

Learning or Self-aligning? Rethinking Instruction Fine-tuning

Mengjie Ren^{1,3*}, Boxi Cao^{1,3}, Hongyu Lin^{1†}, Cao Liu⁵, Xianpei Han^{1,2,4}

Ke Zeng⁵, Guanglu Wan⁵, Xunliang Cai⁵, Le Sun^{1,2,4}

¹Chinese Information Processing Laboratory ²State Key Laboratory of Computer Science
Institute of Software, Chinese Academy of Sciences

³University of Chinese Academy of Sciences

⁴Key Laboratory of System Software, Chinese Academy of Sciences

⁵Meituan

{renmengjie2021,boxi2020,hongyu,xianpei,sunle}@iscas.ac.cn

{liucaozengke02,wanguanglu,caixunliang}@meituan.com

Abstract

Instruction Fine-tuning (IFT) is a crucial phase in building large language models (LLMs). Previous works mainly focus on the IFT's role in the transfer of behavioral norms and the learning of additional world knowledge. However, the understanding of the underlying mechanisms of IFT remains significantly limited. In this paper, we design a knowledge intervention framework to decouple the potential underlying factors of IFT, thereby enabling individual analysis of different factors. Surprisingly, our experiments reveal that attempting to learn additional world knowledge through IFT often struggles to yield positive impacts and can even lead to markedly negative effects. Further, we discover that maintaining internal knowledge consistency before and after IFT is a critical factor for achieving successful IFT. Our findings reveal the underlying mechanisms of IFT and provide robust support for some very recent and potential future works. We release our experimental dataset and codes to facilitate future work¹.

1 Introduction

The advent and evolution of large language models (LLMs) have marked a significant milestone in natural language processing (NLP) (Brown et al., 2020; Touvron et al., 2023b,a). As one of the core steps in the construction of LLMs, instruction fine-tuning (IFT) employs supervised instruction-response pairs to fine-tune LLMs, thereby facilitating the transformation of LLM from a continuous writing model to a question-answering agent (Chung et al., 2024; Iyer et al., 2022; Jang et al., 2023).

* Work was partially done during Ren's internship at Meituan.

† Corresponding Author

¹<https://github.com/renmengjie7/Self-Aligning>

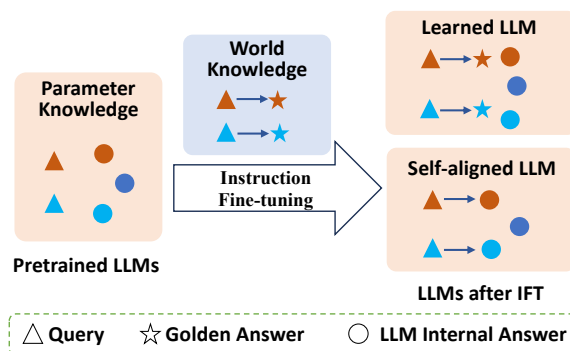


Figure 1: Two potential mechanisms for instruction fine-tuning. 1) Learning, which injects world knowledge in IFT data into LLMs; 2) Self-aligning, which aligns queries with knowledge already in LLMs with similar behavioral norms. Elements with the same color are related.

Despite the crucial role of IFT in the construction of LLMs, there is a significant lack of in-depth research on the mechanisms by which IFT operates. In traditional machine learning, supervised learning aims to fit models to specific tasks and data distributions (Goodfellow et al., 2016; Bishop, 2006), whereas the impact of IFT on LLMs is markedly different. As shown in Figure 1, on one hand, one of the most apparent effects of IFT is its ability to align the output of LLMs more closely with the latent behavioral norms contained within the IFT data, thereby enabling more effective parameter knowledge expression (Zhou et al., 2024; Chen et al., 2023a). On the other hand, many existing studies (Li et al., 2023; Cui et al., 2023; Chen et al., 2023b) aim to facilitate domain-specific adaptation of LLMs through IFT, by injecting the world knowledge contained in IFT data into LLMs.

Unfortunately, both the transfer of behavioral norms and the enhancement of domain knowledge

are closely coupled with the corpus applied in IFT, rendering the analysis of IFT’s true effects exceedingly challenging. Due to the interconnection of these two effects, it is challenging to discern whether the benefits derived are due to the promotion of more effective expression of parameter knowledge or the injection of additional world knowledge. The coupling between the above two factors, along with a lack of in-depth analysis of the IFT’s mechanism, hampers our comprehension of the effectiveness of IFT. This limitation hinders our ability to develop robust strategies for IFT data construction, model training, and model evaluation due to insufficient theoretical support. Therefore, a thorough and comprehensive analysis of the underlying core factors that drive the effectiveness of IFT is crucial for achieving more effective IFT.

To this end, this paper designs a knowledge intervention framework for analyzing the underlying mechanisms of IFT. The main idea of our framework is to control the consistency between the knowledge in IFT data and the existing parameter knowledge of LLM, in order to decouple the injection of domain knowledge from the transfer of behavioral norms during IFT. This allows for a separate analysis of the roles of these two crucial factors. Specifically, we first employ in-context learning (ICL) (Dong et al., 2022; Brown et al., 2020) to probe the internal parameter knowledge of LLMs. Building on this, we intervene in the composition of existing parameter knowledge and the newly introduced world knowledge within IFT data and then observe the differences in model behavior after IFT using different intervention groups. Based on the framework, we conduct an in-depth analysis to answer the following two critical research questions (RQ):

- **RQ1:** *How does the world knowledge within IFT data affect LLMs?*
- **RQ2:** *What is the underlying cause of the above impact?*

For RQ1, we initially discover that significant discrepancies between the world knowledge contained in IFT data and the existing parameter knowledge within LLM can substantially undermine the model’s abilities. Performance derived from a set of IFT data that contains incorrect world knowledge but aligns with the model’s parameter knowledge is significantly better than that from a set containing correct world knowledge but inconsistent

with the model’s internal parameter knowledge. To dive into this phenomenon, we explicitly supply LLMs with the world knowledge necessary for answering the instruction, integrated into the context. This strategy allows the model to focus on transforming the output behavioral norms instead of jointly learning the inconsistent world knowledge. We discover that the detrimental effects caused by the inconsistency between parameter knowledge and world knowledge can be significantly mitigated by explicitly providing such self-contained IFT data points. These two findings indicate that attempting to introduce world knowledge through IFT that is inconsistent with the model’s parameter knowledge can severely undermine the model, which suggests that the injection of world knowledge does not lie at the center of a successful IFT.

For RQ2, we further analyze the model performance under varying degrees of consistency between the knowledge contained in IFT data and the parameter knowledge in the original LLM. We find that while the consistency between the two has a significant impact on model performance, a higher degree of consistency does not necessarily correlate with better model performance. However, our further research reveals a strong correlation between the model’s ultimate performance after IFT and the consistency of the model’s internal knowledge before and after IFT. That is to say, for the model after IFT, if its responses are more consistent with the responses produced by the original model from in-context learning probing, then the performance of the model after IFT is also better. The validity of the finding is independent of whether the test data belongs to the same domain as the training data, and is also unrelated to the original performance of the base LLM. This implies that the phenomenon is solely influenced by the consistency of the knowledge before and after IFT. Furthermore, we discover that using IFT data that is either too consistent or too inconsistent with the original parameter knowledge can lead to a divergence in the model’s internal knowledge before and after IFT, thereby resulting in a decline in performance.

