, but since I can't actually
         see the image, I'll have to go off the description \n
1146
         Step 2 - First, I know that when two lines are parallel, certain properties hold, like corresponding angles being equal, alternate interior angles
         1147
         Step 3 - Looking at the Asymptote code, it defines points A, B, C, D, E, F, X, and Y. A, B, C, D are the main points, with AB and CD being the
1148
         parallel lines. E is another point below AB, and F is above CD. Lines are drawn from E to F, which intersects AB at point X and CD at point
```

EF and AB, and point F is on EF. So angle AXF is at point X, formed by lines AX and FX. Since AB is parallel to CD, and EF is a transversal, there might be some corresponding angles or alternate interior angles that I can use  $\backslash n \backslash n$ 1151 Step 5 - I need to find angle FYD. Point F is on EF, point Y is the intersection of EF and CD, and point D is on CD. So angle FYD is at point

1152 Y, formed by lines FY and DY. That angle is also on the transversal EF, but on the other side of the parallel lines  $\n$ 

Y. So, essentially, line EF is a transversal cutting through the two parallel lines AB and CD n

Step 6 - Since AB is parallel to CD, and EF is a transversal, the corresponding angles at X and Y should be equal. So, if angle AXF is 118 degrees, then the corresponding angle at Y should also be 118 degrees. But wait, angle AXF is on the same side as angle FYD? Or is it on the opposite side? \n\n Wait, angle AXF is above AB, and angle FYD is above CD. Since AB and CD are parallel, the angles formed by the transversal EF should be corresponding. So, angle AXF and angle FYD are corresponding angles. Therefore, they should be equal. So, angle FYD should also be 118 degrees \n\n

Step 4 - Given that, angle AXF is 118 degrees. Let me try to identify where that angle is located. Point A is on AB, point X is the intersection of

Step 7 - But hold on, let me double-check. When you have two parallel lines and a transversal, corresponding angles are equal. So, angle AXF is at the intersection above AB, and the corresponding angle above CD would be angle FYD. So yes, they should be equal  $\n$ 

Step 8 - Alternatively, I can think about the consecutive interior angles. If I consider angle AXF and the angle adjacent to it on the same side, they should add up to 180 degrees because they are consecutive interior angles. So, if angle AXF is 118 degrees, the adjacent angle on the same side would be 180 - 118 = 62 degrees  $\n\$ 

Step 37 - \*\*Final Answer\*\*  $118 < \sinh > \ln [...]$ 

1149

1150

1153

1154

1155

1156

1157

1158

1159 1160

1161 1162

1163 1164

1165

1166

1167

1168

1169

1170

1171

1172

1173

1174

1175

1176

1177

1178

1179

1180

1181

1182

1183 1184

1185 1186 1187

Figure 11: Sample 658 from MATH500 train —— Solution:  $62^{\circ}$  —— k = 30 with DS-Qwen14B

Step 1 - Okay, so I've got this geometry problem here. It says that line AB is parallel to line CD, and the measure of angle AXF is 118 degrees. I need to find the measure of angle FYD. Hmm, let me try to visualize this. There's an Asymptote diagram provided, but since I can't actually see the image, I'll have to go off the description. \n\n First, I know that when two lines are parallel, certain properties hold, like corresponding angles being equal, alternate interior angles being equal, and so on. Maybe I can use those properties here  $\n$  Step 2 - Looking at the Asymptote code, it defines points A, B, C, D, E, F, X, and Y. A, B, C, D are the main points, with AB and CD being the parallel lines. E is another point below AB, and F is above CD. Lines are drawn from E to F, which intersects AB at point X and CD at point

Y. So, essentially, line EF is a transversal cutting through the two parallel lines AB and CD. \n\n Given that, angle AXF is 118 degrees. Let me try to identify where that angle is located. Point A is on AB, point X is the intersection of EF and AB, and point F is on EF. So angle AXF is at point X, formed by lines AX and FX. Since AB is parallel to CD, and EF is a transversal, there might be some corresponding angles or alternate interior angles that I can use  $\n$ 

Step 3 - I need to find angle FYD. Point F is on EF, point Y is the intersection of EF and CD, and point D is on CD. So angle FYD is at point Y, formed by lines FY and DY. That angle is also on the transversal EF, but on the other side of the parallel lines  $\n$ 

Step 4 - Since AB is parallel to CD, and EF is a transversal, the corresponding angles at X and Y should be equal. So, if angle AXF is 118 degrees, then the corresponding angle at Y should also be 118 degrees. But wait, angle AXF is on the same side as angle FYD? Or is it on the opposite side? \n\n Wait, angle AXF is above AB, and angle FYD is above CD. Since AB and CD are parallel, the angles formed by the transversal EF should be corresponding. So, angle AXF and angle FYD are corresponding angles. Therefore, they should be equal. So, angle FYD should also be 118 degrees \n\n

Step 5 - "But hold on, let me double-check. When you have two parallel lines and a transversal, corresponding angles are equal. So, angle AXF is at the intersection above AB, and the corresponding angle above CD would be angle FYD. So yes, they should be equal. \n\n Alternatively, I can think about the consecutive interior angles. If I consider angle AXF and the angle adjacent to it on the same side, they should add up to 180 degrees because they are consecutive interior angles. So, if angle AXF is 118 degrees, the adjacent angle on the same side would be 180 - 118 = 62 degrees. \n\n But angle FYD is on the other side, so maybe it's another corresponding angle. Wait, no, actually, angle FYD is the corresponding angle to angle AXF, so it should be equal, 118 degrees  $\n$ 

Step 12 - Wait, but just to make sure, let me think about the transversal EF. At point X, we have angle AXF = 118 degrees. Since AB is parallel to CD, the corresponding angle at Y should be the same. So, angle FYD is 118 degrees. \n\n Yes, that makes sense. \n\n Alternatively, if I consider triangle FXD, but I don't think that's necessary here because we can solve it using just the properties of parallel lines and transversals.  $\n$  So, I think I'm confident that angle FYD is 118 degrees.  $\n$  \*Final Answer\*\*  $\n$  118  $\n$  <  $\n$ 

Figure 12: Sample 658 from MATH500 train —— Solution:  $62^{\circ}$  —— k = 300 with DS-Qwen14B

#### G.2 Semantic properties implied by the parameter k

**Objective.** To complete the previous section, we adopted a complementary approach to reinforce our selection. We examine the semantic properties of reasoning steps. We show that the choice of the segmentation parameter k also affects the meaningfulness and coherence of each step. This ablation study has two objectives:

- 1. We demonstrate that certain values of *k* enhance the semantic meaning of individual steps, making them more self-contained and interpretable,
- 2. An appropriate value of k improves the accuracy of our step-tagger, since the semantic coherence impact the detection performance of step types using text classifiers.

**Methodology.** To address these objectives, we selected different values of k, namely: 1, 30, 100, 150, 200, 300. We then constructed their corresponding *datasets* by segmenting the same reasoning traces according to the different parameters k that we selected (Section 3.2). We selected the MATH500 train dataset using the seed 42, obtained from the DS-Qwen14B model. To access the ground-truth, we re-labeled each datasets using the same method defined in Section 5.1.

Impact of k on the tag distribution. Figure 13 shows that the parameter k affects both number of steps obtained and the distribution of ground-truth labels step-types. To control this distributional effects, we considered two dataset variants by pre-processing the datasets as follows: (a) **Balanced:** sampling the data to achieve a 50/50 distribution of positive and negative classes, (b) **Downsampling:** down-sampling the datasets to obtain the same number of samples (as per the smallest dataset - i.e. k = 300).

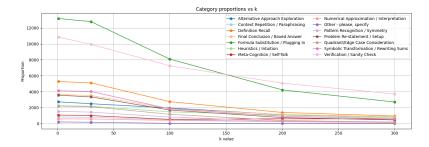


Figure 13: Step-tags distribution as per GPT-40-mini for each values of k - ReasonType taxonomy

For each dataset, we trained binary BERT classifiers with identical hyperparameters (same as in Section 5.1). Similarly to our Step-Taggers, the classification task is step-type detection, where the model predicts whether a given step corresponds to a certain step type  $\tau^*$ . We conducted experimentation on 2 step-types: (a) Validation, (b) Exploration. While *Validation* is the one of the most occurrent step-type, *Exploration* has some interesting sematic properties, and is less frequent.

