

EpiCoDe: Boosting Model Performance Beyond Training with Extrapolation and Contrastive Decoding

Mingxu Tao¹, Jie Hu², Mingchuan Yang², Yunhuai Liu^{3,4},
Dongyan Zhao^{1,5}, Yansong Feng¹✉

¹Wangxuan Institute of Computer Technology, Peking University

²Research Institute of China Telecom

³School of Computer Science, Peking University

⁴Beijing Institute of Big Data Research

⁵National Key Laboratory of General Artificial Intelligence, BIGAI

{thomastao, fengyansong}@pku.edu.cn

Abstract

The remarkable performance of Large language models (LLMs) relies heavily on the availability of abundant high-quality training data. However, the high cost of acquiring annotated data often prevents models from obtaining capabilities to tackle downstream tasks. In this paper, we introduce a novel method, EpiCoDe that boosts model performance in data-scarcity scenarios without extra training. We first employ model extrapolation to enhance a finetuned model with its inferior version, and then adopt contrastive decoding to further reduce predicted errors, by comparing the logit scores given by the extrapolated and the vanilla finetuned model. Experiments across three tasks over four different LLMs show that EpiCoDe consistently outperforms existing methods with significant and robust improvement. We also propose a new theoretical framework to reveal the mechanism behind contrastive decoding in data-scarcity scenarios, which further helps us better understand the effectiveness of EpiCoDe.

1 Introduction

Large language models (LLMs) have demonstrated superior performance across a wide range of downstream tasks, making them an indispensable application in natural language processing (Yang et al., 2024; Team, 2024; DeepSeek-AI et al., 2025). Their impressive capabilities, however, are primarily dependent on the availability of large-scale, high-quality training data. The scarcity of such data, especially in specific domains, has become a major barrier to further advancing LLMs. This challenge is particularly pronounced in fields such as law and health (Huang et al., 2023; Chen et al., 2024), where access to domain-specific data is often limited due to privacy concerns, copyright restrictions, and the high cost of data collection. Furthermore, the expense of high-quality annotations may also remain a bottleneck, restricting the consistent development of model abilities.

Therefore, there has been a growing need to develop methods that can make the most use of limited data without extra training to enhance model performance in such data-scarcity scenarios. Existing efforts include two types of methods, model extrapolation (Zheng et al., 2024) and contrastive decoding (Li et al., 2023), which are developed from the perspectives of parameters and inference, respectively. Model extrapolation collects two checkpoints from different training stages. It obtains a stronger model by comparing the parameters of two checkpoints without extra training. Beyond directly editing model parameters, there is another thread of research focusing on improving the probability to output correct next tokens during inference. Li et al. (2023) hypothesize the LLMs trained on the same datasets can have similar patterns to make mistakes. Meanwhile, models of different sizes exhibit different capabilities. Thus, by comparing the logit scores predicted by a small LLM and a large LLM in the same family, contrastive decoding can narrow the performance gap and reduce their shared errors.

Previous work (Karras et al., 2024) also proposes to improve the quality of generated images by comparing the prediction of a model and its bad version trained with noisy data. These two techniques are both good practice to enhance LLMs without extra training. However, they sometimes fail to bring improvement (O’Brien and Lewis, 2023; Zheng et al., 2024). Furthermore, existing works have not explored whether they are effective in the context of data scarcity.

In this paper, we introduce EpiCoDe, a novel method that leverages model extrapolation and contrastive decoding to improve the performance of a finetuned model without further training in data-scarcity scenarios. We first obtain an extrapolated model of a finetuned LLM by employing its weakened version at early training stage as the anchor. The resulting models as well as the vanilla

finetuned model learn from the same dataset, exhibiting similar capabilities and error patterns. We then compare their logit scores and employ contrastive decoding to further optimize the prediction of next tokens. This frustratingly easy method can take the strengths of both methods to further improve model performance.

Previous efforts mainly provide empirical evidence to explain the success of contrastive decoding, without a proper analyzing tool to help us understand how contrastive decoding contributes to the improvement. In this work, we further introduce a theoretical framework to explain the mechanisms behind contrastive decoding and investigate strategies for model selection in EpiCoDe. Our key insight is that model extrapolation helps maintain model locality across finetuned models and extrapolated models, which can ensure further improvement during the contrastive decoding stage. Specifically, we argue that EpiCoDe brings more stable improvement when selecting the weak models that are closer to the extrapolated model in the parameter space.

We evaluate our method across three distinct tasks (law, math, and logical reasoning) using four LLMs from different families and with various scales. Experimental results show that EpiCoDe consistently outperforms baseline models, which employ either model extrapolation or contrastive decoding, with significant performance gains across all tasks. Furthermore, our experiments also demonstrate the validity of our theoretical framework, highlighting the importance of model locality in the context of EpiCoDe.

Our main contributions are threefold:

- We propose a frustratingly simple yet effective method, EpiCoDe, that benefits from both model extrapolation and contrastive decoding, enhancing model capabilities without additional training, particularly in data-scarcity scenarios.
- We present a theoretical framework to explain the mechanism behind EpiCoDe and the model selection strategies in contrastive decoding, which have been only discussed empirically.
- Experiments on three diverse tasks across four different base models show that our method can bring consistent improvement, significantly outperforming baseline methods.

2 Related Works

The scarcity of high-quality annotated training data often prevents us from effectively finetuning LLMs, necessitating the development of approaches that go beyond conventional training paradigms. Among others, model extrapolation and contrastive decoding are two representative threads of research falling to this type.

Model Extrapolation An intuitive idea to improve model performance is to extrapolate a model to a better position in the parameter space. However, we are unable to precisely locate the optimal point without training. A simple approximate solution is to regard model extrapolation as the reverse process of linear merging (Zheng et al., 2024). They first enhance a finetuned LLM with Reinforcement Learning with Human Feedback (RLHF, Ouyang et al., 2022). To obtain a better RLHF model, they hypothesize the RLHF model is constructed by merging the finetuned model and an unknown ultra-powerful model. As shown in the left part of Figure 1, they argue that extrapolation can alleviate the limitation of insufficient training data, resulting a stronger RLHF model.

Here, we believe model extrapolation might be a possible choice to make the most of limited training data by exploring across different checkpoints. In this paper, we extract an intermediate checkpoint during finetuning as the weak model. We then regard the model trained with more epochs as the strong model, and extrapolate it to obtain a stronger one.