Our experiments reveal the fundamental role of IFT in the construction of LLMs. Essentially, IFT is not a supervised, domain-specific **learning** process, but a process of **self-aligning** instruction with existing internal knowledge of LLMs that can be obtained through few-shot in-context learning probing. Our findings not only provide robust theo-

retical support for the very recently emerged research on self-alignment (Sun et al., 2024) and super-alignment (Burns et al., 2023), etc., but also shed light on the future direction of data construction, model training, and model evaluation for IFT.

2 Related Work

By observing output token distribution shift of models before and after IFT, Lin et al. (2024) found that most shifts occur with stylistic tokens, strongly supporting the superficial alignment hypothesis (Zhou et al., 2024), false promise (Gudibande et al., 2023) and related works on IFT data construction (Chen et al., 2023a; Shen, 2024) and proxy-guided decoding (Liu et al., 2024). While providing intuitive insights, they fall short of providing a comprehensive analysis of IFT’s underlying mechanisms.

Meanwhile, recent efforts have focused on achieving automated alignment, such as self-instruction-tuning (Sun et al., 2024; Guo et al., 2024), self-rewarding (Yuan et al., 2024) and super-alignment (Burns et al., 2023). Despite repeated validations of their effectiveness, there remains limited understanding of their success.

3 Knowledge Intervention Framework

During the process of IFT, the potential transfer of behavioral norms and the injection of world knowledge are coupled together. Consequently, prior research on IFT has struggled to distinguish the relative effects of these two. To further investigate the underlying mechanisms of IFT, this paper designs a knowledge intervention framework to decouple these two factors. The main idea behind our framework is to control the association between knowledge in IFT data and the existing parameter knowledge in LLM, thereby managing the degree of potential world knowledge that would be injected during IFT. Through this, we can decouple the injection of world knowledge and the transfer of behavioral norms, by observing the effects of IFT at varying degrees of world knowledge injection.

Specifically, we select four multi-choice datasets from different domains. For each question in each dataset, we employ in-context learning to probe the internal parameter knowledge for each base LLM. Then, we construct multiple instruction datasets by adjusting the consistency of the knowledge within IFT data and model parameter knowledge for each base model. Finally, we analyze the underlying impact of different degrees of world knowledge

injection on IFT by fine-tuning LLMs under different settings and examining their performance on the homogeneous, in-domain, and out-of-domain test sets. In the following, we will introduce our knowledge intervention framework in detail.

3.1 Domain IFT Corpus Setup

In order to facilitate more efficient knowledge consistency identification and IFT evaluation, we select all IFT corpus in the form of multiple choice from four domains: medicine, history, engineering, and jurisprudence. For medicine, we craft a dataset with 20,000 training, 2,206 testing, and 10 development instances by filtering MedMCQA (Pal et al., 2022) for entries with explanations and one correct answer, and then applying random sampling. For the other three domains, we procure the relevant items from Xiezhi Benchmark (Gu et al., 2024) and held 10 for development and 250 for test ².

To make a comprehensive evaluation, for each domain’s IFT, we construct three types of test sets: **1) homogeneous test set (HOMO)**, which is held out from the same multiple-choice dataset of the domain; **2) in-domain test set (ID)**, including data from MMLU (Hendrycks et al., 2021) that belong to the domain; **3) an out-of-domain test set (OOD)**, comprising instances in MMLU that are from distinct domains. By observing the accuracy performance differences across the three types of test sets, we aim to examine the impact of IFT on various scenarios. Please refer to the Appendix A for more details about our data processing.

3.2 Parameter Knowledge Probing via Few-shot In-context Learning

Our knowledge intervention framework relies on effectively detecting the parameter knowledge of pre-trained LLMs. To this end, this paper leverages few-shot in-context learning (Dong et al., 2022; Brown et al., 2020), which is a widely-used approach for probing the abilities and internal knowledge of pre-trained LLMs (Zhang et al., 2023; Wan et al., 2024), to identify the parameter knowledge of our base LLMs. Specifically, we utilize in-context learning to probe the base model’s response to each data item in domain multi-choice dataset and regard the response as the model’s parameter knowledge for this question.

²The availability of multiple-choice datasets tailored to specific domains is extremely limited. Consequently, from this benchmark, we select the three domains with the largest volume of data.

3.3 Construction of Instruction Data

Upon probing the internal knowledge of pre-trained LLMs, we build instances based on the consistency between world knowledge contained within the domain data and parameter knowledge, thereby constructing different IFT datasets. Specifically, for each domain and base model, we construct IFT data under three settings, including:

- **Harmonious setting**, which consists of data where the embedded world knowledge is consistent with model parameter knowledge. This means the pre-trained LLM can answer correctly under in-context learning. In the learning process under this setting, the model only needs to transfer behavioral norms, without the need to acquire additional world knowledge due to the above consistency.
- **Incompatible setting**, which comprises instances where the pre-trained LLM cannot correctly answer under in-context learning. Due to the complete inconsistency, the model needs to learn both the behavioral norms and the world knowledge during its training phase.
- **Self-aligning setting**, which consists of data in which the queries are exactly the same as those in the incompatible set, but we modify the answers corresponding to each query to match the pre-trained LLM’s internal knowledge. Therefore, under this setting, all responses are incorrect, and the model will not learn any additional world knowledge.

To ensure a fair comparison, we maintain the same size across the three groups of data. Meanwhile, to prevent the potential collapse during model training due to the exclusive use of multiple-choice questions, we generate an explanation for the answer and incorporate an equal proportion of general instruction data sampled from alpaca-gpt4-en (Peng et al., 2023), thereby ensuring a more stable and real IFT.

3.4 Experiment Setup

Base Model We use LLaMA-2-7B, LLaMA-2-13B, LLaMA-2-70B (Touvron et al., 2023b), and Mistral-7B (Jiang et al., 2023) as the base models of our experiments.

Training Details We only calculate loss on outputs, setting epoch to 3, learning rate to $2e^{-5}$ and

batch size to 256. For Mistral-7B, we set learning rate to $1e^{-5}$ ³. We use DeepSpeed ZeRO3 (Rasley et al., 2020) for LLaMA-2-70B and FSDP (Zhao et al., 2023) for the other three. All experiments are implemented on Nvidia A100-80GB GPUs.

4 Exp-I: Does Learning Domain-specific World Knowledge Matter for IFT?

In our first group of experiments, we would like to examine how the additional world knowledge in the IFT corpus affects LLMs. To this end, we conduct experiments under three settings including harmonious, incompatible, and self-aligning. By observing the performance discrepancies among these settings, we analyze the effects of injecting world knowledge into the IFT.

Finding 1. *When encompassing correct world knowledge, IFT data congruent with model parameter knowledge can lead to superior IFT outcomes.*

To show this, we compare the experimental results under two settings: harmonious and incompatible. Note that the datasets of both two settings contain correct world knowledge, meaning each query-response pair aligns with correct knowledge. Therefore, the core distinction between these two settings lies in whether the entailed world knowledge is consistent with the parameter knowledge of the LLM. Specifically, training in the harmonious setting only requires learning behavioral norms without the need for learning any additional knowledge, while training in the incompatible setting requires learning both.