Figure 14 and 15 shows the label distribution of the three variants on the *Verification* and *Exploration* step-type, respectively. While *Downsampling* enable clear comparison of performance for different values of k - same dataset size, *Balancing* labels of datasets investigates the performance of the models with ideal distribution for step-type detection. In contrast, switching off the pre-processing steps preserve the natural distribution of labels, providing insights under more realistic conditions.

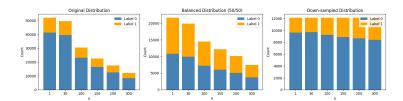


Figure 14: Distribution of dataset variants - Verification step-type

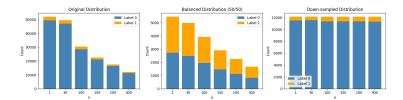


Figure 15: Distribution of dataset variants - Exploration step-type

**Evaluation.** Figure 16(a) and 16(b) present the Micro-F1 and Macro-F1 of the binary BERT classifiers trained on the different dataset types for different values of k and for the *Verification* and *Exploration* step-types, respectively. We first note that the training performed on the Exploration steps lead to higher accuracy than the one achieved on the Verification step (Macro-F1 0.8 - 0.97 vs. 0.8 - 0.87). It could be explain by the nature of the steps. The Verification steps might be more diverse, while the Exploration steps might carry more semantic meaning, making them easier to detect (specifically for the balanced dataset). For imbalanced datasets (Original and Down-sampled), it is worth noting that reporting the Macro-F1 was also important in order to access to the performance of the minority class (positive).

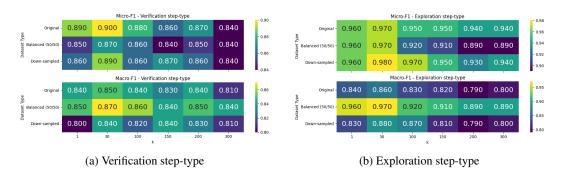


Figure 16: Step-Tagger performance per dataset distribution and step segmentation parameter k

Importantly, we observe a general trend of performance for the different configuration. The parameter k seems to impact the performance of the classifiers. Specifically, the performance seems to decrease when k is increase from values ranging from 30 to 300. For the Original Validation, the Micro-F1 drops from 0.9 for k=30 to 0.84 for k=300, similarly for the Macro-F1 (from 0.85 for k=30 to 0.81 for k=300). Surprisingly, the value k=1 lead to lower performance compared to k=30. We suspect that very small value of k (such as k=1) imply noisy steps. Indeed, Figure 8 in Appendix E confirms this observation. When k is very small, a lot of noisy steps are created, and this could perturb the training.

In addition, for both labels, the Balanced dataset allows for higher performance in term of Macro-F1. It means that classifiers are better at detecting positive classes, at the cost of the performance on the negative class (lower Micro-F1 with respect to the two other datasets). However, same effects on the values of k can be noticed. Moreover, down-sampling the dataset seems to harm slightly the performance, especially for the Verification step-type. It indicates that limiting the dataset size reduces the diversity of training samples for lower values of k.

**Takeaway.** This ablation study supports our selection of k. Based on the semantic properties of the steps, we found that higher values of k lead to a loss of semantic meaning of the steps. Similarly, small values of k implies noise in the step segmentation - which could perturb the monitoring analysis of the reasoning.

#### G.3 Influence of the parameter k on the Step-Tagging Early-Stopping criteria

**Objectives.** In the two previous ablation studies, we have seen that the step segmentation parameter k influenced the amount of information contained in each reasoning step. The following ablation study focus on assessing the impact of the parameter k on the performance of our step-tagging early-stopping criteria. It is composed of two sections: We first evaluate the impact of the parameter k on the constraint threshold  $\delta$  of our criteria (Section G.3.1), and then assess the impact of k on the performance of the criteria (Section G.3.2).

#### G.3.1 ST-ES THRESHOLDS $\delta$

 We have seen that the step segmentation parameter k influences both size and amount of information contained in each steps. To further validate our values of k, we are looking at its influence on the threshold  $\delta$  of our Step-Tagging Early-Stopping constraint  $(\tau, \delta)$ , for each step-types  $\tau$  of our taxonomy.

Impact on the constraint values. To assess the influence of k on the constraint values  $\delta$ , we adopt a different point of view. For each value of k, we re-used the datasets obtained from the reasoning traces of the DS-Qwen14B model on the MATH500 of the training datasets labeled by GPT-40-mini (see Appendix G.2). For each values of  $k \in \{1, 30, 100, 150, 200, 300\}$ , Figures 18 and 18 compare the Accuracy and the Average number of tokens per sample for each constraints  $\delta \in [1, 10]$  for each step-types of the taxonomy, respectively. The positions and speed of convergence will allow us to assess the impact of the k on the constraints values  $\delta$ .

**Evaluation.** First, we observe that lower values of k (e.g.  $k \in \{1,30\}$ ) reveal smoother trade-off for both accuracy and token count. Indeed, the curves converge less quickly to the original traces (plateau when  $\delta \to \infty$ , tending to no-constraints i.e. the standard inference).

In comparison, when k grows, the curves are sharper and tends to converge faster, which offers less trade-off between accuracy and token-count. Furthermore, the curves of the different step-types becomes less distinguishable and overlaps much more than for lower values of k. In the context of our Early-Stopping framework, these observations confirm the importance of keeping lower values of k to enable efficient and interpretable inference through the constraints  $\{\tau, \delta\}$ .

**Takeaway.** This ablation study shows that higher values of k tends to decrease the controllability of the inference using our framework.

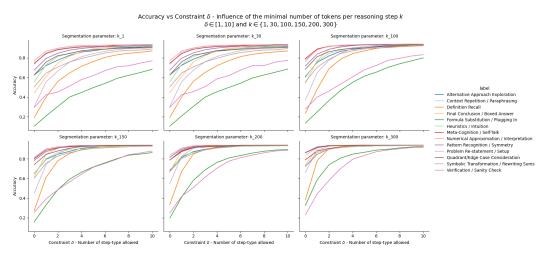


Figure 17: Constraint values  $\delta$  vs. Accuracy for each value of  $k \in \{1, 30, 100, 150, 200, 300\}$  - Qwen14B on MATH500 train (1,000 samples) - Seed 42,  $\delta \in [1, 10]$ 

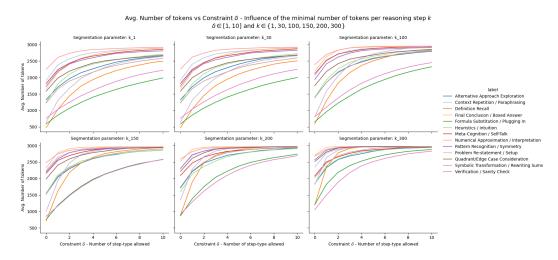


Figure 18: Constraint values  $\delta$  vs. Average Number of Tokens per sample for each value of  $k \in \{1, 30, 100, 150, 200, 300\}$  - Qwen14B on MATH500 train (1,000 samples) - Seed 42,  $\delta \in [1, 10]$ 

#### G.3.2 Influence of the parameter k on the ST-ES performance

We complete this analysis with an ablation study on the influence of the step segmentation parameter k on the calibration of the Step-Tagging Early-Stopping. We are assessing if the parameter k affects the performance of our early-stopping framework (i.e. selecting optimal constraints).

**Methodology.** To address our objective, we performed the calibration experiment presented in Section 5.3, using the datasets obtained on our previous experiment, i.e. for the values of  $k \in [1,30,100,150,200,300]$  (see Section G.2). For each value of k, the experiment resulted in one *Pareto Curve*, corresponding to the most efficient constraint parameters among the set of thresholds  $\delta$  and step-types  $\tau$ . For each value of k, we re-used the datasets obtained from the reasoning traces of the DS-Qwen14B model on the MATH500 of the training datasets labeled by GPT-40-mini.

**Evaluation.** Figure 19 presents the Pareto Curves of the Step-Tagging Early Stopping applied to DS-Qwen14B on MATH500 train, for different segmentation parameters k. To enhance the analysis, Table 4 showcases the AUC of the Pareto curves (restricted to overlapping token ranges between curves). We can observe that the lower values of k (i.e. 1 and 30) lead to higher efficiency (higher accuracy and lower token-count) since their associated Pareto curves are above others for almost every levels of complexity. Table 4 confirms this observation, with  $k \in \{1,30\}$  presenting the higher AUC values.