Contrastive Decoding Different from model extrapolation which edits the parameters to get a new model, contrastive decoding (Li et al., 2023) combines two distinct LLMs’ predictions during inference. The fundamental hypothesis is that, on one hand, models trained on the same data may have similar patterns to make mistakes; while on the other hand, the capabilities of models differ due to model size. Li et al. (2023) proposes to collect the two probability distributions of next token predicted by a strong and a weak model in the same family, respectively. Their difference may represent the capability gap between two models without co-occurrence errors. As shown in the right part of Figure 1, contrastive decoding improves the quality of model outputs by modifying the predicted logit scores with the computed differences.

In our work, we are interested in whether con-

trastive decoding can further enhance the extrapolated models. Different from previous works (Li et al., 2023; O’Brien and Lewis, 2023) that use LLMs in the same family but with different sizes, we here adopt the finetuned model and the extrapolated model for exploration. The effectiveness of contrastive decoding in data-scarcity scenarios, e.g., enhancing a finetuned model with its early version, is still an unexplored research question.

3 Methodology

In this paper, we propose EpiCoDe, which leverages model extrapolation and contrastive decoding in data-scarcity scenarios, to boost model performance without additional training.

3.1 Preliminary

In data scarcity scenarios, the training data available for finetuning is insufficient. We can only obtain under-trained LLMs with limited capabilities. And, model checkpoints collected from the earlier stage of such training conditions can perform unsurprisingly worse.

Model Extrapolation. Following previous work (Zheng et al., 2024), we regard a strong model θ^{strong} as the linearly merged model (Wortman et al., 2022; Choshen et al., 2022) of its weakened version θ^{weak} and an ultra-strong model. Thus, we can obtain the unknown model by linear extrapolation:

$$\theta^{ep} = \theta^{strong} + \mu(\theta^{strong} - \theta^{weak}),$$

where $\mu > 0$ is a hyper-parameter.

Contrastive Decoding For the input sequence $x_{<i} : \langle x_1, \dots, x_{i-1} \rangle$, our aim is to predict the next token x_i within the vocabulary \mathcal{V} . Given the distribution $p_s(x_i|x_{<i})$ predicted by a strong model and $p_w(x_i|x_{<i})$ predicted by a weak model, previous work (Li et al., 2023) defines the *contrastive difference* as:

$$c(x_i|x_{<i}) = \log(p_s(x_i|x_{<i})) - \log(p_w(x_i|x_{<i})),$$

which excludes the co-occurrence errors of the two models.

Let $L_s(x_{<i})$ and $L_w(x_{<i}) \in \mathbb{R}^{\mathcal{V}}$ denote the unnormalized distributions (i.e., logit scores) predicted by the strong and the weak model. We can reformulate the *contrastive difference* as:

$$c(x_{<i}) = L_s(x_{<i}) - L_w(x_{<i}).$$

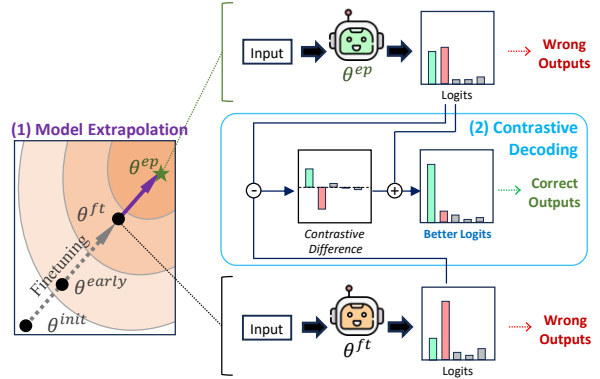


Figure 1: The process of EpiCoDe, which first employs model extrapolation to obtain a stronger LLM θ^{ep} , and then uses contrastive decoding to get better logit scores.

Following previous work (O’Brien and Lewis, 2023), we further define the logit score for next token as:

$$L_{CD}(x_{<i}) = L_s(x_{<i}) + \lambda c(x_{<i}),$$

where $\lambda > 0$ is a hyper-parameter controlling the strength of contrastive decoding.

3.2 Our Method: EpiCoDe

Formally, we denote the initial LLM as θ^{init} and finetune it on the training dataset. We can collect the model from two training stages: (1) the model θ^{early} at an early stage, which may be finetuned with part of the training data or trained with fewer epochs; (2) the model θ^{ft} finetuned on all data with more epochs.

However, in data scarcity cases, the model θ^{ft} may be still undertrained, due to limited training data available. To address this challenge, we develop a novel method EpiCoDe, which integrates model extrapolation and contrastive decoding, without the need of any additional training. Figure 1 illustrates the process of EpiCoDe.

We first employ model extrapolation, where θ^{early} serves as the weak model and θ^{ft} as the strong model, to obtain the extrapolated model $\theta^{ep} = \theta^{ft} + \mu(\theta^{ft} - \theta^{early})$. By selecting an appropriate μ , we can derive an extrapolated model θ^{ep} with better capabilities compared to θ^{ft} .

During inference, we use contrastive decoding for better performance. As shown in the left part of Figure 1, we have several models in hand, whose performance should have following ranking: $\theta^{init} < \theta^{early} < \theta^{ft} < \theta^{ep}$. We intuitively choose θ^{ep} as the strong model in contrastive decoding, while one of the other three models can serve as

the weak model. An important research question is which model should be employed to achieve the best performance. We argue that the similar error patterns between the weak and the strong model play an important role in the effectiveness of contrastive decoding. Thus, we employ θ^{ft} , which locates the nearest to θ^{ep} in the parameter space, as the weak model. The *contrastive difference* can be refined as $c(x_{<i}) = L_{ep}(x_{<i}) - L_{ft}(x_{<i})$, with $L_{ep}(x_{<i})$ and $L_{ft}(x_{<i})$ denoting the logit scores yield by θ^{ep} and θ^{ft} , respectively. Similarly, we represent the logit scores in contrastive decoding as $L_{CD} = L_{ep}(x_{<i}) + \lambda c(x_{<i})$.

Our pilot experiments show that using θ^{ft} as the weak model in contrastive decoding can result in good performance (See details in Section 7.2). We keep this setting for main experiments.

