

OpenHuEval: Evaluating Large Language Model on Hungarian Specifics

Haote Yang^{1*}, Xingjian Wei^{1*}, Jiang Wu^{1*†}, Noémi Ligeti-Nagy^{2*},
Jiaxing Sun^{3*}, Yinfan Wang^{4*}, Zijian Győző Yang^{2*}, Junyuan Gao^{5*}, Jingchao Wang⁶,
Bowen Jiang⁷, Shasha Wang¹, Nanjun Yu⁶, Zihao Zhang⁸, Shixin Hong⁹, Hongwei Liu¹,
Wei Li¹, Songyang Zhang¹, Dahua Lin^{1,10}, Lijun Wu¹, Gábor Prószték², Conghui He^{1‡}

¹Shanghai Artificial Intelligence Laboratory,

²HUN-REN Hungarian Research Centre for Linguistics, ³Wuhan University,

⁴Shanghai University, ⁵University of Chinese Academy of Sciences,

⁶East China Normal University, ⁷Peking University, ⁸Harbin Institute of Technology,

⁹Tsinghua University, ¹⁰Chinese University of Hong Kong

Abstract

We introduce OpenHuEval, the first benchmark for LLMs focusing on the Hungarian language and specifics. OpenHuEval is constructed from a vast collection of Hungarian-specific materials sourced from multiple origins. In the construction, we incorporated the latest design principles for evaluating LLMs, such as using real user queries from the internet, emphasizing the assessment of LLMs’ generative capabilities, and employing LLM-as-judge to enhance the multidimensionality and accuracy of evaluations. Ultimately, OpenHuEval encompasses eight Hungarian-specific dimensions, featuring five tasks and 3953 questions. Consequently, OpenHuEval provides the comprehensive, in-depth, and scientifically accurate assessment of LLM performance in the context of the Hungarian language and its specifics. We evaluated current mainstream LLMs, including both traditional LLMs and recently developed Large Reasoning Models. The results demonstrate the significant necessity for evaluation and model optimization tailored to the Hungarian language and specifics. We also established the framework for analyzing the thinking processes of LLMs with OpenHuEval, revealing intrinsic patterns and mechanisms of these models in non-English languages, with Hungarian serving as a representative example. We will release OpenHuEval at <https://github.com/opendatalab/OpenHuEval>.

1 Introduction

Recent advancements in Large Language Models (LLMs) (Jaech et al., 2024; Team, 2024a) represent significant strides toward artificial general intelligence (AGI). However, notable performance gaps remain between English and other languages in both language-agnostic tasks (e.g., math reasoning, code generation) (Huang et al., 2025; Zhang

et al., 2024) and language-specific tasks (e.g., idiom usage, cultural understanding) (Naous et al., 2023; Sun et al., 2024; Liu et al., 2024c), posing challenges to global AI deployment and equitable development. The disparities in cross-lingual performance arise mainly from two factors: First, the training data, particularly the pre-training data, is heavily skewed toward English. Second, while English evaluation benchmarks are advanced and rapidly evolving, non-English benchmarks are underdeveloped, particularly for language-specific features, limiting the identification of shortcomings in non-English languages and leading to their neglect in research.

This paper focuses on the evaluation of Hungarian language and specifics. Hungarian is spoken by around 14 million people worldwide. Research on the Hungarian language not only improves the user experience for Hungarian speakers but also offers valuable insights for similar studies in other languages and regions. Existing Hungarian evaluation datasets are largely translations of English ones, missing essential Hungary-specific elements such as language nuances, culture, history, and regional context, which are key for Hungarian users. Among the existing evaluation datasets, HuLU (Ligeti-Nagy et al., 2024) is the key benchmark for Hungarian language understanding, but its focus on multiple-choice and true/false questions limits its ability to evaluate broader LLM capabilities, such as language generation, open-domain Q&A, reasoning and instruction-following.

To address this gap, we introduce OpenHuEval, the first evaluation benchmark for LLMs focused on Hungarian language and specifics. The comparison of OpenHuEval with the existing related benchmarks is shown in Table 1. Overall, OpenHuEval has two main distinguishing features:

1) **Hungarian-Specific** Inspired by (Liu et al., 2024b; Sun et al., 2024), we propose eight distinct Hungarian-specific dimensions (see §2.1), cover-

*Equal contribution.

†Project lead.

‡Corresponding author. heconghui@pjlab.org.cn

Benchmark	Real user query	Self-awareness evaluation	Proverb Reasoning	Generative task & llm-as-judge	Hungarian Lang	Comprehensive Hu-specific
WildBench(Lin et al., 2024)	✓	✗	✗	✓	✗	✗
SimpleQA(Wei et al., 2024), ChineseSimpleQA(He et al., 2024)	✗	✓	✗	✓	✗	✗
MAPS(Liu et al., 2024c)	✗	✗	✓	✗	✗	✗
MARC, MMLU et al in (Lai et al., 2023)	✗	✗	✗	✗	✓	✗
BenchMAX(Huang et al., 2025)	✗	✗	✗	✓	✓	✗
MILQA(Novák et al., 2023)	✗	✗	✗	✗	✓	✗
HuLU(Ligeti-Nagy et al., 2024)	✗	✗	✗	✗	✓	✗
OpenHuEval (ours)	✓	✓	✓	✓	✓	✓

Table 1: Comparison of related benchmarks.

