

D-ProtoCoT: Prototype-Based Path Selection for Chain-of-Thought Reasoning

Anonymous ACL submission

Abstract

Chain-of-Thought (CoT) prompting enables large language models to solve complex reasoning tasks by generating explicit multi-step reasoning paths. To improve robustness, self-consistency aggregates multiple sampled CoT paths via majority voting over final answers. However, this strategy operates purely at the answer level and implicitly assumes uniform reasoning quality, overlooking substantial variation among reasoning paths that yield the same answer.

In this work, we propose **D-ProtoCoT**, an inference-time framework for selecting high-quality reasoning paths based on representation-level alignment. D-ProtoCoT first constructs a contrastively aligned embedding space using weak supervision from gold answer matching, in which representations of correct reasoning paths are encouraged to cluster while incorrect ones are separated. At inference time, a dynamic reasoning prototype is formed by aggregating multiple sampled paths in this space, and the reasoning path most aligned with the prototype is selected.

D-ProtoCoT does not modify the underlying language model and requires no additional annotation beyond gold answers for representation alignment. Experiments on CommonsenseQA, GSM8K, and StrategyQA with LLaMA-3.1-8B-Instruct and Qwen3-8B show that D-ProtoCoT consistently outperforms self-consistency across most settings, achieving up to 23.6% absolute improvement on StrategyQA. Further analysis demonstrates that representation alignment is essential, as naïve centroid-based selection with frozen embeddings yields substantially inferior performance, indicating that semantic alignment provides a more reliable signal than answer frequency for reasoning path selection.

1 Introduction

Large language models (LLMs) have shown strong performance on multi-step reasoning tasks when

guided by Chain-of-Thought (CoT) (Wei et al., 2023) prompting, which encourages models to generate explicit intermediate reasoning steps before producing a final answer. To further improve robustness, self-consistency samples multiple CoT reasoning paths and aggregates their final answers via majority voting. This combination has become a standard inference strategy for reasoning benchmarks such as GSM8K, CommonsenseQA, and StrategyQA.

Despite its effectiveness, self-consistency operates solely at the answer level. It implicitly assumes that all reasoning paths leading to the same answer are equally reliable, and that answer frequency is a sufficient proxy for reasoning quality. However, in practice, different reasoning paths may arrive at the same answer through substantially different logical processes, including flawed or internally inconsistent reasoning. As a result, answer-level agreement often fails to reflect the true quality of reasoning.

This limitation naturally raises the question of whether reasoning paths can be evaluated and compared directly, rather than aggregated only through their final answers. A straightforward idea is to embed reasoning chains into a semantic space and perform aggregation using simple statistics such as centroids. However, empirical analysis shows that such naïve semantic aggregation is insufficient. Even with strong pretrained encoders, correct and incorrect reasoning paths are often entangled in the embedding space, as surface-level semantic similarity does not reliably capture logical validity.

This work proposes **D-ProtoCoT** (Dynamic Prototype-based Chain-of-Thought), an inference-time framework for reasoning path selection that operates in a contrastively aligned representation space. The core idea of D-ProtoCoT is to explicitly align reasoning representations such that correct reasoning paths are pulled closer to their corresponding questions, while incorrect paths are

086 pushed away. This alignment is achieved by fine-
087 tuning a lightweight encoder with a contrastive ob-
088 jective derived from gold answer matching, without
089 modifying the underlying language model.

090 Building on this aligned space, D-ProtoCoT con-
091 structs a *dynamic reasoning prototype* for each
092 question by aggregating multiple sampled reason-
093 ing paths with similarity-based weighting. Rather
094 than relying on majority voting, the framework se-
095 lects the reasoning path that is most aligned with
096 the prototype, thereby prioritizing internally co-
097 herent reasoning trajectories. All selection is per-
098 formed at inference time, and no additional supervi-
099 sion or model updates are required beyond encoder
100 alignment.

101 D-ProtoCoT is evaluated on GSM8K, Common-
102 senseQA, and StrategyQA using two represen-
103 tative instruction-tuned LLMs: LLaMA-3.1-8B-
104 Instruct and Qwen3-8B. Experimental results show
105 that D-ProtoCoT consistently outperforms standard
106 CoT decoding and self-consistency across tasks
107 and model families, with particularly large gains
108 on reasoning-intensive benchmarks. Further abla-
109 tion studies and qualitative analyses demonstrate
110 that representation alignment is critical, as naïve
111 centroid-based selection with frozen embeddings
112 yields substantially inferior performance.