Table 1 reveals that models fine-tuned under the harmonious setting outperform those fine-tuned under the incompatible setting across homogeneous, in-domain, and out-of-domain evaluations⁴. Specifically, the harmonious setting yields mean performance gains of 11.27%, 14.58%, and 14.57% over the incompatible setting for homogeneous, in-domain, and out-of-domain tests, respectively. The results indicate that utilizing IFT data, which is consistent with model parameter knowledge and does not inject any additional domain-specific world knowledge, yields superior fine-tuned models. This

³Training loss of Mistral-7B using learning rate $2e^{-5}$ does not converge and even spikes.

⁴Regarding the evaluations of HOMO, ID, and OOD: The homogeneous, in-domain, and out-of-domain evaluations for each domain cover different subsets and the level of difficulty varies across these subsets. Therefore, the absolute performance differences between HOMO, ID, and OOD are not directly comparable.

Eval	Medicine			History			Engineering			Jurisprudence		
	HAR	INC	SELF	HAR	INC	SELF	HAR	INC	SELF	HAR	INC	SELF
LLaMA-2-7B												
HOMO	40.22 _{11.77†}	28.45	<u>37.00</u> _{8.55†}	38.80 _{9.20†}	29.60	<u>33.60</u> _{4.00†}	48.40 _{16.00†}	32.40	<u>32.80</u> _{0.40†}	37.60 _{3.60†}	<u>34.00</u>	33.20 _{0.80↓}
ID	<u>39.82</u> _{2.56†}	37.26	41.46 _{4.20†}	54.30 _{23.22†}	31.08	<u>46.02</u> _{14.94†}	42.07 _{11.04†}	<u>31.03</u>	26.21 _{4.82↓}	38.86 _{3.16†}	35.70	<u>36.34</u> _{0.64†}
OOD	<u>39.97</u> _{3.22†}	36.75	40.94 _{4.19†}	39.64 _{8.95†}	30.69	<u>37.22</u> _{6.53†}	40.38 _{12.12†}	28.26	<u>29.17</u> _{0.91†}	38.49 _{3.93†}	34.56	<u>34.88</u> _{0.32†}
LLaMA-2-13B												
HOMO	40.83 _{4.78†}	<u>36.05</u>	34.41 _{1.64↓}	48.40 _{16.00†}	32.40	<u>43.60</u> _{11.20†}	58.00 _{20.80†}	37.20	<u>55.20</u> _{18.00†}	44.00 _{11.60†}	32.40	<u>37.60</u> _{5.20†}
ID	55.43 _{20.37†}	35.06	<u>52.13</u> _{17.07†}	68.28 _{22.15†}	46.13	<u>64.09</u> _{17.96†}	45.52 _{15.86†}	29.66	<u>40.00</u> _{10.34†}	54.77 _{16.22†}	38.55	<u>52.77</u> _{14.22†}
OOD	54.21 _{18.44†}	35.77	<u>50.98</u> _{15.21†}	51.30 _{13.32†}	37.98	<u>49.06</u> _{11.08†}	52.15 _{16.21†}	35.94	<u>51.12</u> _{15.18†}	50.83 _{11.57†}	39.26	<u>48.27</u> _{9.01†}
LLaMA-2-70B												
HOMO	47.95 _{5.41†}	42.54	<u>46.03</u> _{3.49†}	59.20 _{17.20†}	42.00	<u>51.60</u> _{9.60†}	62.40 _{7.20†}	55.20	<u>57.60</u> _{2.40†}	55.20 _{7.60†}	47.60	<u>51.60</u> _{4.00†}
ID	65.37 _{3.97†}	61.40	<u>63.11</u> _{1.71†}	82.37 _{11.08†}	71.29	<u>81.29</u> _{10.00†}	55.17 _{15.86†}	39.31	<u>54.48</u> _{15.17†}	67.69 _{5.48†}	62.21	<u>67.52</u> _{5.31†}
OOD	65.34 _{4.99†}	60.35	<u>63.93</u> _{3.58†}	63.63 _{5.69†}	57.94	<u>63.54</u> _{5.60†}	65.62 _{6.41†}	59.21	<u>64.75</u> _{5.54†}	61.90 _{4.87†}	57.03	<u>61.45</u> _{4.42†}
Mistral-7B												
HOMO	49.80 _{15.12†}	34.68	<u>35.02</u> _{0.34†}	46.80 _{13.60†}	33.20	<u>40.80</u> _{7.60†}	59.60 _{11.20†}	48.40	<u>55.20</u> _{6.80†}	48.00 _{9.20†}	38.80	<u>43.60</u> _{4.80†}
ID	58.17 _{16.40†}	41.77	<u>51.83</u> _{10.06†}	67.74 _{38.39†}	29.35	<u>50.11</u> _{20.76†}	44.83 _{13.80†}	31.03	<u>42.07</u> _{11.04†}	55.21 _{13.78†}	41.43	<u>49.38</u> _{7.95†}
OOD	54.48 _{14.01†}	40.47	<u>47.81</u> _{7.34†}	53.07 _{20.09†}	32.98	<u>45.07</u> _{12.09†}	50.49 _{8.60†}	41.89	<u>44.51</u> _{2.62†}	52.42 _{11.49†}	40.93	<u>48.88</u> _{7.95†}

Table 1: The performance of the four base LLMs after fine-tuning under harmonious (HAR), incompatible (INC), and self-aligning (SELF) settings. For each domain and base model, models fine-tuned on the harmonious dataset and on the self-aligning dataset achieve superior performance compared to those fine-tuned on the incompatible dataset, across all scenarios including homogeneous (HOMO), in-domain (ID), and out-of-domain (OOD) evaluations. The best/second-best performance for each domain and base model in each evaluation is highlighted in bold/underline. The arrows indicate the differences compared to the incompatible setting.

conclusion holds true across homogeneous, in-domain, and out-of-domain evaluations.

Finding 2. *Using IFT data that aligns with model parameter knowledge yet is erroneous yields better performance than employing those that are correct but incongruent with model parameter knowledge.*

To further investigate the impact of learning domain-specific world knowledge on IFT, we conduct a more direct comparative experiment between the self-aligning and incompatible settings. For each domain and pre-trained LLM, the two settings’ datasets use identical queries that the model cannot answer correctly under in-context learning and have different responses: the incompatible dataset’s responses are correct, reflecting world knowledge, while the self-aligning dataset’s responses represent model parameter knowledge, are incorrect.

Table 1 compares results under the self-aligning and incompatible settings. Surprisingly, despite the self-aligning dataset containing only incorrect answers, models fine-tuned on it significantly outperform those using the incompatible dataset, which requires learning inconsistent world knowledge. The performance difference is notable, with the former achieving an average increase of 5.25%, 9.78%, and 6.97% in homogeneous, in-domain, and out-of-domain evaluations, respectively.

This finding emphatically indicates that injecting additional domain knowledge through IFT also

fails to bring effective improvements to LLMs even in homogeneous and in-domain evaluations. Conversely, maintaining consistency with model parameter knowledge, that is, without injecting any additional world knowledge through the self-aligning setting, can yield superior results. Moreover, this advantage holds true across all evaluations, as well as different model sizes and architectures.