4 Why can EpiCoDe Work?

Previous analysis (Chang et al., 2024) explains that contrastive decoding may play a role similar to reduce the loss. However, a theoretical framework for understanding the underlying mechanisms still remains absent.

Different from previous work (Li et al., 2023; O’Brien and Lewis, 2023), in the contrastive decoding of EpiCoDe, we use models from different stages based on the same LLM, rather than models of different sizes. Our models vary in capabilities due to the quantity of finetuning data or whether they are obtained through model extrapolation.

In the context of EpiCoDe, we are uncertain whether using the extrapolated model θ^{ep} and its inferior version for contrastive decoding can be effective. With θ^{ep} serving as the strong model, we also aim to study which other model should be employed as the weak model. Thus, we develop a theoretical toolkit to explain the role of contrastive decoding in EpiCoDe and how to select the weak model yielding optimal performance.

4.1 Error Analysis of Contrastive Decoding

In this paper, we propose the theory from the perspective of logit scores. Let θ^* denote the hypothetical optimal model which may learn from a very large amount data, while θ denotes a model trained on insufficient data (e.g., our models θ^{init} , θ^{early} , θ^{ft} , and θ^{ep}). Given the input sequence $x_{<i}$ of any test example, we formulate the error in logit score predicted by θ as:

$$\delta(x_{<i}, \theta) = L(x_{<i}|\theta) - L(x_{<i}|\theta^*).$$

In contrastive decoding, we employ a strong θ^s and a weak model θ^w to yield logit scores, respectively. We suppose that $\delta(*, \theta^s)$ follows a normal distribution $\mathcal{N}(0, \epsilon^2)$. Similarly, the error of weak model’s logit scores has $\delta(*, \theta^w) \sim \mathcal{N}(0, (k\epsilon)^2)$, with $k > 1$.

We can then re-formulate logit score of contrastive decoding in Section 3.1 as:

$$L_{CD}(x_{<i}) = L(x_{<i}|\theta^*) + (1 + \lambda)\delta(x_{<i}, \theta^s) - \lambda\delta(x_{<i}, \theta^w).$$

The variance of the new logit score under contrastive decoding depends on the correlation between $\delta(*, \theta^s)$ and $\delta(*, \theta^w)$. We primarily study two typical scenarios: (1) The strong and the weak have learned similar capabilities; (2) The two model are not trained on the same data.

4.2 Locality in EpiCoDe

The variance of logit error can approach its lower bound $(1 - \lambda(k - 1))\epsilon$, when $\delta(*, \theta^s)$ and $\delta(*, \theta^w)$ exhibit perfect positive correlation¹This is an ideal scenario when the strong and the weak model have consistent patterns to make mistakes.

We argue that θ^{ep} and θ^{ft} naturally have **locality**, ensuring their similar error patterns. Given the evaluating loss on downstream tasks \mathcal{L} , the improvement of model extrapolation is $\Delta\mathcal{L} = \mathcal{L}(\theta^{ep}) - \mathcal{L}(\theta^{ft})$. To ensure model extrapolation can bring positive effects, we have to limit $\mu \ll 1$ and estimate the improvement by $\Delta\mathcal{L} = \mu(\theta^{ep} - \theta^{ft}) \cdot \nabla\mathcal{L}(\theta^{ft})$. Since LLMs suffer from insufficient training, this inner product can usually keep less than zero.

Note that θ^{ep} lies within the neighborhood of θ^{ft} . The **locality** in model extrapolation ensures them with similar capabilities and patterns to make mistakes. Therefore, we believe that employing θ^{ft} as the weak model can align closely with this idealized scenario. Contrastive decoding can help to decrease the variance of errors in the predicted logit scores (by up to $\lambda(k - 1)\epsilon$).

4.3 Negative Effect of Inconsistent Errors

We further introduce another idealized scenario where the strong and the weak model learn from different data. The two models make mistakes inconsistently, and we can suppose that $\delta(*, \theta^s)$

¹Denote the probability density of $\delta(*, \theta^s)$ and $\delta(*, \theta^w)$ as $f_s(u)$ and $f_w(u)$, respectively, where $u \in \mathbb{R}^V$. Here, we should have $f_w(u) = kf_s(u)$ for any u .

and $\delta(*, \theta^w)$ are independent of each other. In this case, the error of the logit scores in contrastive decoding follows a normal distribution $\mathcal{N}(0, (1 + \lambda)^2 \epsilon^2 + \lambda^2 k^2 \epsilon^2)$. This means the logit scores given by contrastive decoding even deviate more from the hypothetical optimal model θ^* . The initial model θ^{init} may hardly learn the same capabilities of θ^{ep} . Thus, we argue that using θ^{init} as the weak model may result in worse performance.

4.4 Weak Model Selection

Building on the two scenarios discussed above, we now analyze the effect of using θ^{early} as the weak model, which hinges on two key factors: (1) the performance gap between θ^{early} and θ^{ep} , which affects k ; and (2) the correlation between the errors in the logit scores predicted by θ^{early} and θ^{ep} , determining whether contrastive decoding can successfully bring improvement.

Previous works (O’Brien and Lewis, 2023; Chang et al., 2024) propose to use smaller LLMs as the weak model. Our theory supports such findings, since a larger k helps to degrade the lower bound of error variance $(1 - \lambda(k - 1))\epsilon$. Note that the LLMs with different sizes are trained on sufficient data, resulting in their similar capabilities and error patterns. We argue that, however, in data-scarcity scenarios, several updating steps can significantly separate θ^{ft} from θ^{early} in the parameter space, which violates the **locality** between θ^{early} and θ^{ep} .

We argue that using θ^{early} as the weak model will make the logit error higher than lower bound, and may even result in negative effects like the second idealized scenario. Thus, in EpiCoDe, we consistently choose θ^{ft} as the weak model in our main experiments. We provide more experimental results in Section 7.2 to examine our theory.

5 Experimental Settings

Datasets We evaluate our proposed method in three different tasks that usually lack finetuning data: law, mathematics, and logical reasoning. In the legal domain, we employ JEC-QA (Zhong et al., 2019) as the training set and DISC-Law-Eval (Yue et al., 2023) as the test set. Following previous work (Huang et al., 2023), we convert these multiple-choice legal questions into true-or-false questions. For mathematics and logical reasoning, we use GSM8K_zh (Yu et al., 2024) and LogiQA (Liu et al., 2020) for finetuning and evaluation, respectively.