113 The main contributions of this work are summa-
114 rized as follows:

- 115 • D-ProtoCoT is proposed as an inference-time
116 framework for selecting chain-of-thought reason-
117 ing paths based on representation-level
118 alignment rather than answer-level aggrega-
119 tion.
- 120 • A contrastive representation alignment mech-
121 anism is introduced to enable the construc-
122 tion of dynamic reasoning prototypes, with-
123 out modifying or fine-tuning the underlying
124 language model.
- 125 • The effectiveness of D-ProtoCoT is demon-
126 strated across multiple reasoning benchmarks
127 and model families, with analysis highlighting
128 its advantage over naïve semantic aggregation
129 methods.

130 2 Related Work

131 2.1 Chain-of-Thought Reasoning and 132 Self-Consistency

133 Chain-of-Thought (CoT) prompting enables large
134 language models to solve multi-step reasoning

135 problems by explicitly generating intermediate rea-
136 soning steps (Wei et al., 2023). Building on this
137 paradigm, self-consistency improves robustness by
138 sampling multiple reasoning paths and aggregat-
139 ing final answers through majority voting (Wang
140 et al., 2023). This strategy has become a standard
141 inference-time technique for reasoning benchmarks
142 such as GSM8K, CommonsenseQA, and Strate-
143 gyQA.

144 Despite its empirical success, self-consistency
145 operates purely at the answer level and implicitly
146 assumes that reasoning paths leading to the same
147 answer are equally reliable. Several studies have
148 noted that majority voting may amplify spurious
149 reasoning patterns and fail when correct reason-
150 ing paths are not numerically dominant (Fu et al.,
151 2023; Zhou et al., 2022). D-ProtoCoT departs from
152 answer-level aggregation by directly operating on
153 reasoning representations, enabling path selection
154 based on semantic and structural alignment rather
155 than answer frequency.

156 2.2 Reasoning Path Selection and Verification

157 A related line of work focuses on selecting or
158 validating reasoning paths using external signals.
159 Verifier-based approaches train separate models
160 to assess the correctness of generated reasoning
161 chains or intermediate steps (Cobbe et al., 2021;
162 Lightman et al., 2023). Similarly, reward-model-
163 based methods employ reinforcement learning or
164 preference optimization to encourage correct rea-
165 soning trajectories (Ouyang et al., 2022; Bai et al.,
166 2022).

167 While effective, these approaches typically re-
168 quire additional supervision, costly human anno-
169 tation, or task-specific training, and often modify
170 the base language model. In contrast, D-ProtoCoT
171 is designed as a lightweight inference-time frame-
172 work. It does not attempt to explicitly judge log-
173 ical correctness, nor does it update the reasoning
174 model. Instead, it relies on representation-level
175 alignment to identify reasoning paths that are se-
176 mantically consistent with high-quality reasoning
177 patterns, offering a complementary alternative to
178 explicit verification.

179 2.3 Contrastive Learning for Reasoning 180 Representations

181 Contrastive learning has been widely adopted to
182 improve semantic representation quality in natural
183 language processing, particularly for sentence and
184 text embedding models (Gao et al., 2021; Gunel

et al., 2021). Recent work has explored contrastive objectives for reasoning-related tasks, including contrastive chain-of-thought prompting (Chia et al., 2023) and reasoning-aware representation learning for logical reasoning problems (Yu et al., 2020).

Unlike prior approaches that apply contrastive learning directly to language model outputs or confidence estimation, D-ProtoCoT employs contrastive alignment to explicitly shape a shared embedding space for questions and reasoning paths. Weak supervision from gold answer matching is used to distinguish correct and incorrect reasoning paths, enabling the construction of an aligned representation space without additional annotation or model modification. This aligned space serves as the foundation for prototype-based reasoning path selection.

2.4 Prototype-Based Representation Modeling

Prototype-based methods have a long history in representation learning and pattern recognition, where class prototypes or centroids are used to capture central tendencies of data distributions (Snell et al., 2017). In NLP, prototype-based representations have been applied to text classification, clustering, and semantic matching tasks (Han et al., 2018; Pan et al., 2020).

Recent studies suggest that prototype aggregation can improve robustness by mitigating noise from individual samples (Liu et al., 2022). D-ProtoCoT extends this idea to reasoning path selection by constructing a dynamic, question-specific prototype at inference time. Unlike static class prototypes, the proposed dynamic prototype aggregates multiple reasoning paths with similarity-based weighting, allowing the model to adaptively capture the dominant semantic structure of high-quality reasoning for each input instance.

