

# Stands to Reason: Investigating the Effect of Reasoning on Idiomaticity Detection

Dylan Phelps<sup>1</sup>, Thomas Pickard<sup>1</sup>, Rodrigo Wilkens<sup>2</sup>,  
Edward Gow-Smith<sup>1</sup>, Maggie Mi<sup>2</sup>, Marco Idiart<sup>3</sup> Aline Villavicencio<sup>1,2</sup>

<sup>1</sup>Department of Computer Science, The University of Sheffield, United Kingdom

<sup>2</sup>Institute for Data Science and Artificial Intelligence, The University of Exeter, United Kingdom

<sup>3</sup>Federal University of Rio Grande do Sul, Brazil

{drsphelps1, tmpickard1, egow-smith1, zmi1}@sheffield.ac.uk  
{r.wilkens, a.villavicencio}@exeter.ac.uk, marco.idiart@gmail.com

## Abstract

The recent trend towards utilisation of reasoning models has improved the performance of Large Language Models (LLMs) across many tasks which involve logical steps. One linguistic task that could benefit from this framing is idiomaticity detection, as a potentially idiomatic expression must first be understood in relation to the context before it can be disambiguated. In this paper, we explore how reasoning capabilities in LLMs affect idiomaticity detection performance and examine the effect of model size. We evaluate, as open source representative models, the suite of DeepSeek-R1 distillation models ranging from 1.5B to 70B parameters across four idiomaticity detection datasets. We find the effect of reasoning to be smaller and more varied than expected. For smaller models, producing chain-of-thought (CoT) reasoning increases performance from Math-tuned intermediate models, but not to the levels of the base models, whereas larger models (14B, 32B, and 70B) show modest improvements. Our in-depth analyses reveal that larger models demonstrate good understanding of idiomaticity, successfully producing accurate definitions of expressions, while smaller models often fail to output the actual meaning. For this reason, we also experiment with providing definitions in the prompts of smaller models, which we show can improve performance in some cases.

**Keywords:** large language models, idiomaticity detection, prompting, reasoning, evaluation

## 1. Introduction

Large Language Models (LLMs) have been shown to excel at many tasks across many disciplines (Ouyang et al., 2022; Achiam et al., 2023; Grattafiori et al., 2024), including on tasks involving idiomatic expressions, such as idiomaticity detection and multiword expression identification (Phelps et al., 2024; Smádu et al., 2024). More recently, reasoning models, LLMs which generate a chain of thought (CoT) ‘reasoning’ responses before reaching a final answer (Wei et al., 2022), have been shown to outperform traditional LLMs in a range of tasks (OpenAI, 2025). Notably, DeepSeek-R1 and smaller ‘distilled’ models trained on data produced by it (DeepSeek-AI et al., 2025), have matched and exceeded other reasoning models such as OpenAI’s GPT-o1 whilst being open source and offered at a lower price.

The motivation for such reasoning models is that training them to generate outputs in chain-of-thought format allows them “think” step-by-step, working out a final answer incrementally. This output format also allows for potentially higher explainability, since the CoT becomes part of the context used to make the final classification. However, some work has shown that CoT explanations can be unfaithful, misrepresenting the reason for a model’s prediction (Turpin et al., 2023; Lyu et al., 2023).

In the field of computational idiom processing,

reasoning models are of interest as one could imagine such step-by-step thinking to improve the handling of potentially idiomatic expressions (PIEs), where the meaning is ambiguous and determined by context. Using knowledge of a definition as an indicator of a level of understanding, the CoT also allows analysis of A) how well models can define PIEs in their literal or idiomatic senses, and B) how well models can use subsequent reasoning to work out whether a PIE is used idiomatically or literally in a given context.

Therefore, in this paper, we explore how the distilled DeepSeek-R1 models perform on a range of idiomaticity detection tasks. Additionally, we evaluate the reasoning outputs of each model to further explore how they represent and understand idiomaticity. Our research questions are:

1. Does the production of reasoning chains improve the ability of models to detect idiomaticity, and how does this vary across model scale?
2. Do the reasoning outputs (CoT) reflect understanding of the target idiomatic expressions?
3. Can the definitions produced by the larger models be used as an additional knowledge source, improving the performance of the smaller models?