Task	Number of Examples			Avg. Len.
	Train	Dev	Test	
Law	17,664	5,126	5,126	265.3
Math	7,471	660	659	117.2
Logic	12,672	1,000	1,000	480.4

Table 1: Statistics of the finetuning and evaluation datasets for each task. **Avg. Len.** shows average output length of training examples, counted by character.

For the legal question answering (QA) and logical reasoning tasks, which require multiple reasoning steps to handle, we collect chain-of-thought (CoT, Wei et al., 2022) data to finetune LLMs. For the mathematical task, we utilize the intermediate solution steps from GSM8K_zh for training, despite their brevity. Table 1 summarizes overall statistics of the datasets. We search the optimal hyper-parameters on hold-out validation sets before testing. See more details in Appendix A.

Metrics Given that our models have undergone finetuning, we evaluate their performance in a zero-shot setting. We adopt accuracy for legal QA and logical reasoning tasks, while exact match for the mathematical task.

Models We explore the effectiveness of EpiCoDe on the LLMs from three families with different scales, including Deepseek-7B-Chat (Bi et al., 2024), Qwen2-1.5B-Instruct, Qwen2-7B-Instruct (Yang et al., 2024), and Llama-3.2-3B-Instruct (Team, 2024). We continually finetune the instruct models with our data for two epochs. We regard the checkpoint after one epoch as θ^{early} , while the checkpoint after two epochs as θ^{ft} . Here we use AdamW (Loshchilov and Hutter, 2019) as the optimizer, with β_1 and β_2 are set to 0.9 and 0.95 respectively. The maximum learning rate is set to $3e-5$, and the batch size is set to 128.

Decoding Constraints Contrastive decoding has potential flaws in computing the logit scores of incorrect tokens. Specifically, if the weak model assigns an extremely low logit score to a wrong token $e \in \mathcal{V}$, the *contrastive difference* of this token $c(x_i = e | x_{<i})$ becomes excessively large. The model enhanced with contrastive decoding can be more prone to generating this incorrect token.

To address this issue, following previous work (Li et al., 2023), we introduce a threshold α , which restricts the model must select the next token from those with high logit scores predicted

Method	Accuracy (Δ Acc)			Avg.
	Law	Math	Logic	
<i>Deepseek-7B-Chat</i>				
Finetune	64.78	27.28	57.22	49.76
ME	65.42(+0.64)	27.12(-0.17)	58.89(+1.67)	50.48
CD	65.29(+0.51)	26.88(-0.41)	58.46(+1.24)	50.21
EpiCoDe	65.51(+0.73)	27.81(+0.53)	59.05(+1.83)	50.79
<i>Qwen2-1.5B-Instruct</i>				
Finetune	64.34	39.48	52.68	52.17
ME	64.94(+0.60)	40.80(+1.32)	53.63(+0.95)	53.12
CD	65.22(+0.88)	39.23(-0.26)	53.42(+0.74)	52.62
EpiCoDe	65.38(+1.04)	41.20(+1.71)	53.87(+1.19)	53.48
<i>Qwen2-7B-Instruct</i>				
Finetune	69.03	57.12	66.67	64.27
ME	69.46(+0.43)	57.79(+0.67)	67.60(+0.93)	64.95
CD	69.91(+0.88)	58.62(+1.50)	67.43(+0.76)	65.32
EpiCoDe	70.25(+1.22)	58.71(+1.59)	68.07(+1.40)	65.68
<i>Llama-3.2-3B-Instruct</i>				
Finetune	62.13	48.45	53.45	54.68
ME	62.73(+0.60)	49.74(+1.29)	55.11(+1.66)	55.77
CD	63.38(+1.25)	53.13(+4.68)	56.62(+3.17)	57.71
EpiCoDe	63.79(+1.66)	54.31(+5.86)	57.48(+4.03)	58.53

Table 2: Accuracy of distinct LLMs at each stage on the three tasks. We abbreviate contrastive decoding as CD and model extrapolation as EP.

by the strong model.

For all experiments, we empirically set α to 0.1.

6 Main Results

We adopt the models finetuned for two epochs as baseline. As comparative experiments, we also study the improvements when using only model extrapolation (Zheng et al., 2024) or contrastive decoding (O’Brien and Lewis, 2023). For model extrapolation, we directly evaluate the performance of θ^{ep} . And for contrastive decoding, we employ θ^{early} and θ^{ft} as the weak and the strong model, respectively. We repeat to train the LLMs on each dataset for 10 times with different random seeds and report the average accuracy.

As shown in Table 2, model extrapolation and contrastive decoding are both effective methods to enhance LLMs without training. In general, using only model extrapolation results in an average improvement of 0.68%~1.09% across the three tasks, while contrastive decoding yields improvement of 0.45%~3.03%. However, **our proposed EpiCoDe outperforms these two baselines**, enhancing the finetuned models with the most improvement of 1.03%~3.85%.

We find that the effects of model extrapolation and contrastive decoding vary on different LLMs. For instance, contrastive decoding can bring signif-

\mathcal{H}_0	Model	Significance Level α		
		Law	Math	Logic
EpiCoDe>ME	<i>DS-7B</i>	0.0899	0.0198	0.0821
	<i>QW-1.5B</i>	0.0026	0.0938	0.0711
	<i>QW-7B</i>	0.0000	0.0067	0.0783
	<i>LM-3B</i>	0.0000	0.0000	0.0001
EpiCoDe>CD	<i>DS-7B</i>	0.0190	0.0341	0.0267
	<i>QW-1.5B</i>	0.0342	0.0001	0.0116
	<i>QW-7B</i>	0.0279	0.3344	0.0130
	<i>LM-3B</i>	0.0113	0.0007	0.0006

Table 3: Significance level of mistakenly rejecting the null hypothesis \mathcal{H}_0 in one-tailed tests. The **dark red** scores mean that the confidence to believe \mathcal{H}_0 being true is lower than **95%**. DS-7B, QW-1.5B, QW-7B, and LM-3B represent to employ Deepseek-7B-Chat, Qwen2-1.5B-Instruct, Qwen2-7B-Instruct, and Llama-3.2-3B-Instruct as the initial LLM, respectively.

icant improvement (1.25%~4.68%) to Llama-3.2-3B-Instruct, which is trained on fewer Chinese data, while model extrapolation enhance the accuracy by merely 0.60%~1.66%. However, for the other LLMs which have learned from sufficient Chinese texts, model extrapolation is usually more effective than contrastive decoding. Differently, **EpiCoDe can always leverage the strengths of both methods**, consistently yielding superior improvements across all LLMs.