3 Method

3.1 Problem Setup

Given an input question q , a large language model (LLM) generates multiple chain-of-thought (CoT) reasoning paths $\mathcal{C} = \{c_1, \dots, c_K\}$ via stochastic decoding. Each reasoning path consists of a sequence of intermediate reasoning steps followed by a final answer. The objective is to select a single reasoning path that exhibits strong internal coherence and leads to a reliable final answer.

Throughout this framework, the base language model used for reasoning path generation remains

frozen. Only a lightweight auxiliary encoder is trained for representation alignment. All components operate at inference time, enabling reasoning path selection without modifying or fine-tuning the underlying language model.

3.2 Multi-Granular Representation Extraction

Each reasoning path is encoded using a BERT-base-uncased encoder $E(\cdot)$, which is subsequently fine-tuned through contrastive representation alignment (Section 3.3). Given the tokenized input of a reasoning path, the encoder produces a sequence of hidden states from its final transformer layer.

To capture reasoning information at different granularities, representations are extracted at three levels.

Token-Level Representation. Let $\mathbf{H}_i \in \mathbb{R}^{T \times d}$ denote the final-layer hidden states for path c_i , where T is the number of tokens and d is the hidden dimension. These token embeddings provide the most fine-grained representation of the reasoning process.

Step-Level Representation. Each reasoning path c_i is segmented into M steps $\{u_{i,1}, \dots, u_{i,M}\}$ based on newline delimiters. Let $\mathcal{T}_{i,j} \subseteq \{1, \dots, T\}$ denote the set of token indices belonging to step $u_{i,j}$. The step-level representation is computed as:

$$\mathbf{s}_{i,j} = \frac{1}{|\mathcal{T}_{i,j}|} \sum_{t \in \mathcal{T}_{i,j}} \mathbf{h}_{i,t},$$

where $\mathbf{h}_{i,t} \in \mathbb{R}^d$ is the token embedding at position t .

Path-Level Representation. A single representation for the entire reasoning path is obtained by mean pooling over its step embeddings:

$$\mathbf{z}_i = \frac{1}{M} \sum_{j=1}^M \mathbf{s}_{i,j}.$$

The resulting path-level embedding $\mathbf{z}_i \in \mathbb{R}^d$ serves as the primary unit for contrastive alignment and prototype-based selection.

3.3 Contrastive Representation Alignment

Motivation. naïve semantic representations obtained from off-the-shelf encoders often fail to distinguish logically valid reasoning from superficially plausible but incorrect chains, as they are

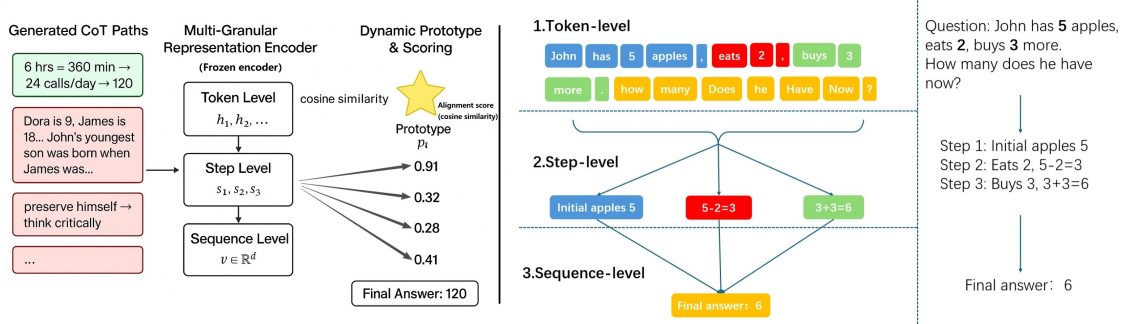


Figure 1: Overview of D-ProtoCoT. The framework encodes multiple chain-of-thought reasoning paths into a contrastively aligned representation space, constructs a dynamic reasoning prototype, and selects the path with the highest prototype alignment score at inference time.

not explicitly trained to align questions with valid reasoning processes. To address this limitation, a contrastive learning stage is introduced to shape the representation space such that correct reasoning paths are pulled closer to their corresponding questions, while incorrect paths are pushed away. This alignment ensures that downstream prototype construction operates in a space where semantic proximity correlates with reasoning quality.