It is worth noting that the disparities between curves are greater when the complexity is lower. We have seen that the complexity increases the verbosity of models. Therefore, more complexity also increases the number of steps, regardless of the value k. For this reason, more complexity increases the opportunity of early-stop, minimizing the impact of k.

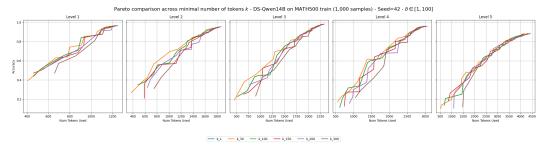


Figure 19: Calibration of Step-Tagging Early-Stopping criteria - Comparison of Pareto Curve for each value of  $k \in \{1, 30, 100, 150, 200, 300\}$  - Qwen14B on MATH500 train (1,000 samples),  $\delta \in [1, 100]$ 

k value			Average across Levels			
70 14140	1	2	3	4	5	riverage across hevels
1	456.15	832.90	1053.85	1376.53	1829.24	1109.74
30	455.42	833.33	1053.79	1376.66	1811.43	1106.13
100	450.06	801.29	1017.66	1327.18	1790.32	1077.30
150	440.02	774.62	1004.62	1304.16	1775.65	1059.81
200	436.50	774.68	993.03	1277.69	1789.72	1054.32
300	421.18	742.81	949.12	1192.54	1726.99	1006.53

Table 4: Area Under the Curve (AUC) of the Pareto Curves - Restricted to overlapping token range between curves

**Takeaway.** The parameter k influences the performance of the Step-Tagging Early-Stopping. Lower values of k seems to increase the performance since these values lead to finer-grained segmentation (and therefore more flexible early-stopping). However, the divergence in performance seems to reduce when the complexity of question (i.e. verbosity of the model) increases.

#### G.4 SUMMARY OF TAKEAWAYS

Our ablation studies contributed to find and validate the parameter k, minimal number of token per step, of our definition of reasoning step. From the Sections of this Appendix, we can formulate three main takeaways:

- 1. The Ideal-Early Stopping  $(\mathcal{IES})$  criteria seems to be a good signal for selecting the segmentation parameter k. We applied this method and compared early-stopped traces from different values of k. For the three LRMs that we selected, we observe that the  $\mathcal{IES}$  accuracy drops when values of k increase. Lower values of k seems to result in steps being more self-contained, including a minimal number of thoughts/conclusion, allowing better early-stopping performances.
- 2. Based on the performance of sentence classifier trained on datasets with different values of k, we found that higher values of k lead to a loss of semantic meaning of the steps. Similarly, small values of k implies noise in the step segmentation which could perturb the monitoring analysis of the reasoning.
- 3. An ablation study on the calibration process to find the most efficient constraints per dataset and models shows that higher values of k tends to decrease the performance and controllability of the inference using our framework.

For our reasoning step definition, these takeaways confirms that users needs to carefully select optimal value of k, since this parameter has direct implication on both information contained and semantic properties of the steps.

#### H VALIDATION OF THE REASONTYPE TAXONOMY

#### H.1 REASON-TYPE TAXONOMY FOR IDENTIFYING REASONING BEHAVIORS

**Objective.** This ablation study is looking at further validating our ReasonType taxonomy. In other words, we are investigating whether our proposed taxonomy captures meaningful distinctions in reasoning steps. We are looking to demonstrate that:

- 1. The ReasonType taxonomy enable semantic distinction of the type of reasoning.
- 2. Our annotation method with the GPT-40-mini model, coupled with the ReasonType taxonomy, is a robust method to access to the ground-truth labels of the reasoning steps.

**Methodology.** To address our objective, we compare the performance of BERT classifiers across Original labels (OG - from GPT-40-mini annotation using the ReasonType taxonomy), and shuffled labels for three step-types, namely: *Verification, Exploration* and *Self-Talk*. For the shuffled labels version, we took the exact same proportion of positive labels as in Original datasets, and used random shuffle with a seed of 42. Each experiment is run on the same training and testing dataset, i.e. the steps obtained with a segmentation parameter k=30, from the MATH500 training dataset on the DS-Qwen14B model. We trained BERT classifiers following the exact same training configuration (see Section 5.1). To compare performances, we report both training loss, and classification metrics (precision and recall on both classes, along with macro and micro average.)

Evaluation. Figure 20 shows the training loss of the Original and Shuffled versions, for the three labels. We observe that models trained on the Original labels presents significant lower losses, and are smoothly decreasing. It demonstrate that the Original datasets contains meaningful patterns between reasoning steps and their labels. In comparison, the models trained on shuffled labels present almost constant loss, relatively higher than the one from the Original labels.

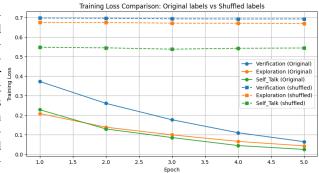


Figure 20: Training losses - ReasonType vs. Shuffled labels

Furthermore, Figures 21(a) and 21(b) show the Precision and Recall classification metrics on the testing dataset, respectively. For Original runs, both classes (0 and 1) achieve good performance despite dataset imbalancity, with Macro average Precision and Recall lying between 0.76 and 0.90 across labels. In comparison, shuffled runs presents poor results, with models failing in predicting positive classes - Precision and Recall of class 1 between 0.00 and 0.06. Along with the training loss, theses metrics highlight that the models trained on shuffles labels cannot learn meaningful relations between steps and labels. In comparison, Original labels (from the ReasonType taxonomy) resulted in satisfying model performance, and smooth training.

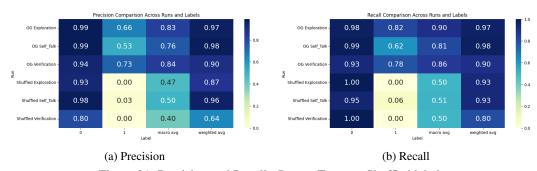


Figure 21: Precision and Recall - ReasonType vs. Shuffled labels

**Takeaway.** Overall, these results comforts our finding that the ReasonType taxonomy labels enable annotation methods to results in reasoning steps carrying semantic meaning.

#### H.2 Tracking step-types to design interpretable early-stopping criteria

**Objective.** This ablation study is looking at motivating our Step-Tagging Early-Stopping (ST-ES) framework, as well as our taxonomy. We have seen that the ST-ES approach is an interpretable early-stopping criteria since the user can select which type of reasoning step to limit to make inference more efficient. However, we are wondering if simpler approaches such as a *simple step-count* is more effective. Despite having less requirements - since no step-tagging is required - only stopping the reasoning based on the number of steps generated could potentially yield to better results.

**Methodology.** To assess this alternative approach is more effective, we lead the same calibration study as in Section 5.3 on the MATH500 training dataset, using the 3 selected LRMs. For clarity, we applied a common threshold across every levels of complexity of the dataset, and reported their average accuracy and token-count. To compare our taxonomy with the *simple step-count* approach, we present two types of runs:

- ReasonType taxonomy: We performed the same experimentation for each single labels of the ReasonType taxonomy. For each tags  $\tau_i \in \mathcal{T}$ , we combined a threshold value  $\delta \in [1, 20]$ .
- Simple Step-count: The same experiment is done without the labels. To do so, we only applied simple step-count for thresholds  $\delta \in [1, 100]$ .

**Evaluation.** Figures 22, 23, and 24 present the experiment on the DS-Llama8B, DS-Qwen14B, and QwQ-32B models, respectively. For both types of runs, each combination of threshold and step-type results in a point (Average Number of Tokens, Accuracy). For the ReasonType taxonomy, each step-types forms a blue curve, with color gradient and different markers to differentiate the tags of the taxonomy. The same experiment with a unique label (equivalent to a simple step-count early-stopping) resulted in a unique curve, print in red.