Significance of EpiCoDe’s Improvement We further study whether EpiCoDe can bring significantly more improvement than the baseline methods. For each task and each random seed, we evaluate the performance of three settings: (1)using only model extrapolation, (2) using only contrastive decoding, and (3) our proposed EpiCoDe method. Then, we compare the performance of the EpiCoDe method with the two baseline methods under the same random seed, conducting a paired t-test. The null hypothesis is that *EpiCoDe performs better than each of the two baseline methods*.

Table 3 illustrates the magnitudes of the significance level α , which means the probability of mistakenly rejecting the null hypothesis. Except for Qwen2-7B-Instruct on Math, we can conclude with more than 90% confidence level to believe that EpiCoDe outperforms the two baseline methods, using only model extrapolation or only contrastive decoding. Furthermore, across the 4 LLMs on 3 tasks, we have over 95% confidence to believe that EpiCoDe surpasses model extrapolation in 7 (/12) experiments and outperforms contrastive decoding in 11 (/12) experiments. Thus, we believe **EpiCoDe is**

Method	Num. of Successful Runs			Total (* / 30)
	Law	Math	Logic	
<i>Deepseek-7B-Chat</i>				
ME	8	6	10	24
CD	8	4	8	20
EpiCoDe	10	8	10	28
<i>Qwen2-1.5B-Instruct</i>				
ME	8	9	7	24
CD	9	4	7	20
EpiCoDe	10	10	9	29
<i>Qwen2-7B-Instruct</i>				
ME	8	8	8	24
CD	10	8	7	25
EpiCoDe	10	10	10	30
<i>Llama-3.2-3B-Instruct</i>				
ME	10	10	10	30
CD	10	10	10	30
EpiCoDe	10	10	10	30

Table 4: The number of times that each method brings improvement to the finetuned models across 10 runs.

significantly better than the two baselines, which achieves optimal performance.

Robustness of EpiCoDe Table 2 also shows that EpiCoDe can consistently enhance all the four LLMs across three tasks. But in the mathematical task, **using only contrastive decoding or only model extrapolation sometimes fails** to bring improvement (e.g., for Deepseek-7B-Chat and Qwen2-1.5B-Instruct). Thus, we further examine the number of times that each method successfully enhances the finetuned model across 10 repeated experiments. As shown in Table 4, **our EpiCoDe can provide very robust improvement for the finetuned LLMs**, with merely three failures in a total of 120 experiments. As a comparison, model extrapolation yields improvement in 80% of the experiments on Deepseek-7B-Chat and the Qwen2 family. Even worse, contrastive decoding exhibits a higher frequency to fail in Deepseek-7B-Chat and Qwen2-1.5B-Instruct, especially on the mathematical task, where it has positive effects to each model merely 4 times. But for Qwen2-7B-Instruct, the two baselines have comparable rates to achieve better accuracy than finetuned models (24 vs. 25 out of 30 runs). Note that as a very strong LLM, Qwen2-7B-Instruct may have learned from various datasets, including those similar to our tasks. We guess that its parameters change less during the training process, which ensures **locality** between θ^{early} and θ^{ft} , an essential premise for contrastive learning to work. See more discussion about local-

ity in Section 7.2.

For 10 runs across all three tasks, EpiCoDe and the two baselines can consistently improve Llama-3.2-3B-Instruct, which suffers extremely from the lack of Chinese training data. Since model extrapolation and contrastive decoding are both effective, we believe that **the combination of two techniques plays a vital role to boost model performance in scenarios with restricted training resources**.

7 Discussions

7.1 Where does the improvement come from?

Table 2 shows that the effectiveness of each method varies across tasks. For instance, in logical reasoning, all three methods can improve the performance of each finetuned LLM, while in legal QA, the improvement becomes relatively less. However, in the mathematical task, only EpiCoDE yields improvement for all LLMs, while using only model extrapolation or only contrastive decoding even degrades performance for models such as Qwen2-1.5B-Instruct and Llama-3.2-3B-Instruct.

We guess that **the difficulty of a task may influence the effectiveness of these methods**. Specifically, GSM8K_zh is a simple task, primarily consisting of grade school math problems. In contrast, legal QA and logical reasoning, as more difficult tasks, require LLMs to possess enough domain-specific knowledge and commonsense, to distinguish similar concepts, and to apply multi-step reasoning. The average length of training examples in Table 1 can represent the difficulties of each task (longer outputs for the more complex task).

To dive deeper into where EpiCoDe brings the improvement, we take legal QA as an example. For each run, we collect the token sequences outputted by the finetuned model (θ^{ft}), and count their lengths. We partition the test set of legal QA into three equal parts according to the length order. We then evaluate the improvement of the three methods on each subset.

Table 5 shows the performance of vanilla finetuned LLMs and the models enhanced by each method. By comparing the results of finetuned models and extrapolated models, we find that **model extrapolation can significantly benefits the hard subsets** (needing the longest CoT), resulting in improvement of +1.02%~+1.67% on Qwen family and Deepseek-7B-Chat. But for the easy part, model extrapolation improves the perfor-

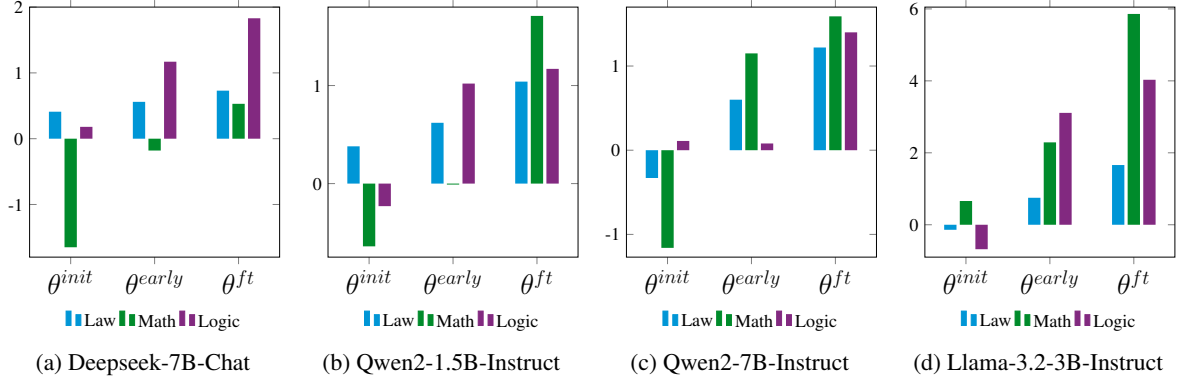


Figure 2: The improvement of using θ^{init} , θ^{early} , or θ^{ft} as the weak model, compared to the performance of vanilla finetuned model θ^{init} .