The above results demonstrate a significant decline of performance in models fine-tuned using data that contain correct world knowledge but conflict with model parameter knowledge, compared to using consistent IFT data aligned with parameter knowledge. This suggests that introducing additional world knowledge through IFT, in cases where it is inconsistent with the parameter knowledge, may not yield the benefits we anticipate. Therefore, the core role of IFT may lie in facilitating the transfer of behavioral norms, rather than injecting additional domain-specific world knowledge. To further validate this conclusion, in the next section, we design a novel method to decouple the conflict knowledge contained in IFT data and present further analysis of this issue.

5 Exp-II: IFT with Contextualized Knowledge

In this section, we introduce a new analysis method called contextualized knowledge decoupling to fur-

Model	IFT	HOMO	ID	OOD	Overall
LLaMA-2-7B	Vanilla	31.11	33.77	32.56	32.48
	Contextualized	37.62	43.70	40.61	40.65
LLaMA-2-13B	Vanilla	34.51	37.35	37.24	36.37
	Contextualized	41.47	49.48	46.59	45.84
Mistral-7B	Vanilla	38.77	35.90	39.07	37.91
	Contextualized	37.22	43.47	44.99	41.89
Average	Vanilla	34.80	35.67	36.29	35.59
	Contextual	38.77	45.55	44.06	42.79

Table 2: The performance comparison between models fine-tuned with vanilla IFT and models fine-tuned with contextualized IFT respectively.

ther investigate the impact of inconsistent knowledge during IFT. This approach involves explicitly providing relevant world knowledge needed to answer a query within the context of the query itself. Under this paradigm, the model no longer needs to learn knowledge during IFT, but only needs to use the knowledge in the context to answer in the expected behavioral norm. This method helps prevent the model from learning additional world knowledge during IFT, which is inconsistent with parameter knowledge, thus separating knowledge injection from behavioral norm transfer.

To this end, we start with the data of the incompatible setting to construct a dataset with contextualized knowledge. Specifically, given an instruction-answer pair in the incompatible group, we employ GPT-3.5⁵ to generate the world knowledge that is required to answer the instruction. The knowledge is then concatenated with the original instruction, as well as the answer, to construct an augmented pair. Finally, we use the constructed data to fine-tune LLMs and compare them with the models fine-tuned with vanilla IFT.

Finding 3. *Ensuring that the model does not learn world knowledge conflicting with parameter knowledge during IFT enhances the effectiveness of IFT.*

Table 2 shows the results of contextualized knowledge decoupling in three base models⁶. From the table, it can be observed that fine-tuning the model with data using explicit contextualized knowledge significantly mitigates the adverse effects caused by inconsistencies between parameter knowledge and world knowledge in IFT data. Compared to vanilla IFT using incompatible data, our

⁵<https://openai.com/>

⁶In this experimental setting, "vanilla" is equivalent to "incompatible" group.

method achieves an average improvement of 8.16% on LLaMA-2-7B, 9.48% on LLaMA-2-13B, and 3.98% on Mistral-7B. Except in rare cases, IFT using data with contextualized knowledge can significantly improve the effect of IFT in the incompatible setting across homogeneous, in-domain, and out-of-domain evaluations.

The results indicate that we should not force the model to learn additional inconsistent knowledge when the world knowledge in IFT data is inconsistent with model parameter knowledge. Instead, by decoupling knowledge learning and the transfer of behavioral norms, the problems caused by the above knowledge conflicts can be effectively alleviated. This further verifies our observation in Section 4 that additional world knowledge injection may be ineffective or even harmful for IFT. During IFT, the model should focus on transferring the behavioral norms relying on the existing parameter knowledge and the regularity in the IFT data, rather than learning additional world knowledge. Therefore, for RQ1, we conclude:

Conclusion 1. *For IFT, there is little, if not even causing additional damage, benefits from the learning of world knowledge incongruent with parameter knowledge.*

6 Exp-III: Is Consistency All You Need?

The above findings appear to suggest a conclusion: For better transfer of behavioral norms, we should employ IFT data that completely aligns with model parameter knowledge without any additional world knowledge. To substantiate this hypothesis, we design a new set of experiments.

Specifically, by adjusting the proportion of samples derived from the incompatible and the self-aligning group, we aim to adjust the ratios of consistency between the world knowledge in the IFT data and model parameter knowledge, thereby observing the variations in IFT outcomes under different consistency ratios.

Finding 4. *Employing IFT data that is fully consistent with model parameter knowledge does not necessarily result in optimal performance.*

Figure 4 displays the results on three base models under different consistency ratios across four domains. From the figure, we can see that: 1) The performance of the incompatible group (i.e.,

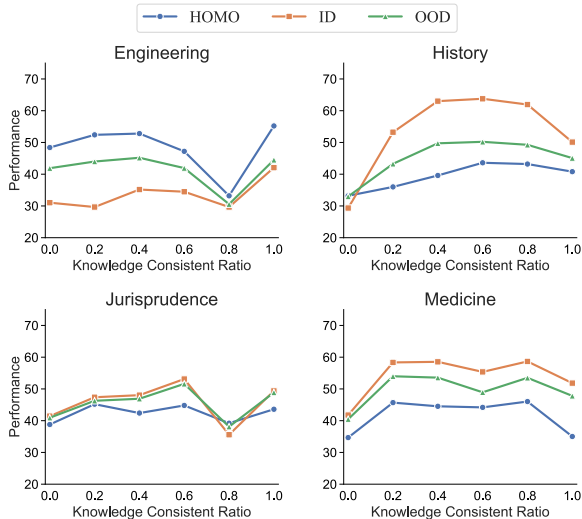


Figure 2: The performance of Mistral-7B fine-tuned with instruction datasets of varying consistency ratios. Each dataset is composed of a mixture of incompatible and self-aligning data, and the consistency ratio represents the proportion of self-aligning samples. Note that a consistency ratio of 0 signifies that all data samples are incompatible, whereas a ratio of 1 indicates exclusively self-aligning data. The results of other base models are presented in the Appendix C due to page limitations.

ratio=0) is indeed poor, which aligns with our previous conclusion; 2) Conversely, relying solely on IFT data that completely aligns with the model’s parameter knowledge (i.e., ratio=1) fails to ensure superior performance across a broad range of scenarios; 3) Optimal performance is most frequently achieved through a balanced integration of incompatible and self-aligning data. The ideal proportion of this combination varies across different base models and domains.

The observations suggest that while there is a significant impact of the consistency between the knowledge within the IFT data and the parameter knowledge of the original model on the performance of the fine-tuned model, this consistency is not the fundamental factor influencing IFT performance. Therefore, a deeper exploration into the mechanisms underlying knowledge consistency is essential to identify the true determinants of IFT’s effectiveness.

7 Exp-IV: Rethinking Consistency: What Really Matters for IFT?

To further explore the underlying mechanisms of IFT, we analyze the knowledge discrepancies of different base LLMs before and after IFT on vari-

Model	HOMO		ID		OOD	
	r	p-value	r	p-value	r	p-value
Mistral-7B	0.78	0.00	0.81	0.00	0.82	0.00
LLaMA-2-7B	0.27	0.14	0.19	0.30	0.21	0.24
LLaMA-2-13B	0.56	0.00	0.78	0.00	0.87	0.00
All	0.43	0.00	0.57	0.00	0.48	0.00

Table 3: The Spearman partial correlation analysis between the model performance after fine-tuning, and the knowledge consistency between base model and fine-tuned model. The analysis is controlled with the base model’s performance on each test set. r and p-value denote partial correlation coefficient and significance respectively. For LLaMA-2-13B and Mistral-7B, p-values significantly lower than 0.05 indicates a high level of confidence.

ous evaluation data to observe the extent of internal knowledge alternation triggered by IFT. Specifically, for each sample within the test data, we initially compute the Pearson correlation coefficient between the ranking of the original model’s predictions on choices through in-context learning probing, and those provided by the fine-tuned model. Building upon this, we calculate the average Pearson correlation coefficient across each test set and subsequently compare it with the performance on the same test set of the fine-tuned model. Through this experimental analysis, we aim to observe how the knowledge changes induced by IFT affect the ultimate effect of IFT.