We observe over the three models that the red curve is almost constantly under the blue curves, in particular for the DS-Qwen14B model. Specifically, for token-count from 1,500, the red curve is bellow the curves of the *Problem-Restatement*, *Exploration*, and *Intuition*, for the DS-Llama8B, DS-Qwen14B, and QwQ-32B models, respectively. Furthermore, for lower token-counts, the red curve is still under the blue curves of the *Problem-Restatement* for both DS-Llama8B, DS-Qwen14B models. Overall, it means that for a given number of tokens, implementing a stopping criteria based on the type of the steps seems to yield to higher accuracy than simply stopping the generation based on the number of steps.

**Takeaways.** Beyond interpretability, this experiment shows that tracking the types of steps yields to better performance than an early-stopping criterion based on the step-count.

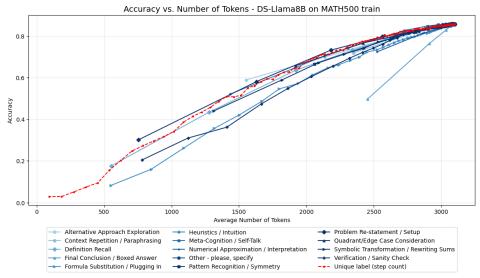


Figure 22: Accuracy vs. Average number of tokens - Step-Tagging Early-Stopping curves per step-types (in blue) and simple step-count (in red) - DS-Llama8B on MATH500 train

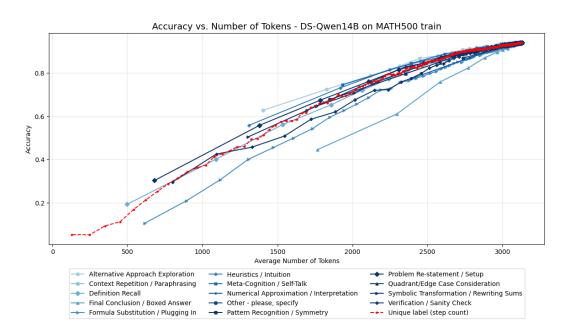


Figure 23: Accuracy vs. Average number of tokens - Step-Tagging Early-Stopping curves per step-types (in blue) and simple step-count (in red) - DS-Qwen14B on MATH500 train

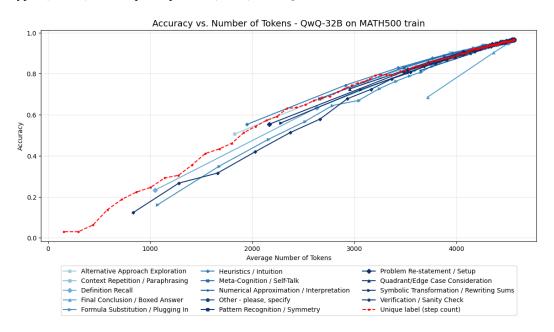


Figure 24: Accuracy vs. Average number of tokens - Step-Tagging Early-Stopping curves per step-types (in blue) and simple step-count (in red) - QwQ32B on MATH500 train

#### H.3 CERTAIN STEP-TYPE ARE BETTER EARLY-STOPPING CRITERIA

**Objective.** In section 5.3, we observed that selecting the step-tagging early-stopping constraint (both threshold and step-types) was not trivial since it depends on models, and complexity of a given problem. This ablation study is looking at showing that certain step-types are better to apply our constraint. Specifically, we are looking to demonstrate that our taxonomy allows us to state that different step-types are leading to different efficiency trade-off.

**Methodology.** Similarly to the sub-section H.2, we are showing the same experiment, but we focus on the curves resulting from the ReasonType taxonomy (in blue). We plot the same Figures, but with one distinct color and marker for each step-type to better differenciate the curves.

**Evaluation.** Figure 25 present the different curves obtained applying our early-stopping framework on the labels of the ReasonType taxonomy, for the DS-Qwen14B on the train MATH500 dataset. We observe that each step-type results in curves with different lengths (token-count range), and widths (accuracy range).

First, the *Exploration* curve seems to be the most efficient for moderate to high token-count range (from approximately 1,400 to 3,000). Indeed, the curve stands above all other curves, meaning that for constraints with the *Exploration* tags results in highest accuracy with equivalent token-budget.

In contrast, other step-types such as *Problem Re-statement*, *Verification* or *Definition Recall* cover larger token-count range (from 500-750 to 3,000). They introduce more flexibility for limited token-budget, but appears less efficient when the token-count grows.

Figures 26 and 27 present similar findings for the DS-Llama8B and QwQ-32B models, respectively.

**Takeaway.** Overall, we can conclude that our taxonomy coupled to our methodology allows us to demonstrate that all reasoning step types are not equally contributing to the reasoning progression.

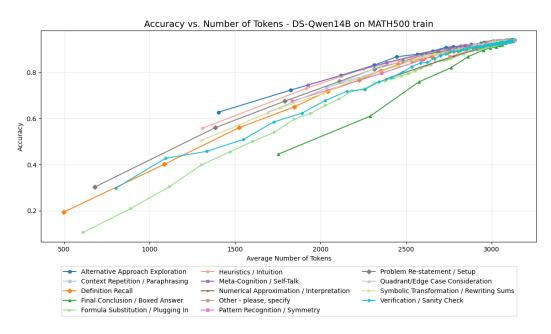


Figure 25: Accuracy vs. Average number of tokens - Step-Tagging Early-Stopping curves per step-types - one color per step-type - DS-Qwen14B on MATH500 train

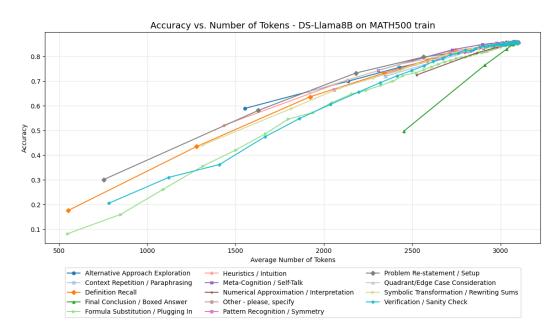


Figure 26: Accuracy vs. Average number of tokens - Step-Tagging Early-Stopping curves per step-types - one color per step-type - DS-Llama8B on MATH500 train

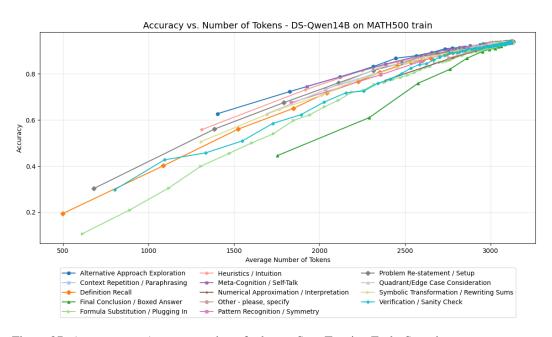


Figure 27: Accuracy vs. Average number of tokens - Step-Tagging Early-Stopping curves per step-types - one color per step-type - QwQ-32B on MATH500 train

H.4 COMPARISON TO ALTERNATIVE TAXONOMY

**Objective.** In this ablation study, we assess how robust our taxonomy is for the Step-Tagging Early-Stopping criteria. To do so, we run the calibration experiment that we conducted on different alternative versions of our original taxonomy.

**Alternative Taxonomies.** Our original taxonomy is wide and fine-grained, containing 13 categories of labels (excluding the placeholder label "Other"). Therefore, we reduce the number of labels in the taxonomy, and grouped similar labels at different levels of abstraction. Table 5 shows resulting the taxonomies, considering from 13, 6, 4, 2, and 1 labels. A unique label represent the simplest form of constraint, where we obtain only a constraint on the number of steps.

Original Taxonomy	6-labels	4-labels	2-labels	1-label
Problem Re-Statement Context Repetition Definition Re-call	Setup	Early Reasoning		
Formula Substitution Symbolic Transformation	Manipulation	Mid Reasoning	Early Reasoning	Reasoning
Edge Case Pattern Recognition	Analysis	ma rousoming		
Verification Heuristic / Intuition	Checking			
Exploration Interpretation Self-Talk	Meta Reasoning	Late Reasoning	Late Reasoning	
Final Conclusion	End Reasoning	End Reasoning		

Table 5: Alternative taxonomies - we regrouped labels at different levels of abstraction to observe the impact of the taxonomy on the Step-Tagging Early-Stopping criteria

**Methodology.** To address our objective, we performed the calibration experiment presented in Section 5.3, using the different taxonomies (i.e. vocabulary of tags  $\mathcal{T}$ ). For each taxonomies, the experiment resulted in *Pareto Curve*, corresponding to the most efficient constraint parameters given a threshold  $\delta$  and a type-step  $\tau$  (lying in the taxonomy tested). For each models, we re-used the MATH500 and GSM8K training datasets labeled by GPT-40-mini using our methodology explicated in Section 5.1. We then merged labels as in Table 5.