To this end, we find that although the performance of the models improves as the models scale, the effect of adding reasoning generation is smaller and more varied. Particularly, for the smaller models we find that producing CoTs on average reduces the performance across our datasets, whereas for the larger models a small increase in performance can be seen in the reasoning variants. Our manual analysis of the dataset shows that the larger models also have a good understanding of idiomaticity and can reliably produce accurate definitions of the given expressions, while the smaller models often fail to do so. However, even the largest models are hindered by their ability to use context and reasoning to disambiguate PIEs.

Our experiments using definitions from the larger models as a knowledge source show that performance of the smallest models improves by an average of 0.069 macro F1 on FLUTE, but does not affect performance significantly for DICE. These results suggest that this methodology has potential as a knowledge distillation technique for certain tasks.

The paper is structured as follows: §2 presents the methodology and §3 presents the results. §4 and §5 analyse the results, with §6 performing further experiments using the generated definitions.

## 2. Methodology

In this section, we introduce the datasets, models, and approach that we use in this work.

### 2.1. Datasets

To enable comparison between the performance of the newer models to those evaluated in [Phelps et al., 2024](#), we choose to evaluate on the same datasets. Additionally, we include the recently released DICE dataset ([Mi et al., 2025](#)). We provide brief descriptions of each dataset.

**SemEval 2022 Task 2a** ([Tayyar Madabushi et al., 2022](#)) is a binary classification task for detecting idiomaticity of noun compounds within context sentences, with examples in English, Portuguese, and Galician. The test set contains 150 PIEs split equally across the three languages, with a total of 2342 examples. To maximize the number of test examples, we follow [Phelps et al., 2024](#) by combining the few-shot and zero-shot test sets. However, when evaluating we don't provide few-shot examples for any of the instances.

**Figurative Language Understanding through Textual Explanations (FLUTE; [Chakrabarty et al. \(2022\)](#))** presents an English-language Natural Language Inference (NLI) task in which models

should predict whether a premise, containing a figurative expression, follows from a hypothesis containing a correct or incorrect paraphrase. We evaluate only the idiom subset of the dataset, which contains a test set of 250 examples with 69 idioms represented.

**MAGPIE** ([Haagsma et al., 2020](#)) is a multi-class idiomaticity detection dataset where a large number of potentially idiomatic expressions in context must be classified as idiomatic, literal, or other with a split of 70/29/1. The entire dataset is in English, with 1134 expressions represented across 4840 examples.

**Dataset for Idiomatic Contrastive Evaluation (DICE; [Mi et al. \(2025\)](#))** aims to assess the ability of models to use context in idiomaticity detection. Existing datasets fail to assess the role of context in idiom interpretation, as literal meanings often stem from grammatical changes to the form of the idiomatic expression, which allows models to rely on surface cues as a reasoning shortcut instead of true comprehension. To avoid this, DICE contains 2066 sentences that are balanced for sense, in which the form of the expression is kept the same across both literal and figurative uses. Being a newer dataset, the test labels were not publically available when the DeepSeek models were trained, so there is no chance of contamination.

### 2.2. Models and Experiments

The models we evaluate are the suite of DeepSeek-R1 distillation ([DeepSeek-AI et al., 2025](#)) models ranging from 1.5-70B parameters, that have been fine-tuned on a reasoning dataset collected from the larger DeepSeek-R1 model (the **reasoning** models). Whilst processing the input, these models produce chain-of-thought reasoning.

Each of the distilled models is a fine-tuned version of another open source model (the **base** models), and so to compare the performance with and without CoT reasoning, we also run our evaluations on the base models. Namely, we evaluate the Qwen2.5 suite of models (1.5B, 7B, 14B, and 32B parameters) ([Qwen, 2024; Qwen et al., 2025b](#)) and their reasoning-tuned DeepSeek-R1 versions, and Llama3.3-70B ([Llama Team, 2024](#)) alongside its reasoning-tuned variant.