Method	Easy		Medium		Hard	
	Acc.	Δ Acc.	Acc.	Δ Acc.	Acc.	Δ Acc.
<i>Deepseek-7B-Chat</i>						
Finetune	68.81	–	65.08	–	60.65	–
ME	68.86	+0.05	65.26	+0.18	62.31	+1.67
EpiCoDe	68.93	+0.13	65.28	+0.19	62.49	+1.84
CD only	68.62	-0.18	65.16	+0.08	62.25	+1.60
<i>Qwen2-1.5B-Instruct</i>						
Finetune	69.05	–	64.37	–	59.68	–
ME	69.31	+0.26	64.89	+0.52	60.70	+1.02
EpiCoDe	69.38	+0.33	65.88	+1.52	60.96	+1.28
CD only	69.25	+0.20	65.40	+1.03	61.07	+1.39
<i>Qwen2-7B-Instruct</i>						
Finetune	75.77	–	69.19	–	62.36	–
ME	75.68	-0.10	69.52	+0.33	63.40	+1.04
EpiCoDe	75.70	-0.07	70.48	+1.28	64.80	+2.44
CD only	75.65	-0.13	69.95	+0.76	64.34	+1.98
<i>Llama-3.2-3B-Instruct</i>						
Finetune	64.74	–	62.36	–	59.37	–
ME	65.15	+0.41	63.07	+0.71	60.04	+0.67
EpiCoDe	65.27	+0.52	63.84	+1.48	62.33	+2.96
CD only	65.23	+0.48	63.58	+1.21	61.39	+2.03

Table 5: Each method’s accuracy on the three subsets of legal QA. For each experiment, we compare the improvement brought by corresponding method on the three subsets. The highest improvement among three subsets is made **bold**, while the second highest is made *italic*.

mance much less, sometimes even degrading the accuracy. However, for Llama-3.2-3B-Instruct which has limited Chinese capabilities, model extrapolation brings comparable improvement to three subsets.

We then examine the performance of EpiCoDe. EpiCoDe surpasses model extrapolation by +0.17%~+1.40% on the hard subset, while merely +0.01~+1.00% on the medium subset. Thus, our EpiCoDe, which is further enhanced by contrastive decoding, also gains improvement primarily from

the hard part.

We also study the effect of using only contrastive decoding, where θ^{ft} and θ^{early} serve as the strong and the weak model, respectively. It also results in the most improvement (+1.39%~+2.03%) on the hard subsets, but relatively less improvement (+0.08%~+1.21%) on the medium ones and even nearly no improvement on the easy ones.

We believe that model extrapolation and contrastive decoding play similar roles in improving LLMs’ capabilities, **especially when generating longer CoT to tackle challenging problems**. We also argue that **EpiCoDe can integrate the strengths of both techniques to achieve optimal performance**.

7.2 Examination of Locality

After model extrapolation, we can obtain four models, θ^{init} , θ^{early} , θ^{ft} , and θ^{ep} . We intuitively employ θ^{ep} as the strong model in contrastive decoding, due to its strongest capabilities. We aim to study the effect of selecting each of the other models as the weak model.

Figure 2 illustrates the improvement of employing each model as the weak model during contrastive decoding. We find that when using the initial model θ^{init} as the weak model, the performance of EpiCoDe often lags behind the vanilla finetuned model θ^{ft} , especially in mathematical tasks. This demonstrates that **the weak model must learn from the same datasets to the extrapolated model; otherwise, contrastive decoding will have a negative effect**, even neutralizing the improvements brought by model extrapolation.

Previous works (O’Brien and Lewis, 2023; Chang et al., 2024) argue that a larger capability gap between the weak and the strong model can

bring more improvement, which indicates θ^{early} may be a better choice than θ^{ft} . However, Figure 2 shows that, if we employ θ^{early} as the weak model in contrastive decoding, EpiCoDe often achieves comparable accuracy to the extrapolated model θ^{ep} , for instance, Deepseek-7B-Chat and Qwen2-1.5B on legal QA and logical reasoning. In the mathematical task, using θ^{early} also makes the improvement of EpiCoDe down to around zero. As shown in Section 4, we empirically demonstrate that **in data-scarcity scenarios, several update steps may break the locality of models, resulting in inconsistent error patterns.**

Comparing the results in Figure 2, we find that employing θ^{ft} as the weak model can consistently achieve the optimal performance for each LLM across all tasks. We are confident that when selecting models from different training stages for contrastive decoding, **similar error patterns between the strong and the weak model is more crucial than a large capability gap.** Therefore, different from previous works (O’Brien and Lewis, 2023; Chang et al., 2024), the best strategy here is to choose models from adjacent stages.

8 Conclusion

In this paper, we propose an easy yet effective method, EpiCoDe, to enhance finetuned LLMs in data scarcity scenarios. Benefiting from both model extrapolation and contrastive decoding, EpiCoDe can operate without additional training. Extensive experiments across different tasks show that EpiCoDe can bring consistent improvement, significantly outperforming baseline models. We also propose a theoretical framework from the perspective of logit errors to analyze the mechanisms of how EpiCoDe can work better. We find that model extrapolation naturally ensures the locality between the finetuned LLM and the extrapolated LLM. During inference, their locality further fulfills the prerequisite that enables contrastive decoding to be effective.

Limitations

Investigated Tasks Due to the limitation of copyrights, there are still no publicly accessible legal QA or logical testing datasets in English. Thus, we employ the legal QA datasets adapted from the judicial examination of China and the logical reasoning dataset from National Public Service Examination in China.