To this end, we employ a partial correlation assessment (Spearman, 1961) to analyze the aforementioned Pearson correlation and model performance. Specifically, we utilize a total of 96 models, which include 72 models from Exp-III, in addition to models fine-tuned with consistency ratios of 0.05 and 0.1, conducting tests on homogeneous, in-domain, and out-of-domain data. To eliminate the potential influence of the base model’s performance on the correlation analysis, ensuring that the results only reflect the differences brought about by varying IFT data, we treat the base model’s performance as a control variable. By employing partial correlation assessment under the constraint of removing this control variable’s influence, we analyze the correlation between the Pearson correlation coefficients of model predictions’ rankings before and after IFT and their ultimate performance. Please refer to the Appendix D for more details about our analysis.

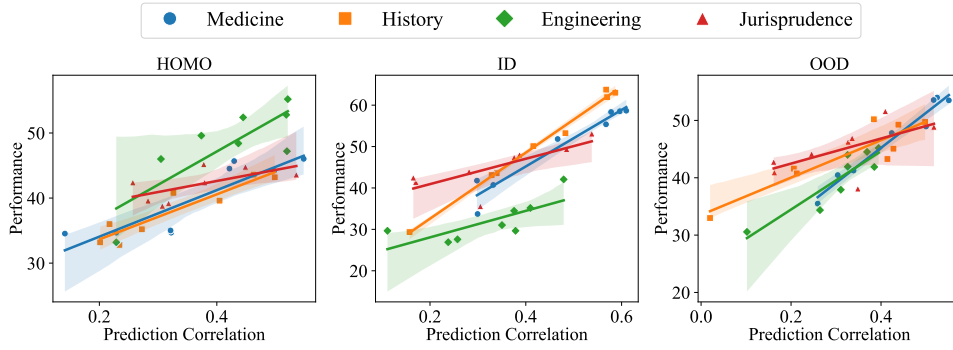


Figure 3: The regression analysis between the model performance after fine-tuning, and the knowledge consistency between base model and fine-tuned model. We show the results of Mistral-7B in three evaluations. The grouped linear regression demonstrate the positive correlations between the model performance after IFT and model internal knowledge consistency before and after IFT. Points in the same regression line indicate the results of the same base model fine-tuned with different IFT data of the same domain on the same test set (HOMO, ID, or OOD).

Finding 5. *The consistency of internal knowledge within a model before and after IFT is the key factor affecting the performance of the fine-tuned model.*

Figure 3 presents the results of grouped regression analysis conducted on Mistral-7b. Table 3 displays the correlation coefficients and significance for the partial correlation analysis across different base models and different evaluations. From the aforementioned figure and table, a significant trend can be observed, namely that the correlation of predictions made by models before and after IFT on a given evaluation has a substantial impact on the final performance of the fine-tuned models on that evaluation. It is noteworthy that this phenomenon holds true across homogeneous, in-domain, and out-of-domain evaluations. This implies that even for homogeneous and in-domain test sets, which are within the same field as the IFT data, the injection of additional world knowledge and the alteration of internal parameter knowledge through IFT do not contribute to enhancing the final performance of the models. Instead, maintaining consistency in the knowledge of models before and after IFT significantly positively influences the performance of the fine-tuned models.

To further investigate whether this finding is the underlying mechanism responsible for Finding 4, we explore the impact of different consistency data on the correlation of model knowledge before and after IFT in Exp-III. We analyze the probability distribution of model outputs under different IFT data conditions. Table 4 presents the related experimental results, which reveal that: IFT using data containing world knowledge completely in-

Model	Best	Self-aligning	Incompatible
Mistral-7B	0.24	0.34	0.37
LLaMA-2-7B	0.16	0.93	0.51
LLaMA-2-13B	0.18	0.68	0.34

Table 4: KL divergence between the prediction distribution of fine-tuned model under zero-shot setting and the prediction distribution of base model using in-context learning probing. "Best" denotes the model that exhibits the best average performance across three evaluations.

consistent with the parameter knowledge evidently leads to a divergence in the internal knowledge of the fine-tuned model from that of the original model, thereby impairing the performance of the fine-tuned model. Furthermore, IFT using data containing world knowledge completely consistent with the parameter knowledge also may result in a divergence in the knowledge distribution between the original and fine-tuned models. Specifically, fine-tuning with only consistent IFT data may steer the model towards a sharp knowledge distribution, whereas the original model’s parameter knowledge exists as a relatively smooth distribution. Conversely, using a middle setting that mixes incompatible and self-aligning data allows the optimization process to maintain model parameter knowledge unchanged while preserving the distribution’s smoothness, thereby enabling the fine-tuned model outputs to more closely resemble those of the original model, ultimately yielding better performance.

The experimental results mentioned above demonstrate that the key to the performance of fine-tuned models lies in the consistency of model parameter knowledge before and after IFT. In fact,

the core mechanism underlying the superior performance observed in both [Finding 1](#) and [Finding 2](#) for the IFT data is attributed to the maintenance of consistency in model knowledge before and after IFT. Therefore, for [RQ2](#) raised above, we conclude:

Conclusion 2. *The essence of an effective IFT lies in maintaining the consistency of model parameter knowledge before and after IFT.*

8 Conclusion and Discussion

Our experiments and conclusions indicate that the core function of IFT is not to learn domain-specific world knowledge. Instead, learning world knowledge that is inconsistent with model parameter knowledge actually undermines the performance of the model in all evaluations from homogeneous, in-domain to out-of-domain. Furthermore, we discover that the consistency of model parameter knowledge before and after IFT (i.e., the knowledge probed through in-context learning before IFT and the knowledge exhibited under zero-shot setting after IFT) plays a crucial role in determining the ultimate performance of the fine-tuned model. These two findings unveil a fundamental mechanism of IFT, that is, IFT is not a supervised, domain-specific knowledge *learning* process, but a process of *self-aligning* instruction with the already existing parameter knowledge of LLMs. Therefore, the ultimate determinant of IFT effect is not the extent of domain knowledge injection, but rather whether the IFT process can facilitate more effective self-aligning, thereby enhancing the expression of the model’s parameter knowledge under the zero-shot question-answering paradigm after IFT.

Our discovery not only provides guidance for future IFT data construction, model training, and model evaluation but also provide robust support for some very recent studies. For instance, super-alignment ([Burns et al., 2023](#)) aims to use a weak model to guide a strong model’s alignment. Our conclusion proves that it is entirely possible to use a weak model with less knowledge to guide a strong, more knowledgeable model for IFT. Our conclusions elucidate the viable underlying factors for self-instruction-tuning ([Sun et al., 2024](#); [Guo et al., 2024](#)), self-rewarding ([Yuan et al., 2024](#)) and consistent alignment ([Wan et al., 2024](#)) and provide a solid foundation for the future development of these studies.