Positive and Negative Path Construction. For each question q , K reasoning paths are generated and labeled as correct or incorrect based on exact match between the final answer and the ground-truth answer. Paths whose final answer matches the gold label are treated as positive samples, while all remaining paths are treated as negatives. This supervision signal is obtained automatically and requires no additional annotation.

Representation Encoding. The question q and each reasoning path c_i are encoded using the same encoder $E(\cdot)$. Let $\mathbf{z}_q = E(q)$ denote the pooled question representation, and $\mathbf{z}_i = E(c_i)$ denote the path-level representation. All encoder parameters are fine-tuned during contrastive alignment.

Contrastive Objective. An InfoNCE-style contrastive loss is employed to align question and reasoning representations:

$$\mathcal{L}_{\text{con}} = -\log \frac{\sum_{i \in \mathcal{P}} \exp(\text{sim}(\mathbf{z}_q, \mathbf{z}_i)/\tau)}{\sum_{j \in \mathcal{P} \cup \mathcal{N}} \exp(\text{sim}(\mathbf{z}_q, \mathbf{z}_j)/\tau)},$$

where \mathcal{P} and \mathcal{N} denote the sets of correct and incorrect reasoning paths, $\text{sim}(\cdot, \cdot)$ is cosine similarity, and τ is a temperature hyperparameter.

3.4 Dynamic Prototype Construction

After contrastive alignment, reasoning paths are embedded into a representation space where correct

paths are encouraged to align more closely with their corresponding questions. Based on this space, a dynamic reasoning prototype is constructed for each question at inference time.

Given a question q and its K reasoning paths with embeddings $\{\mathbf{z}_1, \dots, \mathbf{z}_K\}$, the dynamic prototype \mathbf{p}_q is defined as a weighted combination of all path representations:

$$\mathbf{p}_q = \sum_{i=1}^K w_i \mathbf{z}_i. \quad 319$$

The weight w_i reflects the alignment between a reasoning path and the question representation:

$$w_i = \frac{\exp(\text{sim}(\mathbf{z}_q, \mathbf{z}_i))}{\sum_{j=1}^K \exp(\text{sim}(\mathbf{z}_q, \mathbf{z}_j))}. \quad 322$$

This similarity-based weighting assigns higher influence to paths that are more aligned with the question, while retaining information from all sampled paths.

3.5 Prototype-Based Alignment and Path Selection

Given the dynamic prototype \mathbf{p}_q , an alignment score is computed for each reasoning path using cosine similarity:

$$s_i = \text{sim}(\mathbf{z}_i, \mathbf{p}_q). \quad 332$$

The final reasoning path is selected by maximizing the alignment score:

$$c^* = \arg \max_{c_i \in \mathcal{C}} s_i. \quad 335$$

The final answer is extracted from the selected top-1 reasoning path c^* .