The DeepSeek-R1 1.5B and 7B parameter variants follow a different training pipeline. Rather than direct reasoning-tuning from base Qwen2.5 models. These variants first undergo intermediate training on approximately 1 trillion tokens of mathematical data, with prompting designed to elicit reasoning behaviors ([Qwen et al., 2025a](#)). This produces three model variants for each size: the original

	Model	Flute	SemEval	MAGPIE	DICE
Base	Qwen2.5-1.5B	0.849	0.458	0.430	0.366
	Qwen2.5-7B	0.921	<b>0.737</b>	0.786	0.710
	Qwen2.5-14B	0.924	0.586	0.823	0.800
	Qwen2.5-32B	0.914	0.612	0.888	<b>0.873</b>
	Llama-70B	0.921	0.658	0.778	0.816
Math	Qwen2.5-Math-1.5B	0.551	0.484	0.495	0.482
	Qwen2.5-Math-7B	0.691	0.404	0.482	0.507
Reasoning	DeepSeek-R1 Qwen-1.5B	0.577	0.533	0.516	0.499
	DeepSeek-R1 Qwen-7B	0.812	0.585	0.626	0.462
	DeepSeek-R1 Qwen-14B	0.929	0.573	0.863	0.863
	DeepSeek-R1 Qwen-32B	<b>0.948</b>	0.641	<b>0.890</b>	<b>0.866</b>
	DeepSeek-R1 Llama-70B	<b>0.947</b>	0.628	0.873	0.857

Table 1: Results of both the base and the reasoning models on the four datasets. Reported is the mean Macro F1 averaged across 5 runs, using different random seeds. The best score(s) per dataset shown in bold.

base model (Qwen2.5), the math-specialized intermediate model (Qwen2.5-Math), and the final reasoning-tuned model (DeepSeek-R1 Qwen). We evaluate all three variants to understand the training progression’s impact on capabilities. While we anticipate that math-specialization may compromise general domain performance, we include these intermediate models in our evaluation to examine whether subsequent reasoning-tuning with broader domain data can recover the lost general capabilities.

All models are run using the vLLM library (Kwon et al., 2023) utilizing Q6\_K\_M quantizations of each model (Frantar et al., 2023). This allows all the models to be run on a single A100 80GB GPU.

As the datasets represent two types of task, we use two different prompts when interfacing with the LLMs. For the FLUTE dataset, we use the following prompt:

"You will be given two sentences, a premise and a hypothesis. Respond with either 'entailment' if the hypothesis follows from the premise, or 'contradiction' otherwise."

For the other datasets, we use:

"Predict whether the MWE given in the sentence being used idiomatically or literally. Respond 'idiomatic' or 'literal' respectively."

The reasoning models generate a CoT before the output, which we split off for analysis. Where the models do not output the required label in an easily parseable format, we use GPT-4o (OpenAI, 2024) to extract the returned label. We evaluate

the models using Macro F1 for all datasets.

### 3. Model Performance

We present the results from our experiments in Table 1. As expected, we see the larger models achieving the highest performance on the idiomacity detection tasks. The best performing model is DeepSeek-R1 Qwen-32B, which slightly outperforms the larger DeepSeek-R1 Llama-70B on all of the datasets. The only model that surpasses comparable larger models is the base version of Qwen2.5-7b, which outperforms the larger non-reasoning models on FLUTE, and, surprisingly, gets the highest score on SemEval-2022, being the only model to score over 0.7.

In relation to the first question, if reasoning improves the ability of models to detect idiomacity, the results observed paint a mostly positive picture. In Table 2 we show the difference between the base and reasoning variants of each model across the datasets, as well as the difference between the base and intermediate math models: Qwen2.5-Math-1b and Qwen2.5-Math-7b.

Math-tuning produces the large drops in performance we expected for both the Qwen2.5-1b and Qwen2.5-7b models. This is slightly less pronounced on the 1.5b model, though this primarily reflects floor effects as performance on SemEval, MAGPIE, and DICE was already near-random levels, leaving little room for further deterioration. Conversely, the reasoning tuning has a large positive effect when moving from the Math variants to the DeepSeek variants, with overall positive improvements of 0.028 and 0.093, an effect that can especially be seen through single task improvements for the 7b parameter model (0.121, 0.181, 0.115).