Such tasks are challenging enough, since they require LLMs to possess domain-specific knowledge and commonsense and to obtain final answers through complex multi-step reasoning. Existing LLMs still can not perform well on such tasks (Yue et al., 2023; Liu et al., 2023), which aligns the need of developing training-free methods to enhance model capabilities.

Investigated Models Due to the restrictions of computational resources, we mainly study the LLMs with scales from 1.5B to 7B. The results of LLMs from difference families across all three tasks show that our proposed EpiCoDe can consistently outperform the baselines. We believe that EpiCoDe can be effective for larger LLMs, which suffer more from data scarcity.

Acknowledgments

This work is supported in part by NSFC (62161160339) and Beijing Science and Technology Program (Z231100007423011). We thank Di He, Xiao Liu, Huishuai Zhang, and the anonymous reviewers for their helpful discussions and suggestions. For any correspondence, please contact Yansong Feng.

References

- Xiao Bi, Deli Chen, Guanting Chen, Shanhuang Chen, Damai Dai, Chengqi Deng, Honghui Ding, Kai Dong, Qiushi Du, Zhe Fu, Huazuo Gao, Kaige Gao, Wenjun Gao, Ruiqi Ge, Kang Guan, Daya Guo, Jianzhong Guo, Guangbo Hao, Zhewen Hao, Ying He, Wenjie Hu, Panpan Huang, Erhang Li, Guowei Li, Jiashi Li, Yao Li, Y. K. Li, Wenfeng Liang, Fangyun Lin, A. X. Liu, Bo Liu, Wen Liu, Xiaodong Liu, Xin Liu, Yiyuan Liu, Haoyu Lu, Shanghao Lu, Fuli Luo, Shiron Ma, Xiaotao Nie, Tian Pei, Yishi Piao, Junjie Qiu, Hui Qu, Tongzheng Ren, Zehui Ren, Chong Ruan, Zhangli Sha, Zhihong Shao, Junxiao Song, Xuecheng Su, Jingxiang Sun, Yaofeng Sun, Minghui Tang, Bingxuan Wang, Peiyi Wang, Shiyu Wang, Yaohui Wang, Yongji Wang, Tong Wu, Y. Wu, Xin Xie, Zhenda Xie, Ziwei Xie, Yiliang Xiong, Hanwei Xu, R. X. Xu, Yanhong Xu, Dejian Yang, Yuxiang You, Shuiping Yu, Xingkai Yu, B. Zhang, Haowei Zhang, Lecong Zhang, Liyue Zhang, Mingchuan Zhang, Minghua Zhang, Wentao Zhang, Yichao Zhang, Chenggang Zhao, Yao Zhao, Shangyan Zhou, Shunfeng Zhou, Qihao Zhu, and Yuheng Zou. 2024. [Deepseek llm: Scaling open-source language models with longtermism](#). Preprint, arXiv:2401.02954.
- Haw-Shiuan Chang, Nanyun Peng, Mohit Bansal, Anil Ramakrishna, and Tagyoung Chung. 2024. [Explaining and improving contrastive decoding by extrapolation](#).

- olating the probabilities of a huge and hypothetical LM. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 8503–8526, Miami, Florida, USA. Association for Computational Linguistics.
- Junying Chen, Zhenyang Cai, Ke Ji, Xidong Wang, Wanlong Liu, Rongsheng Wang, Jianye Hou, and Benyou Wang. 2024. [Huatuogpt-o1, towards medical complex reasoning with llms](#). *Preprint*, arXiv:2412.18925.
- Leshem Choshen, Elad Venezian, Noam Slonim, and Yoav Katz. 2022. [Fusing finetuned models for better pretraining](#). *Preprint*, arXiv:2204.03044.
- DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui Ding, Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang, Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, R. L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, S. S. Li, Shuang Zhou, Shaoqing Wu, Shengfeng Ye, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wanjia Zhao, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong Liu, Xiaohan Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue Jin, Xiaojin Shen, Xiaosha Chen, Xiaowen Sun, Xiaoxiang Wang, Xinnan Song, Xinyi Zhou, Xianzu Wang, Xinxia Shan, Y. K. Li, Y. Q. Wang, Y. X. Wei, Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Wang, Yi Yu, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yuduan Wang, Yue Gong, Yuheng Zou, Yujia He, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanhong Xu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Ying Tang, Yukun Zha, Yuting Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen Huang, Zhipeng Xu, Zhongyu Zhang, and Zhen Zhang. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#). *Preprint*, arXiv:2501.12948.
- Quzhe Huang, Mingxu Tao, Chen Zhang, Zhenwei An, Cong Jiang, Zhibin Chen, Zirui Wu, and Yansong Feng. 2023. [Lawyer llama technical report](#). *Preprint*, arXiv:2305.15062.
- Tero Karras, Miika Aittala, Tuomas Kynkäänniemi, Jaakko Lehtinen, Timo Aila, and Samuli Laine. 2024. [Guiding a diffusion model with a bad version of itself](#). In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Xiang Lisa Li, Ari Holtzman, Daniel Fried, Percy Liang, Jason Eisner, Tatsunori Hashimoto, Luke Zettlemoyer, and Mike Lewis. 2023. [Contrastive decoding: Open-ended text generation as optimization](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12286–12312, Toronto, Canada. Association for Computational Linguistics.
- Hanmeng Liu, Jian Liu, Leyang Cui, Zhiyang Teng, Nan Duan, Ming Zhou, and Yue Zhang. 2023. [Logiqa 2.0—an improved dataset for logical reasoning in natural language understanding](#). *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 31:2947–2962.
- Jian Liu, Leyang Cui, Hanmeng Liu, Dandan Huang, Yile Wang, and Yue Zhang. 2020. [Logiqa: A challenge dataset for machine reading comprehension with logical reasoning](#). In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 3622–3628. International Joint Conferences on Artificial Intelligence Organization. Main track.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *International Conference on Learning Representations*.
- Sean O’Brien and Mike Lewis. 2023. [Contrastive decoding improves reasoning in large language models](#). *Preprint*, arXiv:2309.09117.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Gray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). In *Advances in Neural Information Processing Systems*.
- Llama Team. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed H. Chi, Quoc V Le, and Denny Zhou. 2022. [Chain of thought](#)

[prompting elicits reasoning in large language models](#). In *Advances in Neural Information Processing Systems*.