The process of merging labels artificially increases the number of labels per datasets. Therefore, to encompass a wider range of early-stopping values, we performed the calibration for threshold values  $\delta \in [1,100]$ . To evaluate the impact of downsizing the original taxonomy, we look at the position of the Pareto Curves compared to each other. The taxonomy giving the most efficient constraints is the one located at the top left (maximizing the accuracy while minimizing the number of tokens).

**Evaluation.** We can note that downsized tag vocabularies outperformed our Original taxonomy for low complexity queries (e.g. 1 and 2), specifically for 1 and 2-labels taxonomies. It can be justified by the nature of the low-label taxonomies. For the purple curve, it acts like a simple step count constraint. For low-level complexity, the reasoning traces are less verbose (less tokens), meaning that it is easier to find a simple form of constraint.

However, the Original and Label-6 taxonomy seem to perform well on higher level of complexity. For higher level of complexity, it seems that the semantic meaning of the steps plays a role in the determination of the final constraint. Nevertheless, it is worth noting that the differences observed are relatively small. Pareto Curves of Original and Label-6 are sensibly close, meaning that our approach seems to be robust to the granularity of the taxonomy.

**Takeaway.** Our framework seems to be robust with regards to the taxonomy selected. Nevertheless, fine-grained taxonomies seems to lead to better performance for higher degree of complexity of questions, while simpler taxonomies might be more adapted to simpler problems (i.e. less verbose).

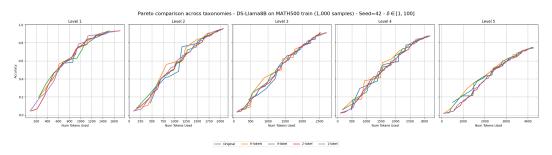


Figure 28: Calibration of Step-Tagging Early-Stopping criteria - Comparison of Pareto Curve for each taxonomies - Llama8B on MATH500 train (1,000 samples)

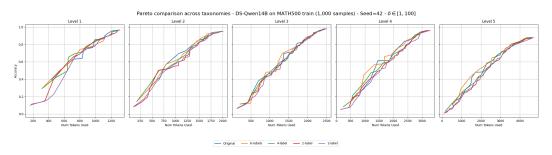


Figure 29: Calibration of Step-Tagging Early-Stopping criteria - Comparison of Pareto Curve for each taxonomies - Qwen14B on MATH500 train (1,000 samples)

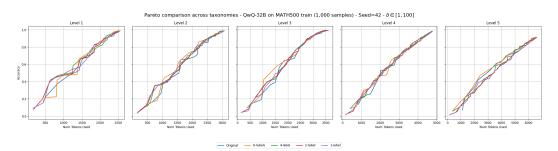


Figure 30: Calibration of Step-Tagging Early-Stopping criteria - Comparison of Pareto Curve for each taxonomies - QwQ-32B on MATH500 train (1,000 samples)

#### H.5 SUMMARY OF TAKEAWAYS

Overall, our ablation studies validate our taxonomy. Key takeaways are the followings:

- 1. The ReasonType taxonomy labels enable annotation methods to results in reasoning steps carrying semantic meaning.
- 2. Beyond interpretability, tracking the types of steps using the ReasonType taxonomy yields to better performance than an early-stopping criterion based on simple step-count.
- 3. The ReasonType taxonomy coupled to our methodology allows us to demonstrate that all reasoning step types are not equally contributing to the reasoning progression.
- 4. Our framework seems to be robust with regards to the taxonomy selected. Nevertheless, fine-grained taxonomies such as ReasonType seems to lead to better performance for higher degree of complexity of questions, while simpler taxonomies might be more adapted to simpler problems (i.e. less verbose inferences).

#### I ANALYSIS OF REASONING STEP TYPES

#### I.1 PROMPTS

Figures 31 and 32 show the two prompts used on GPT-40-mini to build the taxonomy and generate the step tags on the traces of the LRMs, respectively.

#### **Prompt Taxonomy**

Below is a reasoning trace of a reasoning language model, split by steps. In these examples, can you please identify the different type of steps? Suggest some reasoning-type labels for each of them.

- Step 1: {step\_1}
- [...]
  - Step t: {step\_t}

Figure 31: Prompt used to generate the Taxonomy

#### **Prompt Taxonomy**

input=[ "role": "system", "content": "Classify the following reasoning step into one of the categories defined. Classes = {taxonomy}", "role": "user", "content": step ],

Figure 32: Prompt used to monitor the steps

Section I contains an additional analysis of the reasoning steps and tags issues to complete our work.

#### I.2 STATISTICS ON THE REASONING STEPS COUNT AND TYPES

Table 6 presents statistics on the number of steps and GPT-40-mini annotation for each models on both datasets we selected. Results are averaged for the seed 42 on test datasets.

Dataset	Model	# Tok. / Steps	# Steps	Runtime
MATH500	DS-Llama8B	85.29	44.25	42.18
	DS-Qwen14B	71.13	46.92	46.32
	QwQ-32B	216.70	21.15	19.55
GSM8K	DS-Llama8B	78.69	6.83	5.80
	DS-Qwen14B	74.52	7.16	7.48
	QwQ-32B	150.29	13.33	17.63

Table 6: Avg. # of steps and annotation runtime per sample

**Step occurrence.** Figure 33 presents the average number of consecutive steps of the same category. Notably, several high-frequent step types (such as *Formula Substitution*, *Verification* and *Formula Substitution*) tend to appear multiple times consecutively. Conversely, some steps tends to appear only once, where their average consecutive step is close to 1.

These observations support our design selection. Firstly, the sequence of repeated labels increases the robustness of the step classification and detection using our Step-Tagging module. Indeed, if a label appears multiple times in a row, local misclassifications are less likely to impact the overall framework. Secondly, the fact that certain steps tend to appear only one at a time (such as Final Conclusion, Interpretation, or Context Repetition) justifies our use of a token threshold k for step delimitation. In fact, some categories often appear as single instances, suggesting that the step encapsulates a single type of thought.

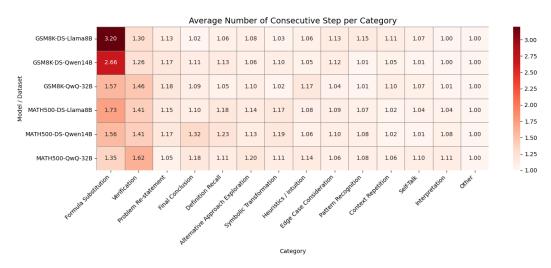


Figure 33: Avg. # of consecutive steps of same category

**Reasoning patterns.** We observe that our Step-Tagging framework allows the monitoring process to clearly follow the reasoning progression of the model. Figure 34 presents the step-types of the reasoning traces of the LRMs for a single sample of the MATH500 dataset. We selected samples that resulted in approximately the same number of steps across the models to allow a fair comparison between the models.

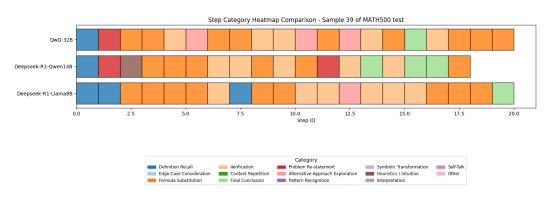


Figure 34: Reasoning patterns - Sample 109 of MATH500

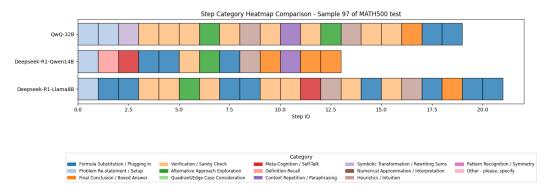


Figure 35: Reasoning patterns - Sample 97 of GSM8K

We observe a pattern in the reasoning traces. For both datasets, LRMs exhibit the same type of steps at the beginning of their generated output (e.g., *Definition Recall* or *Problem Re-statement*). Then we

notice a heavy use of *Verification* and *Formula Substitution*, helping the model to find and refine its current answer. Later, we observe the appearance of transition steps such as *Self-Talk* or *Alternative Approach Exploration*, which seems to lead to different answers and increases the diversity of the model's answers. We also observe that *Final Conclusion* steps appear in the middle of the reasoning traces, meaning that the model tends to draw intermediate conclusions but still pursues its reasoning, presumably because it is uncertain of the validity of the current solution.