Larger models (14B, 32B, and 70B) demon-

	Model	FLUTE	SemEval				MAGPIE	DICE	Mean Diff.
			ALL	EN	PT	GL			
Base	1.5b	-0.272	0.075	0.075	0.068	0.082	0.086	0.132	0.006
	7b	-0.109	-0.153	-0.107	-0.138	-0.077	-0.160	-0.247	-0.167
	14b	0.005	-0.014	-0.015	0.051	0.011	0.040	0.064	0.024
	32b	0.034	0.039	0.039	0.081	-0.010	0.002	-0.007	0.017
	70b	0.026	-0.021	-0.066	-0.085	-0.019	0.095	0.042	0.035
Math	1.5b	0.026	0.049	0.050	0.050	0.057	0.021	0.017	0.028
	7b	0.121	0.181	0.212	0.231	0.206	0.115	-0.044	0.093

Table 2: Difference in performance of the reasoning models compared to ‘Base’ and ‘Math’ variants, in absolute difference in Macro F1. A negative value means reasoning hurts performance, whilst a positive value means it improves performance. The “Mean Diff.” column reports the average improvement from using a reasoning model across the four datasets.

Type	Label	Score
<b>Understanding</b>	No definition	0
	Focus on single word	0
	Inaccurate definition	0
	Partial definition	2
	Accurate definition	3
<b>Reasoning</b>	Nonsensical reasoning	0
	OK reasoning	2
	Good reasoning	3
	Meta-reasoning (inaccurate)	2
	Meta-reasoning (accurate)	3

Table 3: The labels used for the manual labelling of the models CoT, with the scores assigned to each label when calculating averages.

strate consistent benefits from CoT reasoning integration, with performance improvements reaching 0.095 for Llama-70B on MAGPIE. However, results vary in magnitude across tasks, and some performance decreases occur on the SemEval dataset—particularly for Portuguese and Galician languages, where the effect of reasoning is less clear.

## 4. Manual Error Analysis

As an initial step in exploring the understanding of the models, we manually inspect samples of the reasoning outputs and evaluate them. We split the evaluation into two facets: whether the model understands the relevant PIE (whether it generates a valid idiomatic definition, either wholly or in parts throughout the CoT), and whether the reasoning is valid (whether the model successfully disambiguates the PIE in context within its CoT). This categorization allows us to separate the effects of the model’s understanding of the expression and context and its reasoning ability.

### 4.1. Labelling Setup

Three of the authors, acting as annotators, each annotate the same 30 responses (15 correct and 15 incorrect) for each reasoning model. We initially chose responses to the MAGPIE dataset as this is the largest dataset, with the largest variance in expressions used. However, whilst labelling we found the dataset to be quite noisy (mislabelling, example sentences which do not actually contain the target expression, etc.) and so we additionally label examples from DICE. Overall, 300 examples were labelled 3 times (once by each annotator) for understanding and reasoning ability on a 5-point scale. Each point on the scale was awarded a score, which is averaged across annotators and examples, to allow us to get a representative score for the model. The full set of labels and scores can be seen in Table 3.

For understanding, labels range from No definition, Focus on one word, and Inaccurate definition (all scored 0 for lacking accurate information) to Partial definition (scored 2) and Accurate definition (scored 3). Reasoning uses a similar scale: Nonsensical reasoning (0), OK reasoning that partially attends to expression and context (2), and Good

	% Judged Valid	
	DICE	MAGPIE
DeepSeek-R1 Qwen-1.5B	0%	0%
DeepSeek-R1 Qwen-7B	0%	6.7%
DeepSeek-R1 Qwen-14B	40%	60%
DeepSeek-R1 Qwen-32B	46.7%	33.3%
DeepSeek-R1 Llama-70B	40%	33.3%

Table 4: For each model, for DICE and MAGPIE, the percentage of sampled “incorrect” predictions (where the model prediction does not match the gold label), where the model prediction was judged to be valid by the annotators, due to a mislabelled or ambiguous example.

reasoning with mostly correct connections (3). We label meta-reasoning, which occurs (rarely) when the model doesn’t correctly identify the expression in the context, instead relying on meta-knowledge typical expression usage. We use 0-2-3 scoring instead of 1-5 to emphasize the gap between inaccurate and accurate responses, with some labels being assigned the same score as they identify different scenarios of similar quality.

## 4.2. Labelling Findings

### 4.2.1. Mislabelling in the dataset

The first thing we observe whilst manually labelling is that, of the examples where the model is judged to be incorrect, there are many labels which are either incorrect or ambiguous. In Table 4, we show the proportion of examples where the gold label and prediction do not match, for which our annotators agreed that the gold label was either incorrect or ambiguous. There is some overlap between the examples, as the idiomatic subset of DICE (which accounts for most of these examples) is taken from MAGPIE. For the larger models, this effect is particularly pronounced, with a large proportion of their “incorrect” predictions actually being judged valid by human annotators. This indicates that accuracy for the larger models may be higher than suggested by the results in Table 1, which further reinforces the high level of understanding shown by the larger models. We emphasise that this does not imply a large proportion of these datasets are incorrectly labelled or ambiguous, instead the accuracy of the large models means that when they are incorrect they have a high chance of picking out the few examples that are ambiguous or incorrect.