Mitchell Wortsman, Gabriel Ilharco, Samir Ya Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, and Ludwig Schmidt. 2022. [Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time](#). In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 23965–23998. PMLR.

An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jianxin Yang, Jin Xu, Jingren Zhou, Jinze Bai, Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang, Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao, Runji Lin, Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Xuejing Liu, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, Zhifang Guo, and Zhihao Fan. 2024. [Qwen2 technical report](#). Preprint, arXiv:2407.10671.

Longhui Yu, Weisen Jiang, Han Shi, Jincheng YU, Zhengying Liu, Yu Zhang, James Kwok, Zhenguo Li, Adrian Weller, and Weiyang Liu. 2024. [Metamath: Bootstrap your own mathematical questions for large language models](#). In *The Twelfth International Conference on Learning Representations*.

Shengbin Yue, Wei Chen, Siyuan Wang, Bingxuan Li, Chenchen Shen, Shujun Liu, Yuxuan Zhou, Yao Xiao, Song Yun, Xuanjing Huang, and Zhongyu Wei. 2023. [Disc-lawllm: Fine-tuning large language models for intelligent legal services](#). Preprint, arXiv:2309.11325.

Chujie Zheng, Ziqi Wang, Heng Ji, Minlie Huang, and Nanyun Peng. 2024. [Weak-to-strong extrapolation expedites alignment](#). In *ICML 2024 Workshop on Models of Human Feedback for AI Alignment*.

Haoxi Zhong, Chaojun Xiao, Cunchao Tu, Tianyang Zhang, Zhiyuan Liu, and Maosong Sun. 2019. [Jecqa: A legal-domain question answering dataset](#). Preprint, arXiv:1911.12011.

A Hyper-Parameter Selection

For model extrapolation, we empirically choose the hyper-parameter μ from $1e-4$ to 0.8 by logarithmic interval, within $\{1, 2, 4, 6, 8\} \times \{10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}\}$. And for contrastive

Magnitude of μ	DS-7B	QW-1.5B	QW-7B	LM-3B
0 (vanilla finetuned)	57.22	52.68	66.67	53.45
$1e-4 \sim 8e-4$	<u>57.58</u>	52.87	67.19	54.00
$1e-3 \sim 8e-3$	58.39	53.43	<u>66.90</u>	55.11
$1e-2 \sim 8e-2$	57.29	<u>52.99</u>	66.82	<u>55.10</u>
$1e-1 \sim 8e-1$	57.06	51.38	66.53	54.20

Table 6: Performance of extrapolated LLMs on Logic under different μ . The highest results are made **bold**, with the second underlined. DS-7B, QW-1.5B, QW-7B, and LM-3B share the same meanings with those in Table 3.

Magnitude of λ	DS-7B	QW-1.5B	QW-7B	LM-3B
0 (only using ME)	58.89	53.63	67.60	55.11
0.1	58.86	53.64	67.59	56.27
0.2	58.97	53.57	68.06	56.68
0.4	58.77	53.79	68.03	56.80
0.6	58.46	<u>53.74</u>	67.97	<u>57.16</u>
0.8	58.55	53.48	67.81	57.33
1.0	58.34	52.98	67.49	57.06

Table 7: Performance of contrastive decoding on Logic under different λ . The highest results are made **bold**, with the second underlined. DS-7B, QW-1.5B, QW-7B, and LM-3B share the same meanings with those in Table 3.

decoding, we choose the hyper-parameter λ within $\{0.1, 0.2, 0.4, 0.6, 0.8, 1.0\}$.

For all experiments, we search the optimal hyper-parameters on development sets, and then employ the same hyper-parameters to evaluate models on hold-out test sets.

Since EpiCoDe can be regarded as a two-stage method, we do NOT use grid search to concurrently optimize μ and λ . Instead, we first select the optimal μ in model extrapolation. We then search λ with μ frozen, ensuring all experiments are fair between EpiCoDe and only using model extrapolation.

We find that these hyper-parameters have similar effects across distinct LLMs and runs. As μ or λ increases from small to large, the improvement brought by model extrapolation or contrastive decoding exhibits convexity. For example, Table 6 shows the effect of μ with different magnitudes in model extrapolation on the Logic task (average over 10 runs). We also use Figure 7 to illustrate the effect of different λ in contrastive decoding on Logical Reasoning (average over 10 runs). Similar to previous works (O’Brien and Lewis, 2023; Zheng et al., 2024), when μ or λ is small, both model extrapolation and contrastive decoding can bring more improvement as the μ or λ grows larger.

Weak Model	Law	Math	Logic
θ^{ft} of Qwen2-1.5B	69.60 _(+0.57)	49.39 _(-7.73)	67.91 _(+1.24)
θ^{init} of Qwen2-7B	68.70 _(-0.33)	55.96 _(-1.16)	66.78 _(+0.11)
θ^{early} of Qwen2-7B	69.63 _(+0.60)	58.27 _(+1.15)	66.75 _(+0.08)
θ^{ft} of Qwen2-7B	70.25 _(+1.22)	58.71 _(+1.59)	68.07 _(+1.40)

Table 8: The results of different weak models in EpiCoDe. For all experiments, we employ θ^{ep} of Qwen2-7B-Instruct as the strong model.

But if we use very large μ or λ , the two methods will fail to improve model performance.

B LLMs with Different Sizes in CD

Inspired by previous work (Li et al., 2023; O’Brien and Lewis, 2023) which employs a small and a large model in contrastive decoding, we also wonder during inference, whether EpiCoDe is still effective if using two models with different numbers of parameters. Thus, we employ the extrapolated Qwen2-7B-Instruct as the strong model and the finetuned Qwen2-1.5B-Instruct as the weak model.

Table 8 shows that contrastive decoding can sometimes be effective when we employ a smaller LLM, which is finetuned on the same data, as the weak model. However, such setting brings less improvement than EpiCoDe where we employ the same-size finetuned LLM as weak model.

Using LLMs of different sizes is also not robust enough to enhance model performance. For example, in two runs on Math, if we employ θ^{ep} of Qwen2-1.5B as the weak model, the accuracy scores decrease from around 57% to around 16%, resulting in catastrophic deterioration.

From these additional results, we find that the effectiveness of contrastive decoding may be highly dependent on the **locality**. Beyond existing methods which employ a large LLM as the strong model and a small LLM as the weak model, our proposed EpiCoDe can achieve the optimal performance.