338	3.6 Inference-Time Properties		
339	D-ProtoCoT operates entirely at inference time.		
340	No parameters are updated during reasoning path		
341	selection, and the framework introduces minimal		
342	computational overhead beyond reasoning genera-		
343	tion and embedding computation. The method is		
344	compatible with any language model capable of		
345	generating chain-of-thought reasoning paths.		
346	4 Experiments		
347	This section evaluates the effectiveness of D-		
348	ProtoCoT on three representative multi-step reason-		
349	ing benchmarks: GSM8K, CommonsenseQA, and		
350	StrategyQA. These datasets emphasize arithmetic		
351	reasoning, commonsense inference, and implicit		
352	semantic reasoning respectively, enabling a com-		
353	prehensive assessment of reasoning path selection		
354	under diverse reasoning characteristics.		
355	4.1 Benchmarks and Backbone Models		
356	D-ProtoCoT is evaluated using two instruction-		
357	tuned large language models: LLaMA-3.1-8B-		
358	Instruct and Qwen3-8B. These models differ sub-		
359	stantially in pretraining data, architectural design,		
360	and instruction-following behavior, enabling an as-		
361	essment of robustness and generalization across		
362	heterogeneous model families.		
363	4.2 Reasoning Path Generation		
364	For each input question, 10 independent chain-		
365	of-thought reasoning paths are sampled using		
366	temperature-controlled decoding. The same set of		
367	sampled reasoning paths is shared across all com-		
368	pared methods to ensure fair comparison. Final		
369	answers are extracted from each reasoning path		
370	following the official evaluation protocols of each		
371	benchmark.		
372	4.3 Baselines		
373	D-ProtoCoT is compared with standard inference-		
374	time reasoning baselines:		
375	Standard CoT selects the answer from a single		
376	sampled reasoning path.		
377	Self-Consistency aggregates multiple reasoning		
378	paths via majority voting over final answers.		
379	C-CoT (Chia et al., 2023) selects reasoning paths		
380	based on confidence estimation.		
	4.4 Implementation Details		381
	Reasoning paths are encoded using a pretrained		382
	BERT encoder. Token-level representations are		383
	extracted from the final transformer layer . Step-		384
	level representations are obtained by mean-pooling		385
	token embeddings within each reasoning step,		386
	while path-level representations are computed by		387
	mean-pooling all step embeddings of a reasoning		388
	path, following the representation definitions in		389
	Section 3.2.		390
	For D-ProtoCoT, the encoder is contrastively		391
	aligned using reasoning paths with weak supervi-		392
	sion derived from gold answer matching (exact		393
	match), which identifies positive and negative rea-		394
	soning paths. The contrastive objective fine-tunes		395
	the entire encoder, rather than a separate projection		396
	head. During inference, all reasoning paths are en-		397
	coded using the aligned encoder, and a dynamic		398
	prototype is constructed via similarity-weighted		399
	aggregation of all paths, where weights are deter-		400
	mined by their alignment with the question repre-		401
	sentation.		402
	All experiments are conducted on a single Nvidia		403
	RTX 4090 GPU. All methods use identical genera-		404
	tion settings, the same number of sampled reason-		405
	ing paths, and the same evaluation protocol. No		406
	additional fine-tuning is performed on the back-		407
	bone LLMs.		408
	4.5 Main Results		409
	Table 1 shows that D-ProtoCoT consistently out-		410
	performs standard CoT decoding and achieves		411
	competitive or superior results compared to self-		412
	consistency across most settings. Improvements		413
	are particularly pronounced on tasks where cor-		414
	rect reasoning paths do not dominate numerically,		415
	indicating that representation-level reasoning align-		416
	ment provides a complementary signal beyond an-		417
	swer frequency.		418
	4.6 Inference Cost		419
	D-ProtoCoT introduces modest inference overhead		420
	due to multi-path encoding and prototype scoring.		421
	However, the dominant cost remains in LLM genera-		422
	tion. The additional representation encoding and		423
	similarity computation incur negligible latency in		424
	practice. Empirically, D-ProtoCoT exhibits compa-		425
	erable inference time to self-consistency when using		426
	the same number of sampled paths.		427

Table 1: Main results (%) on CommonsenseQA, GSM8K, and StrategyQA. Results are reported using 10 sampled reasoning paths.

Method	LLaMA-3.1-8B-Instruct			Qwen3-8B		
	CSQA	GSM8K	StrategyQA	CSQA	GSM8K	StrategyQA
Standard CoT	68.23	43.80	68.60	75.60	90.40	66.80
Self-Consistency	77.40	63.20	62.60	86.98	90.65	76.60
C-CoT	64.40	87.00	64.40	68.20	75.60	62.70
D-ProtoCoT	79.80	78.20	86.20	87.71	94.15	72.60

5 Analysis

D-ProtoCoT is designed as an inference-time reasoning path selection framework, rather than a training-time modification of language models. This design choice prioritizes modularity and ease of integration, allowing the method to be applied to a wide range of pretrained models without additional fine-tuning or supervision.

Unlike verifier-based or reward-model-based approaches, D-ProtoCoT does not explicitly judge logical correctness. Instead, it relies on representation-level alignment to identify reasoning paths that are structurally and semantically consistent with high-quality reasoning patterns. This makes D-ProtoCoT a lightweight middleware situated between naïve aggregation strategies and heavy verification pipelines.