### 4.2.2. Error Patterns by Model Size

Figure 1 shows the manual label scores for each model in the case where the model is correct and

incorrect. We observe that mistakes made by different model sizes often fall into different categories of errors. The larger models appear to have an understanding of most of the expressions in the data, and therefore as part of their reasoning produce a definition of the idiom. The smaller models, however, often produce an incorrect definition or struggle to identify or recognize the expression all together. This then leaves a poor basis for the rest of the reasoning, leaving the label as almost random.

The 1.5B parameter model shows a low level of both understanding and reasoning ability, with most of its correct outputs being due to chance. The 7B parameter model has on average higher levels of both reasoning and understanding, which correlates with the results in Table 1. The main factor to whether the model is correct or incorrect appears to be the reasoning performance, as this varies from close to the larger models when correct, but closer to the 1.5B model when incorrect. The larger models understanding of the expressions is generally high when it is both correct and incorrect, with more variation being identified in the reasoning capabilities. However, the overall reasoning scores are still relatively high, which may indicate that they are incorrect on the more nuanced examples, something we noted whilst labelling.

## 5. Chain-of-Thought Length

We investigate the potential relationship between the Chain of Thought (CoT) size and model accuracy. The DeepSeek paper showed that as model training progressed and model performance improved, average CoT length went up (DeepSeek-AI et al., 2025). To explore the correlation between reasoning length and performance in idiomaticity detection, we split the CoTs into two categories for incorrect and correct predictions, and then tokenise them (using spaCy + en\_core\_web\_sm) to give lengths. We show two histograms of CoT length for correct and incorrect predictions in Figure 2. We can see that the CoT length distribution appears similar across both cases. Next, we calculate the point-biserial correlation between CoT size and prediction correctness for each model size category (1.5B, 7B, 14B, 32B and 70B), which can be seen in Table 5. Additionally, we compute the pseudo- $R^2$  for a logistic regression model that predicts accuracy based on CoT length.

The results reveal that, across various datasets, there is no consistent pattern suggesting a significant relationship between CoT size and prediction accuracy. For instance, the most considerable observed correlation (-30.50%) in the EN-FLUTE corpus corresponds to a modest  $R^2$  value of 0.11, meaning only 11% of the variance in prediction outcomes is explained by CoT size. While models of

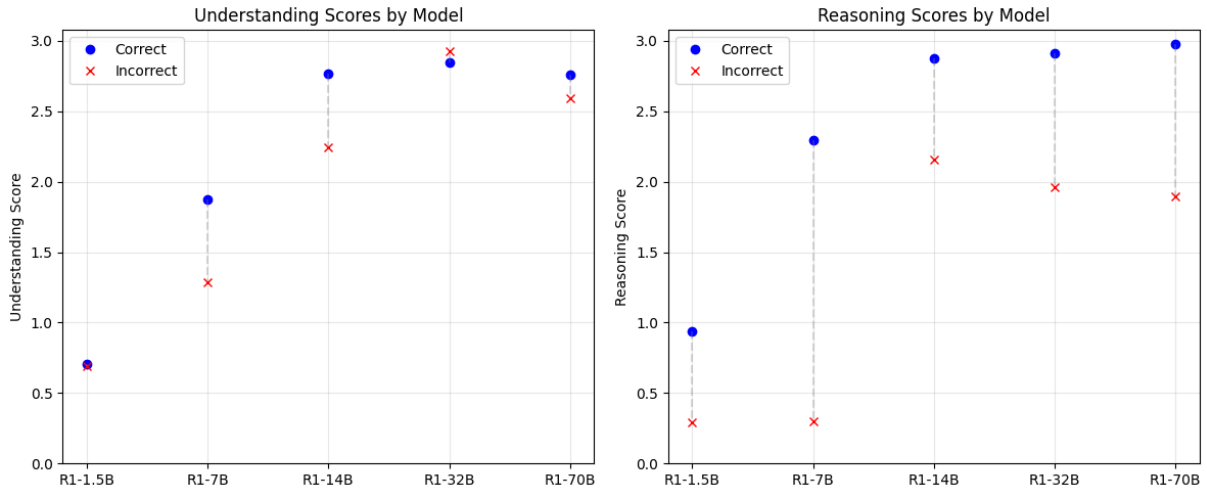


Figure 1: The average manually labelled reasoning and understanding scores (between 0 and 3) for each model, for 15 correct and 15 incorrect examples, averaged across DICE and MAGPIE.