Limitations

In order to facilitate probing model parameter knowledge, we currently focus on multiple-choice questions. In the future, we plan to extend our framework to free-style generation. Besides, due to hardware limitations, the vast majority of experiments are conducted on models with about 10B parameters, we only explore 70B models in several experiments. Repeating our study on larger models in more domains will contribute to a deeper understanding of the IFT for larger models.

Acknowledgements

We sincerely thank all anonymous reviewers for their insightful comments and valuable suggestions. This research work is supported by CAS Project for Young Scientists in Basic Research (Grant No.YSBR-040), the National Natural Science Foundation of China under Grants no. 62122077 and 62106251, and the Basic Research Program of ISCAS, Grant No. ISCAS-JCZD-202303.

References

- Christopher M Bishop. 2006. *Pattern recognition and machine learning*, volume 4. Springer.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, and Jared D. Kaplan. 2020. [Language models are few-shot learners](#). *Advances in Neural Information Processing Systems*, 33:1877–1901.
- Collin Burns, Pavel Izmailov, Jan Hendrik Kirchner, Bowen Baker, Leo Gao, Leopold Aschenbrenner, Yining Chen, Adrien Ecoffet, Manas Joglekar, Jan Leike, et al. 2023. [Weak-to-strong generalization: Eliciting strong capabilities with weak supervision](#). *arXiv preprint arXiv:2312.09390*.
- Hao Chen, Yiming Zhang, Qi Zhang, Hantao Yang, Xi-aomeng Hu, Xuetao Ma, Yifan Yanggong, and Junbo Zhao. 2023a. [Maybe only 0.5% data is needed: A preliminary exploration of low training data instruction tuning](#). *arXiv preprint arXiv:2305.09246*.
- Yirong Chen, Xiaofen Xing, Jingkai Lin, Huimin Zheng, and Zhenyu Wang. 2023b. [Soulchat: Improving llms’ empathy, listening, and comfort abilities through fine-tuning with multi-turn empathy conversations](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 1170–1183, Singapore. Association for Computational Linguistics.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2024. [Scaling instruction-finetuned language models](#). *Journal of Machine Learning Research*, 25(70):1–53.

- Jiaxi Cui, Zongjian Li, Yang Yan, Bohua Chen, and Li Yuan. 2023. [Chatlaw: Open-source legal large language model with integrated external knowledge bases](#). *arXiv preprint arXiv:2306.16092*.
- Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, and Zhifang Sui. 2022. [A survey on in-context learning](#). *arXiv preprint arXiv:2301.00234*.
- Lingyu Gao, Aditi Chaudhary, Krishna Srinivasan, Kazuma Hashimoto, Karthik Raman, and Michael Bendersky. 2023. [Ambiguity-aware in-context learning with large language models](#). *arXiv preprint arXiv:2309.07900*.
- Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep Learning*. MIT Press.
- Zhouhong Gu, Xiaoxuan Zhu, Haoning Ye, Lin Zhang, Jianchen Wang, Yixin Zhu, Sihang Jiang, Zhuozhi Xiong, Zihan Li, Weijie Wu, et al. 2024. [Xiezhi: An ever-updating benchmark for holistic domain knowledge evaluation](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 18099–18107.
- Arnav Gudibande, Eric Wallace, Charlie Snell, Xinyang Geng, Hao Liu, Pieter Abbeel, Sergey Levine, and Dawn Song. 2023. [The false promise of imitating proprietary llms](#). *arXiv preprint arXiv:2305.15717*.
- Hongyi Guo, Yuanshun Yao, Wei Shen, Jiaheng Wei, Xiaoying Zhang, Zhaoran Wang, and Yang Liu. 2024. [Human-instruction-free llm self-alignment with limited samples](#). *arXiv preprint arXiv:2401.06785*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring massive multitask language understanding](#). In *International Conference on Learning Representations*.
- Srinivasan Iyer, Xi Victoria Lin, Ramakanth Pasunuru, Todor Mihaylov, Daniel Simig, Ping Yu, Kurt Shuster, Tianlu Wang, Qing Liu, Punit Singh Koura, et al. 2022. [Opt-impl: Scaling language model instruction meta learning through the lens of generalization](#). *arXiv preprint arXiv:2212.12017*.
- Joel Jang, Seungone Kim, Seonghyeon Ye, Doyoung Kim, Lajanugen Logeswaran, Moontae Lee, Kyungjae Lee, and Minjoon Seo. 2023. [Exploring the benefits of training expert language models over instruction tuning](#). In *Proceedings of the 40th International Conference on Machine Learning*, pages 14702–14729.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. [Mistral 7b](#). *arXiv preprint arXiv:2310.06825*.
- Yunxiang Li, Zihan Li, Kai Zhang, Ruilong Dan, Steve Jiang, and You Zhang. 2023. [Chatdoctor: A medical chat model fine-tuned on a large language model meta-ai \(llama\) using medical domain knowledge](#). *Cureus*, 15(6).
- Bill Yuchen Lin, Abhilasha Ravichander, Ximing Lu, Nouha Dziri, Melanie Sclar, Khyathi Chandu, Chandra Bhagavatula, and Yejin Choi. 2024. [The unlocking spell on base LLMs: Rethinking alignment via in-context learning](#). In *The Twelfth International Conference on Learning Representations*.
- Alisa Liu, Xiaochuang Han, Yizhong Wang, Yulia Tsvetkov, Yejin Choi, and Noah A Smith. 2024. [Tuning language models by proxy](#). *arXiv preprint arXiv:2401.08565*.
- Sewon Min, Xinxu Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. [Rethinking the role of demonstrations: What makes in-context learning work?](#) In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11048–11064.
- Ankit Pal, Logesh Kumar Umapathi, and Malaikannan Sankarasubbu. 2022. [Medmcqa: A large-scale multi-subject multi-choice dataset for medical domain question answering](#). In *Conference on Health, Inference, and Learning*, pages 248–260. PMLR.
- Baolin Peng, Chunyuan Li, Pengcheng He, Michel Galley, and Jianfeng Gao. 2023. [Instruction tuning with gpt-4](#). *arXiv preprint arXiv:2304.03277*.
- Jeff Rasley, Samyam Rajbhandari, Olatunji Ruwase, and Yuxiong He. 2020. [Deepspeed: System optimizations enable training deep learning models with over 100 billion parameters](#). In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD '20*, pages 3505–3506, New York, NY, USA. Association for Computing Machinery.
- Ming Shen. 2024. [Rethinking data selection for supervised fine-tuning](#). *arXiv preprint arXiv:2402.06094*.
- C. Spearman. 1961. *The Proof and Measurement of Association Between Two Things*. Studies in Individual Differences: The Search for Intelligence. Appleton-Century-Crofts, East Norwalk, CT, US.
- Zhiqing Sun, Yikang Shen, Qinzhong Zhou, Hongxin Zhang, Zhenfang Chen, David Cox, Yiming Yang, and Chuang Gan. 2024. [Principle-driven self-alignment of language models from scratch with minimal human supervision](#). *Advances in Neural Information Processing Systems*, 36.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023a. [Llama: Open and efficient foundation language models](#). *arXiv preprint arXiv:2302.13971*.

- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023b. [Llama 2: Open foundation and fine-tuned chat models](#). *arXiv preprint arXiv:2307.09288*.
- Fanqi Wan, Xinting Huang, Leyang Cui, Xiaojun Quan, Wei Bi, and Shuming Shi. 2024. [Mitigating hallucinations of large language models via knowledge consistent alignment](#). *arXiv preprint arXiv:2401.10768*.
- Aiyuan Yang, Bin Xiao, Bingning Wang, Borong Zhang, Ce Bian, Chao Yin, Chenxu Lv, Da Pan, Dian Wang, Dong Yan, et al. 2023. [Baichuan 2: Open large-scale language models](#). *arXiv preprint arXiv:2309.10305*.
- Weizhe Yuan, Richard Yuanzhe Pang, Kyunghyun Cho, Sainbayar Sukhbaatar, Jing Xu, and Jason Weston. 2024. [Self-rewarding language models](#). *arXiv preprint arXiv:2401.10020*.
- Hanning Zhang, Shizhe Diao, Yong Lin, Yi R Fung, Qing Lian, Xingyao Wang, Yangyi Chen, Heng Ji, and Tong Zhang. 2023. [R-tuning: Teaching large language models to refuse unknown questions](#). *arXiv preprint arXiv:2311.09677*.
- Yanli Zhao, Andrew Gu, Rohan Varma, Liang Luo, Chien-Chin Huang, Min Xu, Less Wright, Hamid Shojanazeri, Myle Ott, Sam Shleifer, et al. 2023. [Pytorch fsdp: experiences on scaling fully sharded data parallel](#). *arXiv preprint arXiv:2304.11277*.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. 2024. [Judging llm-as-a-judge with mt-bench and chatbot arena](#). *Advances in Neural Information Processing Systems*, 36.
- Chunting Zhou, Pengfei Liu, Puxin Xu, Srinivasan Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, Lili Yu, et al. 2024. [Lima: Less is more for alignment](#). *Advances in Neural Information Processing Systems*, 36.

A Data Construction

A.1 Train

As mentioned in Section 3, we collect four multi-choice question datasets from different domains: medicine, history, engineering, and jurisprudence. The quantity and split of the dataset for each domain are detailed in Table 5.

Domain	Dev	Test	Train
Medicine	10	1462	20000
History	10	250	8605
Jurisprudence	10	250	6510
Engineering	10	250	4805

Table 5: The number of instances in the development, test, and train sets for each domain.

For each item in each train set, we employ five demonstrations from the corresponding dev set for in-context learning to assess the parameter knowledge of each base model. The probing prompt is shown in Table 6.

<p>The following are multiple choice questions about <domain>. Please choose the correct answer.</p> <p><i>5-shot</i></p> <p><question1> <question1 options> Answer:<question1 answer></p> <p><demo2> . . . <demo5></p> <p><question> <question options> Answer:</p>
--

Table 6: Prompt design based on Yang et al. (2023) for probing model parameter knowledge.

Because model’s responses are influenced by the selection and order of demonstrations in in-context learning (Gao et al., 2023; Min et al., 2022), we regard responses with a confidence level exceeding 0.5 as reflective of the model’s parameter knowledge to ensure the reliability of identification. For every domain and base model combination, we establish train datasets on the three different settings as described in Section 3. We report the amount of training data for each combination in Table 7. We employ the prompt in Table 8 for constructing instruction-response pairs.

Model	Engineering	History	Jurisprudence	Medicine
LLaMA-2-7B	738	996	1033	2507
LLaMA-2-13B	668	1236	676	1782
LLaMA-2-70B	712	1072	1002	2518
Mistral-7B	534	838	639	1400

Table 7: The number of train sets for each domain and base model combination.

For <explanation> in the prompt, if the answer is consistent with the model’s parameter knowledge but incorrect, we employ prompt in Table 9 for prompting the base model to generate explanations for its

Input

The following are multiple choice questions about domain. Please choose the correct answer.

<question>

<question options>

Output

<answer>

Explanation: <explanation>

Table 8: Format of instruction-response pair for Vanilla IFT.

choice; if the answer is correct, we directly use the explanation in MedMCQA for medicine and employ prompt in Table 10 for prompting GPT3.5 to generate explanations for the golden answer for the other three domains.

For stable and real IFT, we incorporate an equal proportion of general instruction data sampled from alpaca-gpt4-en (Peng et al., 2023) which is identical for each domain and base model combination. We use vicuna-v1.5 (Zheng et al., 2024) format to train models.

A.2 Test

To compare the performance across three settings for each domain and base model combination, we devise three evaluation types: homogeneous, in-domain, and out-of-domain tests. The first evaluation involves a holdout from the same multi-choice question set, while the subsequent two, in-domain and out-of-domain, are derived from MMLU’s splits. The specific details of the splits for each domain are detailed in Table 11.

B Details about Contextualized IFT

In Section 5, we use the prompt in Table 12 to employ GPT-3.5 to provide the knowledge required for the inconsistent IFT data. For contextualized IFT, we employ the prompt in Table 13 for building instruction-response pairs to train models. Note we do not provide the knowledge required for the instruction during the testing phase.

C Performance of Models Fine-tuned in Different Consistency Ratios

In Exp-III, due to page limitations, we only report the performance of Mistral-7B fine-tuned with instruction datasets of varying consistency ratios. In Figure 4, we supplement the results of the other two base LLMs: LLaMA-2-7B and LLaMA-2-13B.

D Details about Spearman Partial Correlation Analysis

Partial correlation measures the degree of association between two variables (here, the model performance after IFT and the knowledge consistency between the base model and fine-tuned model) while controlling the effect of one or more additional variables (the accuracy on the same test set of the base model in our case).

To calculate the Spearman partial correlation between the model performance after IFT and the knowledge consistency between the base model and fine-tuned model, we first calculate the knowledge consistency, measured by the correlation coefficient of model prediction ranking before and after IFT using the following formula.

$$pr = \frac{1}{n} \sum_{i=1}^{|D_{\text{test}}|} \text{pearson}(\text{sort}(m_{\text{base}}, q_i, c_i), \text{sort}(m_{\text{tuned}}, q_i, c_i))$$

Instruction

Below is a list of conversations between a human and an AI assistant (you).
Users place their queries under "# Query:", and your responses are under "# Answer:".
You are a helpful, respectful, and honest assistant. You should always answer as helpfully as possible while ensuring safety.
Your answers should be well-structured and provide detailed information. They should also have an engaging tone.
Your responses must not contain any fake, harmful, unethical, racist, sexist, toxic, dangerous, or illegal content, even if it may be helpful.
Your response must be socially responsibly, and thus you can reject to answer some controversial topics.

Query:

```Can you tell me some common types of renewable energy sources? ```

**# Answer:**

```

Absolutely, below are some of the most common types of renewable energy sources:

1. Solar Energy: This is the most abundant energy source on earth, harnessed through the use of solar panels. These panels convert sunlight into electricity without any moving parts, noise, pollution, or damage to the environment.
2. Wind Energy: Wind turbines convert the kinetic energy in the wind into mechanical power. This mechanical power can be used for specific tasks (such as pumping water) or converted into electricity to power homes, businesses, and schools.
3. Hydropower: Generated by using electricity generators to capture the energy from falling or fast-moving water. This renewable source can come from various ways, including a large dam on a river, tidal or wave energy from the ocean, or using small scale turbines in streams.
4. Geothermal Energy: This type of energy is generated from the heat deep within the Earth. This heat can be used directly for heating buildings or to generate electricity. It is continuously produced inside the Earth and is nearly as reliable as the tides.
5. Biomass Energy: Biomass is organic material that comes from plants and animals, and it contains stored energy from the sun. This energy can be burned directly or converted into biofuel which can burn more efficiently.