5.1 Why naïve Centroid Selection Fails

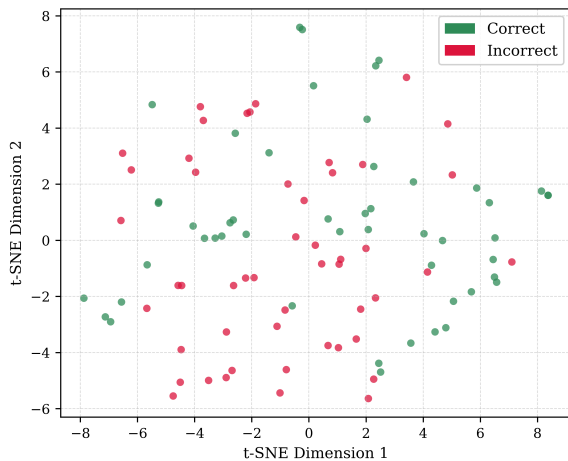


Figure 2: t-SNE visualization of BERT embeddings from 100 chain-of-thought paths generated for a StrategyQA example. Green points indicate correct reasoning paths, while red points indicate incorrect ones. Strong overlap suggests that surface-level semantic similarity alone is insufficient to distinguish reasoning quality.

A natural baseline for reasoning path selection

Table 2: Comparison on GSM8K (LLaMA-3.1-8B-Instruct) between naïve centroid-based selection and D-ProtoCoT.

Method	Accuracy (%)
Raw BERT + Centroid	21.00
D-ProtoCoT (ours)	78.20

is to compute embeddings for each reasoning path and select the one closest to the centroid of all paths. As shown in Figure 2, frozen BERT embeddings of correct and incorrect reasoning paths are heavily intermixed, indicating that surface-level semantic similarity is insufficient to separate logically valid reasoning from flawed ones. This limitation is further confirmed quantitatively in Table 2, where naïve centroid-based selection with raw BERT embeddings performs poorly compared to D-ProtoCoT.

This result demonstrates that without representation alignment, centroid-based aggregation is vulnerable to selecting semantically similar but logically incorrect reasoning paths. In contrast, D-ProtoCoT benefits from contrastively aligned representations that better reflect reasoning coherence.

5.2 Case Study

Figure 3 illustrates a representative failure case of self-consistency. Although the majority of reasoning paths lead to the same incorrect answer, D-ProtoCoT selects a path with higher internal coherence by aligning with the reasoning prototype.

5.3 Effect of Representation Granularity

D-ProtoCoT encodes reasoning paths at multiple levels of granularity. Token-level representations capture surface variations, while path-level representations summarize global semantics but may obscure localized logical errors. Step-level representations provide an intermediate abstraction by

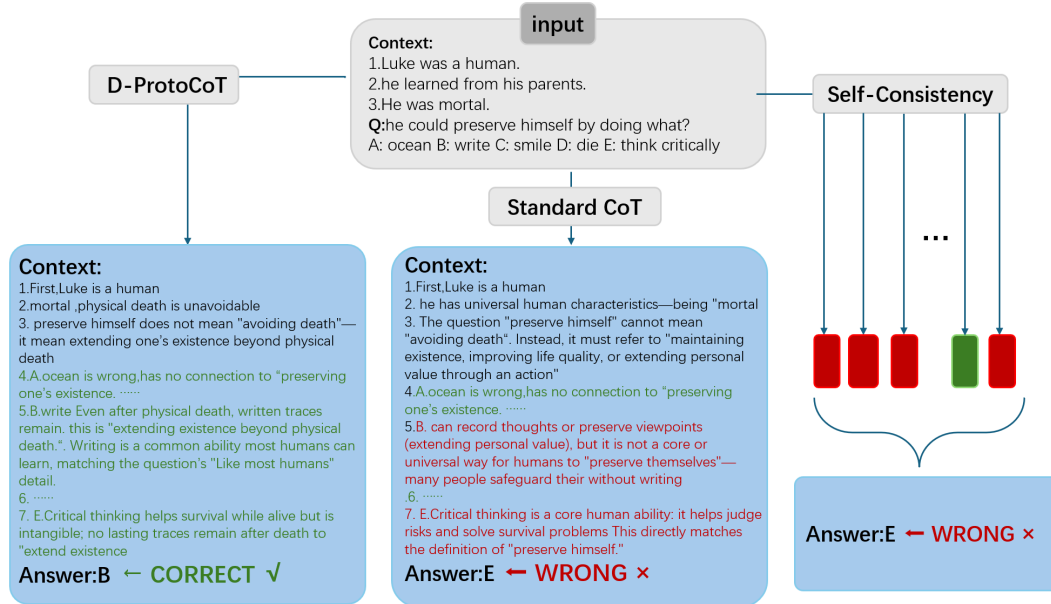


Figure 3: Case study on StrategyQA. While self-consistency selects an incorrect answer due to majority voting, D-ProtoCoT identifies a reasoning path that correctly interprets the implicit semantics of the question.

modeling semantically meaningful inference units.