Model	Flute	SemEval			MAGPIE	DICE
		EN	PT	GL		
DeepSeek 1.5B	-0.038	0.009	-0.037	-0.061	0.000	-0.011
DeepSeek 7B	<b>-0.151</b>	<b>-0.075</b>	-0.001	<b>-0.089</b>	<b>-0.057</b>	-0.006
DeepSeek 14B	-0.022	0.020	-0.036	0.000	<b>-0.135</b>	<b>-0.177</b>
DeepSeek 32B	<b>-0.246</b>	-0.003	-0.064	<b>0.126</b>	<b>-0.119</b>	<b>-0.147</b>
DeepSeek 70B	<b>-0.305</b>	-0.036	-0.002	0.028	<b>-0.132</b>	<b>-0.156</b>

Table 5: Correlation between CoT length and model correctness. Values in bold indicate a p-value  $< 0.05$ .

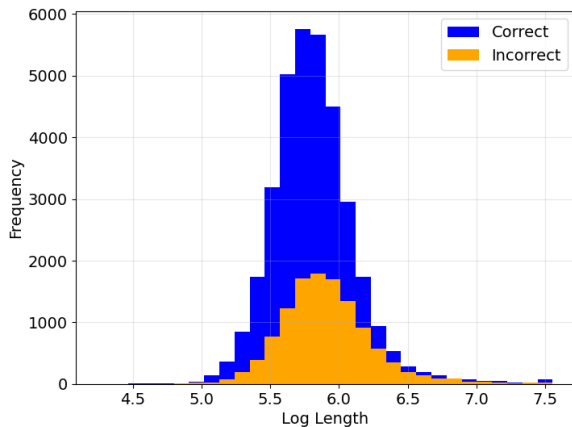


Figure 2: The distribution of reasoning output log lengths for correct (blue) vs. incorrect (orange) model predictions, aggregated across all models and datasets.

larger sizes (32B and 70B) show slight improvements in  $R^2$  for certain corpora, such as FLUTE and DICE, the overall findings suggest no clear trend linking CoT length to predictive success.

## 6. Definition Generation as Distillation

The results from our manual labelling reveal that the 32B and 70B parameter models can accurately produce definitions of the expressions represented in the dataset. Given the lack of comprehensive resources for idiomatic expression definitions, this finding presents an opportunity to use model-generated definitions as a novel distillation technique.

As an initial exploration, we investigate whether providing the model generated definitions from the larger models to the smaller reasoning models can improve performance. In our manual analysis, we found that the 1.5B and 7B parameter reasoning models were not able to accurately reproduce definitions, which may suggest that they do not contain the required knowledge to do so, partly explaining their poor performance on the detection tasks. By providing this information in the prompt, this knowledge is able to be used by the models whilst making predictions.

For this experiment, we use DeepSeek-R1 32B to generate definitions for all the expressions in

PIE	Generated Definition
all hell broke loose	“All hell broke loose” means that a situation suddenly became extremely chaotic, noisy, or violent.
off the hook	To be off the hook means to be free from a responsibility, obligation, or troublesome situation, often after being expected to handle it.
against the grain	To go against the grain is to act in a way that is contrary to the usual approach, often resulting in resistance.
play with fire	‘Play with fire’ means to engage in a risky or dangerous situation that could lead to negative consequences.
make a killing	To make a killing means to achieve great success or earn a lot of money.

Table 6: Examples of the definitions generated by DeepSeek-R1 32B for given PIEs from the DICE dataset.