**Visualization of reasoning step types.** To further assess the quality of the annotation, we computed the t-SNE projection of BERT-encoded reasoning steps from 500 DS-Llama8B traces (see Figure 36). For clarity, we excluded the most frequent step-types, *Verification* and *Formula Substitution*, as well as *Other* (since it does not contain any semantic meaning). Some step types (e.g., *Definition Recall, Problem Re-Statement*) are dispersed, which is likely due to question-specific semantics. However, we can observe distinct semantic clusters, particularly for *Interpretation, Self-Talk* and *Exploration*. Notably, the *Self-Talk* cluster lies semantically close to *Exploration*, reflecting their conceptual overlap: internal dialogue that often prompts the model to question its assumptions, encouraging alternative approaches.

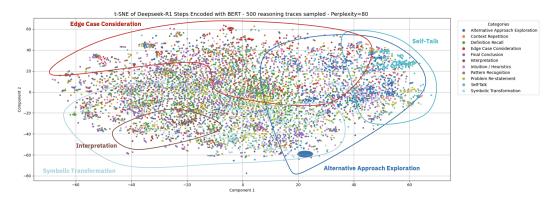


Figure 36: t-SNE visualization of 500 Deepseek-R1 reasoning steps encoded with bert-base-nli-mean-tokens (PPL = 80).

1998 **Example of Step-Tagging reasoning step annotation by GPT-40-mini.** Figure 37 present an example of reasoning trace annotated by GPT-40-mini using our step-tags taxonomy. Figure 1 in 2000 the introduction is based on this example. This example is extracted from the DS-Qwen14B model 2001 on the sample 39 of the MATH500 dataset, using a seed of 42. 2002 Step 1 - Okay, so I need to find the greatest common factor (GCF) of 6432 and 132, and then increase 2003 that GCF by 11. Hmm, let me remember how to find the GCF of two numbers. I think the best way is to 2004 use the Euclidean algorithm because it's systematic and works well for larger numbers.\n\n Step 2 - First, let me write down the two numbers: 6432 and 132. Since 6432 is much larger than 132, 2006 I'll start by dividing 6432 by 132 and find the remainder. Then, I'll repeat the process with 132 and that 2007 remainder until the remainder is zero. The last non-zero remainder will be the GCF.\n\n Step 3 - So, let's divide 6432 by 132. Hmm, I need to figure out how many times 132 goes into 6432. 2008 Maybe I can simplify this division. Let me see, 132 times 50 is 6600, which is more than 6432. So, it 2009 must be a little less than 50. Let me try 48 times 132. 48 times 100 is 4800, 48 times 32 is 1536, so adding 2010 those together, 4800 + 1536 is 6336. Hmm, 6336 is less than 6432. So, 132 times 48 is 6336. Subtracting 2011 that from 6432: 6432 - 6336 is 96. So, the remainder is 96.  $\n$ 2012 Step 4 - Alright, now I need to find the GCF of 132 and 96. Let's apply the same method. Divide 132 by 96. 96 goes into 132 once, which is 96. Subtracting, 132 - 96 is 36. So, the remainder is 36.\n\n 2013 Step 5 - Next step, find the GCF of 96 and 36. Dividing 96 by 36. 36 times 2 is 72, which is less than 96. 2014 Subtracting 72 from 96 gives 24. So, the remainder is  $24.\n\$ 2015 Step 6 - Now, find the GCF of 36 and 24. Dividing 36 by 24, it goes once with a remainder. 36 - 24 is 12. 2016 So, remainder is  $12.\n\n$ 2017 Step 7 - Continuing, find the GCF of 24 and 12. Dividing 24 by 12, it goes exactly 2 times with no remainder. Since the remainder is zero, the last non-zero remainder is 12. Therefore, the GCF of 6432 2018 and 132 is 12.\ $\mathbf{n}$ \ $\mathbf{n}$ 2019 Step 8 - Wait, let me double-check that to make sure I didn't make a mistake. Starting from the top: 6432 2020 divided by 132 is 48 with a remainder of 96. Then, 132 divided by 96 is 1 with a remainder of 36. Then, 2021 96 divided by 36 is 2 with a remainder of 24. Then, 36 divided by 24 is 1 with a remainder of 12. Then, 2022 24 divided by 12 is 2 with no remainder. Yep, that seems right. So, the GCF is indeed 12.\n\n Step 9 - Now, the problem says to increase this GCF by 11. So, 12 plus 11 is... 23. Is that right? 12 + 11 2023 is 23. Hmm, okay, that seems straightforward.  $\n$ 2024 Step 10 - Just to be thorough, let me verify the GCF another way. Maybe by prime factorization. Let's 2025 factor both numbers into primes.  $\n$ 2026 Step 11 - Starting with 132. 132 is even, so divide by 2: 132/2 = 66. 66 is also even, so divide by 2: 66/33. 33 is divisible by 3: 33/3 = 11.11 is a prime. So, the prime factors of 132 are  $2^2 * 3^1 * 11^1 \cdot \mathbf{n} \cdot \mathbf{n}$ 2027 **Step 12 -** Now, factorizing 6432. Let's see, 6432 is even, so divide by 2: 6432 / 2 = 3216. 3216 is even, 2028 divide by 2: 3216/2 = 1608. 1608 is even, divide by 2: 1608/2 = 804. 804 is even, divide by 2: 804/22029

= 402. 402 is even, divide by 2: 402 / 2 = 201. Now, 201 is odd. Let's check divisibility by 3: 2 + 0 + 1 = 3, which is divisible by 3. So, 201 / 3 = 67. 67 is a prime number. So, putting it all together, the prime factors of 6432 are  $2^5 * 3^1 * 67^1$ .\n\n

Step 13 - Now, to find the GCF, we take the lowest power of the common prime factors. The common primes between 6432 and 132 are 2 and 3. In 132, 2 is squared, and in 6432, 2 is to the fifth power. So, we take the lower one, which is  $2^2$ . For 3, both have  $3^1$ . So, GCF is  $2^2 * 3^1 = 4 * 3 = 12$ . Yep, that

matches my earlier result.\n\n

Step 14 - 'Therefore, the GCF is definitely 12. Adding 11 to that gives 23. So, the result is 23.\n\n'

Step 15 - I think that's solid. I don't see any mistakes in my calculations, and both methods gave me the same GCF. So, increasing it by 11 is straightforward. 12 + 11 is 23.\n\n

Step 16 - '\*\*Final Answer\*\* The result is boxed{23}. i/think; To find the greatest common factor (GCF) of 6432 and 132, we use the Euclidean algorithm: 1. Divide 6432 by 132, which gives a quotient of 48 and a remainder of 96. 2. Divide 132 by 96, which gives a quotient of 1 and a remainder of 36. 3. Divide 96 by 36, which gives a quotient of 2 and a remainder of 24. 4. Divide 36 by 24, which gives a quotient of 1 and a remainder of 12. 5. Divide 24 by 12, which gives a quotient of 2 and a remainder of 0.\n\n

Step 17 - The last non-zero remainder is 12, so the GCF of 6432 and 132 is 12.\n\n

Step 18 Next we increase this GCF by 11, 12, 11 = 22. Thus, the result is beyond (22)

**Step 18 -** Next, we increase this GCF by 11: 12 + 11 = 23. Thus, the result is boxed  $\{23\}$ .

#### Legend:

2030

2031

2032

2033

2034

2035

2036

2037

2038

2039

2040

2041

2042

2043 2044

2045

2046

2047

2048

2049 2050 Definition Recall, Problem Re-statement, Heuristics / Intuition, Formula Substitution, Verification, Final Conclusion

Figure 37: Step annotation as per GPT-4o-mini using our step-tags taxonomy - Deepseek-R1-Qwen14b-Distil on MATH500 sample 39.