Although we do not present separate quantitative experiments for each granularity, qualitative analysis suggests that step-level representations are particularly effective for multi-step reasoning tasks, where errors often arise from a single flawed intermediate step.

5.4 Impact of Backbone Model Strength

We observe that when the backbone model is a strong instruction-tuned LLM such as Qwen3-8B, self-consistency can already achieve competitive performance with a limited number of sampled reasoning paths. In this setting, the performance gap between self-consistency and D-ProtoCoT narrows.

This observation suggests that the benefit of prototype-based reasoning selection is most pronounced when reasoning path quality exhibits higher variance, or when a larger number of candidate reasoning paths are available.

5.5 Scope and Limitations

D-ProtoCoT is not intended to replace formal verification or symbolic reasoning systems. Instead, it serves as a lightweight and modular component that can be easily integrated into existing reasoning pipelines. Its effectiveness depends on the diversity and quality of sampled reasoning paths, and its benefits diminish when most paths are already coherent and correct.

6 Conclusion

This work investigates inference-time reasoning path selection for large language models beyond answer-level aggregation. To this end, D-ProtoCoT is proposed—a prototype-based framework that selects chain-of-thought paths by measuring their alignment with a dynamically constructed reasoning prototype in a contrastively aligned representation space. By prioritizing internal semantic coherence over surface-level answer agreement, the approach offers a complementary strategy to existing decoding and aggregation methods.

Empirical results across multiple reasoning benchmarks and model families demonstrate that D-ProtoCoT consistently improves performance in settings where majority voting is suboptimal, particularly when correct reasoning paths are semantically concentrated while incorrect ones are scattered. Further analysis highlights the insufficiency of naïve semantic aggregation and underscores the importance of structured representation learning for reliable reasoning path selection.

Although D-ProtoCoT does not guarantee logical correctness in all cases, it provides evidence that representation-level alignment serves as a more informative signal than answer frequency alone. Future work may explore finer-grained reasoning representations, integration with formal verification, or adaptation to settings without gold answers.

535 Limitations

536 D-ProtoCoT has several limitations that warrant
537 discussion. First, the method assumes access to
538 gold answers during the *training phase* of the aux-
539 iliary encoder for contrastive alignment (3.3). Con-
540 sequently, it is not directly applicable to fully un-
541 supervised settings where no labeled answers are
542 available, which may limit its use in certain real-
543 world deployment scenarios.

544 Second, the effectiveness of reasoning path se-
545 lection depends on the diversity and correctness
546 distribution of the sampled reasoning paths. When
547 most generated paths are incorrect or semantically
548 homogeneous (e.g., due to low decoding tempera-
549 ture), the dynamic prototype may fail to capture a
550 meaningful signal, thereby reducing potential per-
551 formance gains. This observation is consistent with
552 our empirical findings that D-ProtoCoT provides
553 smaller improvements on tasks with inherently low
554 path diversity.

555 Third, D-ProtoCoT introduces additional infer-
556 ence overhead due to the use of an auxiliary en-
557 coder for representation extraction and similarity
558 computation. Although this overhead is modest
559 compared to the cost of large language model gen-
560 eration, it still results in higher latency than vanilla
561 self-consistency. Moreover, the method may inherit
562 biases or domain mismatches from the encoder’s
563 pretraining data.

564 Finally, all experiments are conducted on
565 English-language benchmarks. The effectiveness
566 of D-ProtoCoT on low-resource languages or spe-
567 cialized domains, such as medical or legal reason-
568 ing, remains unexplored and warrants further inves-
569 tigation.

570 Acknowledgments

571 References

572 Yuntao Bai, Saurav Kadavath, Sandipan Kundu,
573 Amanda Askell, Jackson Kernion, Andy Jones, Anna
574 Chen, Anna Goldie, Azalia Mirhoseini, Cameron
575 McKinnon, Carol Chen, Catherine Olsson, Christo-
576 pher Olah, Danny Hernandez, Dawn Drain, Deep
577 Ganguli, Dustin Li, Eli Tran-Johnson, Ethan Perez,
578 and 32 others. 2022. *Constitutional ai: Harmlessness
579 from ai feedback*. *Preprint*, arXiv:2212.08073.

580 Yew Ken Chia, Guizhen Chen, Luu Anh Tuan, Soujanya
581 Poria, and Lidong Bing. 2023. *Contrastive chain-of-
582 thought prompting*. *Preprint*, arXiv:2311.09277.