Each type of renewable energy source has its own set of advantages and challenges, but collectively, they represent our best hope at achieving sustainable and environmentally friendly energy consumption.

```

**# Query:**

```

Below is a multiple-choice question and the answer. Please give the explanation.

Question: <question>

Choices: <question options>

Answer: <answer>

```

**# Answer:**

Table 9: Prompt design based on URIAL (Lin et al., 2024) for employing base model to generate explanations for its predicted answer.

The following is a multi choice question about <domain>.

<question>

<choice>

The answer is "<answer>". Please explain why.

Table 10: Prompt design for employing GPT-3.5 to generate explanations for the golden answer.

where  $q_i$  and  $c_i$  represent the  $i$ -th question and its corresponding set of candidate choices. The terms  $m_{\text{base}}$  and  $m_{\text{tuned}}$  denote the base model and its fine-tuned model, respectively.  $D_{\text{test}}$  represents any test set like HOMO, ID, or OOD.  $|D_{\text{test}}|$  denotes the size of the test set. The ‘sort’ function ranks the candidate choices based on the predicted probability distribution from the input model.  $pr$  reflects the average

Domain	In-domain Subcategories
Engineering	["electrical_engineering"]
History	["high_school_european_history", "high_school_us_history", "high_school_world_history", "prehistory"]
Jurisprudence	["econometrics", "high_school_geography", "high_school_government_and_politics", "high_school_macroeconomics", "high_school_microeconomics", "high_school_psychology", "human_sexuality", "international_law", "jurisprudence", "professional_law", "sociology", "public_relations", "professional_psychology", "security_studies", "us_foreign_policy"]
Medicine	["anatomy", "clinical_knowledge", "college_medicine", "human_aging", "medical_genetics", "nutrition", "professional_medicine", "virology"]

Table 11: In-domain and out-of-domain split of MMLU subcategories for our four experimental domains. The out-of-domain subcategories comprise the remaining subcategories not included in the in-domain classification.

<p>Given a multi-choice question and the answer, please write a short piece of evidence to support it so that a layman who has read the evidence you give can answer the question correctly. If your response contains words "listed", "option" or "choice" like "among the listed/given options", you will be penalized.</p> <p><b>Question:</b> &lt;question&gt; &lt;question options&gt;</p> <p><b>Answer:</b> &lt;answer&gt;</p> <p><b>Evidence:</b></p>
------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Table 12: Prompt design for employing GPT-3.5 to generate the knowledge required for the instruction.

<p><b># Input</b> The following are multiple choice questions about &lt;domain&gt;. Given the context. Please choose the correct answer.</p> <p>&lt;context&gt; &lt;question&gt; &lt;question options&gt;</p> <p><b># Output</b> &lt;answer&gt;</p>
---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Table 13: Format of instruction-response pair for our contextualized IFT.

Pearson correlation coefficient between the rankings predicted by the base model and the fine-tuned model of each question.

For each base model, its every fine-tuned model, and each test set, we calculate the corresponding above  $pr$  and organize them along with the corresponding accuracy on the same test set of the base model and fine-tuned model to three lists, denoted as  $pr_s$ ,  $p_{base}$  and  $p_{tuned}$ . Then we compute the Spearman partial correlation. The formulas are:

$$\begin{aligned}
 e(pr_s) &= pr_s - \hat{pr}_s(p_{base}) \\
 e(p_{tuned}) &= p_{tuned} - \hat{p}_{tuned}(p_{base}) \\
 r, p\text{-value} &= \text{spearman}(e(pr_s), e(p_{tuned}))
 \end{aligned}$$

Here,  $\hat{pr}_s(p_{base})$  and  $\hat{p}_{tuned}(p_{base})$  present the estimated value of  $pr_s$  and  $p_{tuned}$  predicted based on

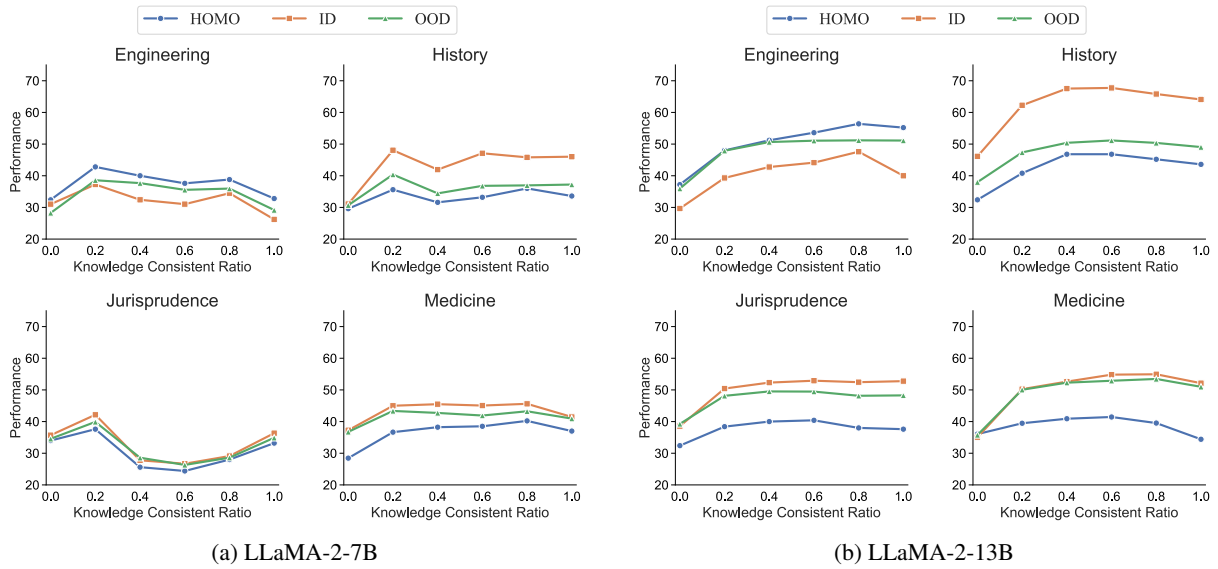


Figure 4: The performance of LLaMA-2-13B and LLaMA-2-7B fine-tuned with instruction datasets of varying consistency ratios. Each dataset is composed of a mixture of incompatible and self-aligning data, and the consistency ratio represents the proportion of self-aligning samples.

$p_{\text{base}}$ , respectively. The residuals of  $pr_s$  and  $p_{\text{tuned}}$  relative to  $p_{\text{base}}$  are denoted as  $e(pr_s)$  and  $e(p_{\text{tuned}})$ , respectively. We finally calculate Spearman’s correlation coefficient and the corresponding p-values between the residuals, denoted as  $r$  and  $p\text{-value}$ .

In this way, we can find the direct relationship between the model performance after IFT, and the knowledge consistency between the base model and fine-tuned model by removing the influence of base model performance that affects them both.