#### J BASELINES

2052

20532054

2055

2056 2057

2073

2074

2075

207620772078

2079

20802081

2082

2083 2084 2085

2086

2088 2089

2090

20912092

2093

20942095

2096

2097

20982099

2100

2101

2102 2103 2104

2105

**Ideal Early-Stopping**  $\mathcal{IES}$  **Algorithm.** Algorithm 3 presents the Ideal Early-Stopping  $\mathcal{IES}$  baseline algorithm.

#### **Algorithm 3** Ideal Early-Stopping $\mathcal{IES}$

2058 **Require:** Prompt x; reasoning delimiter  $\alpha \in V$ ; minimal step size  $k \in \mathbb{N}$ ; Ground-Truth Answer  $y_{\text{gold}}$ ; Answer checker  $\beta$ (prediction, ground\_truth); Reasoning Language Model  $\mathcal{M}$ ; tokenizer  $\mathcal{T}$ ; EOS token  $\gamma$ ; 2060 1:  $y \leftarrow \mathcal{T}(x)$  ▶ Tokenize the input 2061 2:  $S_{running} \leftarrow [];$ ▷ Initialize output 3:  $t \leftarrow 0$ 2062 4:  $b \leftarrow True$ ▶ Initialize stopping criteria 2063 5: while b do 2064 Generate step  $s_i$  using  $\mathcal{M}$ ,  $\alpha$ , where  $|s_i| > k$ 2065 7: if  $\beta(y, y_{\text{gold}})$  then  $b \leftarrow False$ 2066 8: 9: else 10: Continue the generation 2068 11: end if 2069 12:  $t \leftarrow t + 1$ 13: end while 2071 14: return y

**Prompt engineering**  $\mathcal{P}_{guided}$ . Figure 38 presents the different prompt variations as baseline that we defined.

```
Prompt Baselines \mathcal{P}_{guided}
  User Prompt - \mathcal{P}_{\text{user}}^{(0)}
  User Prompt: Please do not reason extensively, be succinct, and put your final answer within boxed { }. {question}
  System Prompt \mathcal{P}_{\text{system}}^{(0)}
  System Prompt: Respond concisely and confidently. Skip validations and over-verification steps
  User Prompt: {question}
  System Prompt: Respond concisely and confidently. Skip validations and over-verification steps. Here is an examples: Example 1: {FS_1}
  User Prompt: {question}
  System Prompt: Respond concisely and confidently. Skip validations and over-verification steps. Here are some examples: Example 1: {FS_1} Example 2:
   {FS_2} Example 3: {FS_3}
  User Prompt: {question}
  Example 1 - Verification step
  Wait, let me double-check. If I plug in x=-3 into the denominator, (-3)^2+(-3)-6=9-3-6=0. Yep, that works. For x=2:
  2^2 + 2 - 6 = 4 + 2 - 6 = 0. Correct. So both roots are valid.
  Example 2 - Verification step
  Therefore, the graph of y=2(x^2+x-6) has vertical asymptotes at x=-3 and x=2, so that's two vertical asymptotes. I don't think there's any
  chance that I made a mistake here, but maybe I should check by graphing the function or plugging in values close to -3 and 2 to see if the function does go
  Example 3 - Verification step
  Another test with n=3. Let's compute manually. All non-empty subsets: Single elements: \{1\}, \{2\}, \{3\} with sums 1,2,3. Pairs: \{1,2\} \rightarrow 2-1=1 \{1,3\} \rightarrow 3-1=2 \{2,3\} \rightarrow 3-2=1. Triple: \{1,2,3\} \rightarrow 3-2+1=2. Total sum: 1+2+3+1+2+1+2=12. Using the formula:
  contributions from each k: k=3: 3*2^2*1=3*4=12. k=1 and k=2 contribute 0. So total sum 12, which matches.
```

Figure 38: Prompt baselines

#### K CALIBRATION OF THE ST-ES CRITERIA

Figures 39, 40 and 41 present the number of tokens vs. accuracy of every tag-types with values of threshold ranging from 0 to 20, for the DS-Qwen14B and QwQ-32B models on our train MATH500 and GSM8K datasets using the synthetic tags, respectively. Constraints selected using our methodology for each model and dataset are reported in Table 7.

Model	Dataset	Level	Tag / Constraints						
			ST-ES (95%)	ST-ES (90%)	ST-ES (85%)	ST-ES (Router)			
		1	Context Repetition / 0	Context Repetition / 0	Self-Talk / 0	Self-Talk / 0			
		2	Final Answer / 0	Symbolic Transformation / 2	Self-Talk / 0	Self-Talk / 0			
DOLL OD	MATH500	3	Interpretation / 1	Intuition / 1	Exploration / 3	Intuition / 2			
DS-Llama8B		4	Context Repetition / 1	Context Repetition / 1	Intuition / 1	Intuition / 2			
		5	Exploration / 3	Exploration / 3	Edge Case / 1	Intuition / 2			
	GSM8K	-	Verification / 3	Verification / 1	Verification / 0	-			
	1		Self-Talk / 0	Verification / 5	Exploration / 0	Self-Talk / 0			
		2	Exploration / 3	Self-Talk / 0	Exploration / 1	Self-Talk / 0			
DC 014D	MATH500	3	Edge Case / 2	Edge Case / 2	Exploration / 1	Exploration / 3			
DS-Qwen14B		4	Exploration / 6	Exploration / 4	Intuition / 2	Exploration / 3			
		5	Intuition / 6	Self-Talk / 2	Exploration / 4	Exploration / 3			
	GSM8K	-	Verification / 0	Formula Substitution / 4	ormula Substitution / 4 Formula Substitution / 3				
		1	Exploration / 2	Formula Substitution / 2	Verification / 4	Exploration / 2			
		2	Final Answer / 2	Intuition / 1	Exploration / 1	Exploration / 2			
OO 22B	MATH500	3	Verification / 6	Verification / 5	Verification / 4	Intuition / 2			
QwQ-32B		4	Intuition / 3	Final Answer / 4	Exploration / 2	Intuition / 2			
		5	Self-Talk / 3	Self-Talk / 1	Self-Talk / 1	Intuition / 2			
	GSM8K	_	Intuition / 2	Intuition / 1	Exploration / 0	=			

Table 7: Overview of tag/constraints determined using the training datasets, for each models.

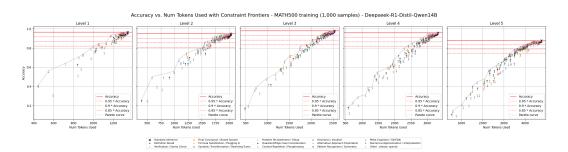


Figure 39: Early-Stopping selection using Pareto - Deepseek-R1-Qwen14B-Distil on MATH500 train (1,000 samples)

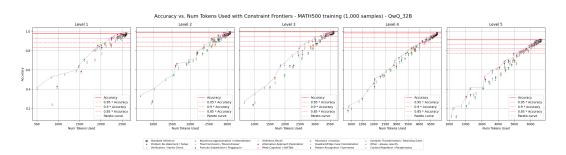


Figure 40: Early-Stopping selection using Pareto - QwQ-32B on MATH500 train (1,000 samples)

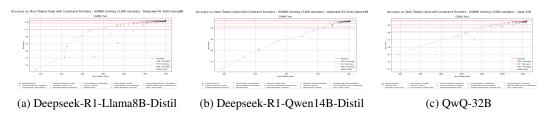


Figure 41: Early-Stopping selection using Pareto - GSM8K train (3,000 samples)

### L LLM-ROUTER: PROMPT COMPLEXITY CLASSIFICATION

Table 8 report the classifier performance of the LLM-Router module between classes  $\{1,2\}$  and  $\{3,4,5\}$  of the MATH500 dataset.

Metric	Validation	Test
Micro-F1	0.785	0.784
Macro-F1	0.734	0.739

Table 8: Micro-F1 and Macro-F1 scores for validation and test sets

#### M STEP-TAGGING PERFORMANCE

Figures 42 and 43 presents the performance of the binary step-taggers trained on the training traces of the DS-Qwen14B and QwQ-32B, respectively.