583 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian,
584 Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias
585 Plappert, Jerry Tworek, Jacob Hilton, Reiichiro

Nakano, Christopher Hesse, and John Schulman. 2021. *Training verifiers to solve math word prob-
587 lems*. *Preprint*, arXiv:2110.14168. 588

Yao Fu, Hao Peng, Ashish Sabharwal, Peter Clark, and
589 Tushar Khot. 2023. *Complexity-based prompting for
590 multi-step reasoning*. *Preprint*, arXiv:2210.00720. 591

Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. *592
593 SimCSE: Simple contrastive learning of sentence em-
594 beddings*. In *Proceedings of the 2021 Conference
595 on Empirical Methods in Natural Language Process-
596 ing*, pages 6894–6910, Online and Punta Cana, Do-
597 minican Republic. Association for Computational
598 Linguistics.

Beliz Gunel, Jingfei Du, Alexis Conneau, and Veselin
599 Stoyanov. 2021. *Supervised contrastive learning for
600 pre-trained language model fine-tuning*. In *9th In-
601 ternational Conference on Learning Representations,
602 ICLR 2021, Virtual Event, Austria, May 3-7, 2021*.
603 OpenReview.net. 604

Xu Han, Hao Zhu, Pengfei Yu, Ziyun Wang, Yuan Yao,
605 Zhiyuan Liu, and Maosong Sun. 2018. *FewRel: A
606 large-scale supervised few-shot relation classification
607 dataset with state-of-the-art evaluation*. In *Proceed-
608 ings of the 2018 Conference on Empirical Methods
609 in Natural Language Processing*, pages 4803–4809,
610 Brussels, Belgium. Association for Computational
611 Linguistics. 612

Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri
613 Edwards, Bowen Baker, Teddy Lee, Jan Leike,
614 John Schulman, Ilya Sutskever, and Karl Cobbe.
615 2023. *Let’s verify step by step*. *Preprint*,
616 arXiv:2305.20050. 617

Yang Liu, Jinpeng Hu, Xiang Wan, and Tsung-Hui
618 Chang. 2022. *Learn from relation information: To-
619 wards prototype representation rectification for few-
620 shot relation extraction*. In *Findings of the Associ-
621 ation for Computational Linguistics: NAACL 2022*,
622 pages 1822–1831, Seattle, United States. Association
623 for Computational Linguistics. 624

Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Car-
625 roll L. Wainwright, Pamela Mishkin, Chong Zhang,
626 Sandhini Agarwal, Katarina Slama, Alex Ray, John
627 Schulman, Jacob Hilton, Fraser Kelton, Luke Miller,
628 Maddie Simens, Amanda Askell, Peter Welinder,
629 Paul Christiano, Jan Leike, and Ryan Lowe. 2022. *630
631 Training language models to follow instructions with
632 human feedback*. *Preprint*, arXiv:2203.02155. 632

Liangming Pan, Yuxi Xie, Yansong Feng, Tat-
633 Seng Chua, and Min-Yen Kan. 2020. *Semantic
634 graphs for generating deep questions*. *Preprint*,
635 arXiv:2004.12704. 636

Jake Snell, Kevin Swersky, and Richard Zemel. 2017. *637
638 Prototypical networks for few-shot learning*. In
639 *NeurIPS*, pages 4077–4087.

Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le,
640 Ed Chi, Sharan Narang, Aakanksha Chowdhery, and
641

642 Denny Zhou. 2023. [Self-consistency improves chain](#)
643 [of thought reasoning in language models](#). *Preprint*,
644 arXiv:2203.11171.

645 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten
646 Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and
647 Denny Zhou. 2023. [Chain-of-thought prompting elic-](#)
648 [its reasoning in large language models](#). *Preprint*,
649 arXiv:2201.11903.

650 Weihao Yu, Zihang Jiang, Yanfei Dong, and Ji-
651 asheng Feng. 2020. [Reclor: A reading comprehen-](#)
652 [sion dataset requiring logical reasoning](#). *Preprint*,
653 arXiv:2002.04326.

654 Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei,
655 Nathan Scales, Xuezhi Wang, Dale Schuurmans,
656 Olivier Bousquet, Quoc Le, and Ed Chi. 2022. [Least-](#)
657 [to-most prompting enables complex reasoning in](#)
658 [large language models](#).