Model	Original Prompt		Definition Prompt	
	FLUTE	DICE	FLUTE	DICE
DeepSeek-R1 Qwen-1.5B	0.577	0.499	<b>0.663</b>	0.487
DeepSeek-R1 Qwen-7B	0.812	0.462	<b>0.864</b>	0.461
DeepSeek-R1 Qwen-14B	0.929	0.863	0.908	0.842

Table 7: The results using both the original and definition appended prompt, when evaluating the smaller models on FLUTE and DICE. We report the mean Macro F1 across 5 runs with different random seeds, and bold any significantly improved results for each model/dataset pair.

the FLUTE and DICE datasets (examples shown in Table 6). We select DeepSeek-R1 32B as this shows the highest overall ‘understanding’ score in the manual analysis (see Figure 1). We then append the definition to the prompt used in our main experiments and rerun the evaluation for the 1.5B, 7B, and 14B parameter reasoning models. These models represent a range of ‘understanding’ scores, allowing us to analyse the effectiveness of the definitions across different performance and ‘understanding’ levels. We expect that the models with the lowest performance on the detection tasks and ‘understanding’, e.g. the 1.5B and 7B parameter models, will benefit the most from the additional information in the prompt.

## 6.1. Results

The results from the experiment (shown in Table 7) show that adding definitions to the prompt significantly increases performance for the 1.5B and 7B parameter models on the FLUTE dataset. However, for DICE, and for the 14B parameter model on both datasets, no significant difference can be seen between adding the definition to the prompt or not.

We suggest that the difference in effectiveness on FLUTE and DICE comes from the fact that, as suggested by (Mi et al., 2025), DICE requires models to attend to the surrounding context to make decisions. Adding the definition will not improve the models’ ability to do this, and so no increase in

performance is seen.

However, the lack of any decrease in performance is promising for the technique as it means that adding definitions to the prompt can safely be done without risk of reducing performance. While the 14B parameter model shows slightly larger absolute decreases on both datasets, these differences are not statistically significant.

## 7. Conclusion

In this study, we have investigated the effect that reasoning has on the idiomaticity detection ability of LLMs. On four idiomaticity datasets (Flute, SemEval, MAGPIE, DICE), we have evaluated the performance of the distilled DeepSeek-R1 models across a range of sizes, in addition to their corresponding non-reasoning base models, and intermediate Math variants for the small parameter count models. To further explore the results, we have performed both a manual and a quantitative analysis on the reasoning outputs of the models in order to gain some insight into the models’ understanding of idiomaticity, as well as their ability to reason with the expression in context. As a result, we have three main findings:

1. The effect of reasoning on idiomaticity detection is relatively small, although it varies across model size and per dataset. For the smaller models, training on math specific reasoning data greatly reduces performance, but subse-

quent training on more general domain reasoning data can restore performance. Larger models show slight increases in performance when trained on general domain reasoning data.

2. Our manual analysis shows the larger models are consistently able to produce high quality definitions of the expressions, something which the smaller models struggle to do. However, for all models the quality of the reasoning is the most prominent factor which decides whether a model will be correct or not. This highlights that even the largest models struggle to use the expression and the context to disambiguate the meaning.
3. Our further experiments utilising the definitions produced by the larger models as knowledge sources for smaller models show promise as a knowledge distillation technique. Appending the definitions to the prompt improves the performance in some cases, whilst having no significant effect in others, implying this technique can be applied generally without risk of regression.

Our results show that while CoT reasoning offers small benefits for idiomaticity detection, the modest improvements suggest that current reasoning approaches may not fully address the underlying contextual disambiguation challenges inherent in idiomaticity detection. Even if the reasoning generated appears to indicate idiomatic understanding, further examination indicates otherwise. Future work should explore alternative reasoning frameworks specifically designed for figurative language processing, perhaps incorporating more structured approaches to context interpretation. Additionally, it could explore other ways to utilise the high-quality definitions generated by the models, either through in-context learning or as fine-tuning data.

## Acknowledgments

This work was partly supported by UKRI (grants MR/U506734/1, EP/T02450X/1, and EP/S023062/1), CNPq (406926/2025-5) and EQUATE. It also received support from the CA21167 COST action UniDive, funded by COST (European Cooperation in Science and Technology).

## 8. Limitations

Whilst we do on average see a small improvement in performance for idiomaticity detection in the larger models, this result isn't wholly consistent, as the reasoning variant of Llama-3.3 70B performs

worse by 0.037 F1 on the English subset of SemEval compared to the non-reasoning variant. The average performance increasing from reasoning is also lower for this model than Qwen2.5 14B and Qwen2.5 32B (0.013 compared to 0.039). Further work could investigate a reason for this inconsistency.