Figure 42: Step-Tagger performance - DS-Qwen14B

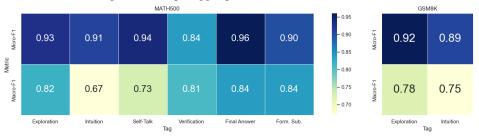


Figure 43: Step-Tagger performance - QwQ-32B

#### N IDEAL-EARLY-STOPPING

Figure 44 shows the number of steps and the accuracy of the standard vs. Ideal Early-Stopping  $\mathcal{IES}$  criteria. Results are averaged over the 5 seeds, on the MATH500 test dataset.

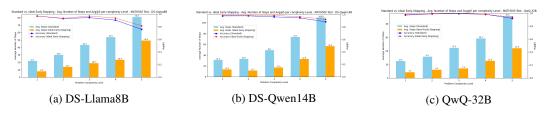


Figure 44: Standard vs.  $\mathcal{IES}$  baseline - number of steps and Avg@5 across the 5 complexity levels of MATH500 test dataset - Results averaged over the 5 seeds

#### O PERFORMANCE OF THE ST-ES CRITERIA

Table 9 reports all the token-usage, the proportion of saved number of tokens, the Avg@5, the Pass@5 and the Cons@5 for all configurations. Results are averaged over the 5 seeds we used. We also show in Figure 45 and 46 the average token count against the Pass@5 and Con@5, respectively, for the three LRMs on the MATH500 and GSM8K datasets.

Model	Config.		I	MATH500					GSM8K		
		# Tokens	Saved (%)	Avg@5	Pass@5	Cons@5	# Tokens	Saved (%)	Avg@5	Pass@5	Cons@5
	Standard	3655.0	-	0.878	0.970	0.726	958.3	-	0.829	0.943	0.651
	Basel. $\mathcal{IES}$	1916.6	47.56	0.911	0.980	0.780	385.3	59.79	0.847	0.952	0.726
	Basel. $\mathcal{P}_{user}^{(0)}$	2989.6	18.21	0.866	0.952	0.722	525.8	45.13	0.771	0.917	0.579
DS-8B	Basel. $\mathcal{P}_{\text{system}}^{(0)}$	2634.4	27.92	0.817	0.960	0.592	456.9	52.32	0.763	0.895	0.574
	Basel. $\mathcal{P}_{\text{system}}^{(1)}$	2139.5	41.46	0.782	0.942	0.526	560.8	41.48	0.754	0.914	0.537
	Basel. $\mathcal{P}_{\text{system}}^{(3)}$	2565.3	29.81	0.789	0.952	0.540	830.5	13.34	0.748	0.904	0.541
	ST-ES (95%)	3260.5	10.79	0.883	0.972	0.730	673.8	29.69	0.818	0.933	0.663
	ST-ES (90%)	2949.3	19.31	0.859	0.964	0.666	568.5	40.67	0.799	0.931	0.604
	ST-ES (85%)	2413.9	33.95	0.801	0.940	0.556	492.1	48.65	0.745	0.923	0.474
	ST-ES Router	2656.2	27.33	0.848	0.956	0.686					
	Standard	3388.8	-	0.923	0.980	0.836	662.9	-	0.910	0.952	0.843
	Basel. $\mathcal{IES}$	1655.9	51.14	0.950	0.990	0.884	316.5	52.26	0.931	0.971	0.871
	Basel. $\mathcal{P}_{user}^{(0)}$	2691.5	20.58	0.933	0.982	0.834	505.1	23.80	0.856	0.956	0.662
DS-14B	Basel. $\mathcal{P}_{\text{system}}^{(0)}$	2346.2	30.77	0.886	0.966	0.754	470.9	28.96	0.873	0.949	0.710
	Basel. $\mathcal{P}_{\text{system}}^{(1)}$	2211.4	34.74	0.873	0.974	0.708	566.5	14.54	0.838	0.952	0.629
	Basel. $\mathcal{P}_{\text{system}}^{(3)}$	2535.0	25.19	0.879	0.968	0.748	839.6	-26.65	0.841	0.952	0.631
	ST-ES (95%)	3113.2	8.13	0.923	0.980	0.824	480.0	27.59	0.884	0.951	0.763
	ST-ES (90%)	2989.7	11.78	0.906	0.976	0.794	497.9	24.89	0.838	0.940	0.656
	ST-ES (85%)	2330.2	31.24	0.841	0.966	0.670	452.5	31.74	0.754	0.918	0.528
	ST-ES Router	2545.4	24.89	0.870	0.968	0.734	2075 7		0.052	0.065	0.024
	Standard	4475.3	-	0.954	0.984	0.898	2075.7	-	0.953	0.965	0.934
	Basel. $\mathcal{IES}$	2213.2	50.55	0.970	0.992	0.940	842.9	59.39	0.976	0.986	0.963
	Basel. $\mathcal{P}_{user}^{(0)}$	2908.8	35.00	0.955	0.986	0.916	988.0	52.40	0.952	0.968	0.937
QwQ-32B	Basel. $\mathcal{P}_{\text{system}}^{(0)}$	3201.1	28.47	0.932	0.976	0.852	833.3	59.85	0.940	0.974	0.869
	Basel. $\mathcal{P}_{\text{system}}^{(1)}$	3182.4	28.89	0.925	0.974	0.856	871.2	58.02	0.943	0.975	0.876
	Basel. $\mathcal{P}_{\text{system}}^{(3)}$	3665.5	18.09	0.926	0.974	0.858	1387.3	33.16	0.935	0.974	0.855
	ST-ES (95%)	3679.4	17.78	0.921	0.980	0.786	1608.0	22.53	0.945	0.968	0.909
	ST-ES (90%)	3459.6	22.69	0.903	0.972	0.768	1506.3	27.43	0.935	0.967	0.888
	ST-ES (85%)	3218.3	28.09	0.878	0.976	0.690	1368.7	34.06	0.929	0.967	0.855
	ST-ES Router	3623.5	19.03	0.904	0.982	0.714					

Table 9: Performance of Step-Tagging Early stopping - 5 seeds (40, 41, 42, 43, 44)

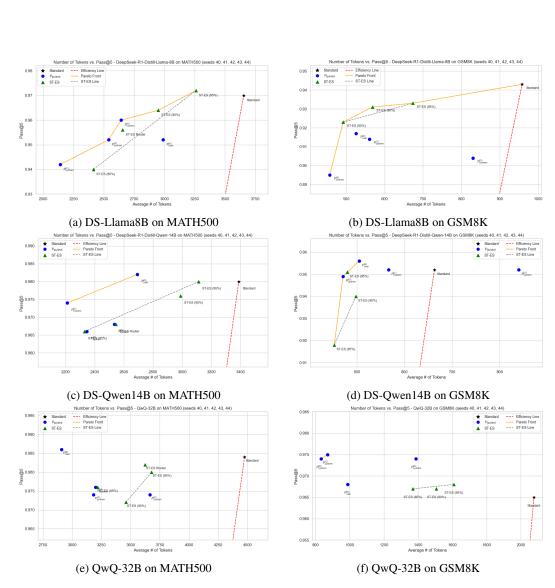


Figure 45: Number of Tokens vs. Pass@5 -  $\mathcal{P}_{guided}$  Baselines and Step-Tagging Early-Stopping (ST-ES) criteria

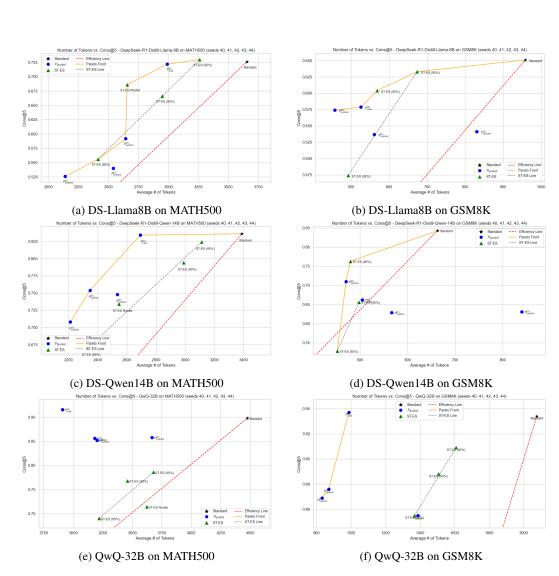


Figure 46: Number of Tokens vs. Cons@5 -  $\mathcal{P}_{guided}$  Baselines and Step-Tagging Early-Stopping (ST-ES) criteria