On our human analysis of the model CoTs, we break the evaluation into "understanding", which we take as the ability of a model to define the relevant PIE in its idiomatic sense, and "reasoning", which we take as the ability of a model to use the context and the PIE to determine whether it's literal or idiomatic. This provides a useful analysis, but a more fine-grained evaluation using other features would perhaps be more useful, for example looking at whether a model is logically consistent, whether it's hallucinating, etc.

## 9. Bibliographical References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altmenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- 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, Wanxia 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](#). ArXiv:2501.12948 [cs].
- Elias Frantar, Saleh Ashkboos, Torsten Hoefler, and Dan Alistarh. 2023. [GPTQ: Accurate Post-Training Quantization for Generative Pre-trained Transformers](#). ICLR 2023. ArXiv:2210.17323 [cs].
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. 2023. [Efficient memory management for large language model serving with pagedattention](#). In *Proceedings of the 29th Symposium on Operating Systems Principles, SOSP '23*, page 611–626, New York, NY, USA. Association for Computing Machinery.
- Llama Team. 2024. [The Llama 3 Herd of Models](#). ArXiv:2407.21783 [cs].
- Qing Lyu, Shreya Havaldar, Adam Stein, Li Zhang, Delip Rao, Eric Wong, Marianna Apidianaki, and Chris Callison-Burch. 2023. [Faithful chain-of-thought reasoning](#). In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 305–329, Nusa Dua, Bali. Association for Computational Linguistics.
- OpenAI. 2024. [Gpt-4o system card](#).
- OpenAI. 2025. [Learning to reason with LLMs](#).
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, 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 *Proceedings of the 36th International Conference on Neural Information Processing Systems, NIPS '22*, Red Hook, NY, USA. Curran Associates Inc.
- Dylan Phelps, Thomas M. R. Pickard, Maggie Mi, Edward Gow-Smith, and Aline Villavicencio. 2024. [Sign of the Times: Evaluating the use of Large Language Models for Idiomaticity Detection](#). In *Proceedings of the Joint Workshop on Multiword Expressions and Universal Dependencies (MWE-UD) @ LREC-COLING 2024*, pages 178–187, Torino, Italia. ELRA and ICCL.
- Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. 2025a. [Qwen2.5 technical report](#).
- Qwen. 2024. [Qwen2.5: A Party of Foundation Models!](#) Section: blog.
- Qwen, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. 2025b. [Qwen2.5 Technical Report](#). ArXiv:2412.15115 [cs].

Răzvan-Alexandru Smădu, David-Gabriel Ion, Dumitru-Clementin Cercel, Florin Pop, and Mihaela-Claudia Cercel. 2024. [Investigating Large Language Models for Complex Word Identification in Multilingual and Multidomain Setups](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 16764–16800, Miami, Florida, USA. Association for Computational Linguistics.

Miles Turpin, Julian Michael, Ethan Perez, and Samuel Bowman. 2023. [Language models don't always say what they think: Unfaithful explanations in chain-of-thought prompting](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 74952–74965. Curran Associates, Inc.

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 *Proceedings of the 36th International Conference on Neural Information Processing Systems, NIPS '22*, Red Hook, NY, USA. Curran Associates Inc.

United States. Association for Computational Linguistics.

## 10. Language Resource References

Tuhin Chakrabarty, Arkadiy Saakyan, Debanjan Ghosh, and Smaranda Muresan. 2022. [FLUTE: Figurative Language Understanding through Textual Explanations](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 7139–7159.

Hessel Haagsma, Johan Bos, and Malvina Nissim. 2020. [MAGPIE: A large corpus of potentially idiomatic expressions](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 279–287. European Language Resources Association.

Maggie Mi, Aline Villavicencio, and Nafise Sadat Moosavi. 2025. [Rolling the DICE on idiomaticity: How LLMs fail to grasp context](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7314–7332, Vienna, Austria. Association for Computational Linguistics.

Harish Tayyar Madabushi, Edward Gow-Smith, Marcos Garcia, Carolina Scarton, Marco Idiart, and Aline Villavicencio. 2022. [SemEval-2022 task 2: Multilingual idiomaticity detection and sentence embedding](#). In *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*, pages 107–121, Seattle,