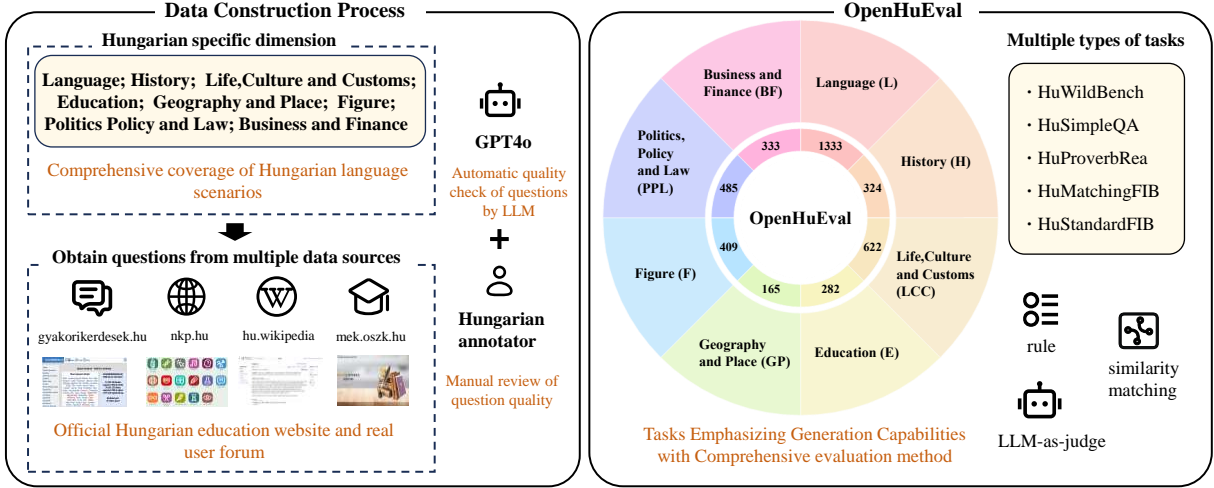


Figure 1: Overview of OpenHuEval.

ing a variety of scenarios that users may encounter when querying in Hungarian. Guided by these dimensions, we collected a vast amount of Hungarian specific material from multiple sources and used this to construct the corresponding evaluation tasks.

2) **Keeping up with the latest advances in LLM evaluation** Significant progress has been made in LLM evaluation, with query sources shifting from manual or rule-based constructions to real-world internet questions (Lin et al., 2024), enhancing practical relevance. Question formats evolved from multiple-choice to open-ended Q&A (Wei et al., 2024), better reflecting actual usage. Evaluation methods transitioned from rule-based approaches to LLM-as-judge and subjective assessments, improving accuracy and objectivity (Li et al., 2024a). However, these advancements primarily apply to English datasets and not Hungarian. Thus, when creating OpenHuEval, we incorporated these principles and methodologies from English evaluations.

Based on OpenHuEval, we evaluated the performance of mainstream LLMs on Hungarian language and specifics. We compared the performance differences of these models on the typical datasets

of OpenHuEval with corresponding datasets in other languages. The results indicate a significant necessity for evaluation and model optimization specifically for Hungarian language and specifics.

Moreover, Large Reasoning Models (LRMs), like o1, mark a new direction in LLM development. Through extensive reasoning, self-reflective negation, and exploring multiple reasoning paths, they greatly improve reasoning abilities, adhering to the test-time scale law—a crucial step toward AGI. Recent studies (Wang et al., 2025) have analyzed these reasoning processes, offering insights for optimization, but have largely focused on English-language contexts, neglecting Hungarian language and specific scenarios. Building on OpenHuEval, we developed the framework for dissecting the reasoning processes of LRMs. Using Hungarian as the example, we uncovered intrinsic patterns of the representative LRMs in non-English languages. These findings provide valuable insights for the research community to further advance the development of LRMs.

In summary, the contributions of this paper include the following three points:

- We developed OpenHuEval, the first bench-

mark for LLMs focusing on the Hungarian language and specifics. OpenHuEval incorporates the latest design principles for evaluating LLMs, such as using real user queries from the internet, emphasizing the assessment of LLMs’ generative capabilities, and employing LLM-as-judge to enhance the multidimensionality and accuracy of evaluations.

- We conducted the comprehensive evaluation of current mainstream LLMs, including traditional LLMs and recently developed LRMs. The results highlight the significant necessity for evaluation and model optimization tailored to Hungarian language and specifics.

- We established the framework for analyzing the thinking processes of the cutting-edge LRMs, revealing the intrinsic patterns and mechanisms of these models in the Hungarian language and specifics, and providing a reference for research related to other non-English languages.

2 OpenHuEval

OpenHuEval is the benchmark specifically designed to evaluate the performance of LLM in handling Hungarian language and specifics. The overview of OpenHuEval is in Figure 1. Examples are shown in Figure 2.

2.1 Hungarian-specific dimensions and OpenHuEval tasks

Inspired by (Liu et al., 2024b; Sun et al., 2024), OpenHuEval encompasses eight Hungarian-Specific Dimensions (HuSpecificDim), as shown in Table 2: *Language (L)*, *History (H)*, *Life, Culture, and Customs (LCC)*, *Education and Profession (EP)*, *Geography and Place (GP)*, *Figure (F)*, *Politics, Policy, and Law (PPL)*, and *Business and Finance (BF)*. These dimensions comprehensively cover a wide range of scenarios encountered by users when utilizing Hungarian as the query language. As a result, they enable the systematic evaluation of the performance of LLMs in tasks related to the Hungarian.

Bearing the above HuSpecificDim in mind, the first step in building OpenHuEval involves gathering corpora rich in Hungarian specifics. Inspired by (Liu et al., 2024c; Li et al., 2023; He et al., 2024; Arora et al., 2024), we collected data from sources such as Hungarian proverbs, exam questions, forums, and Wikipedia. Through processes including filtering, refinement, construction, and quality assurance, we developed a total of five eval-

uation tasks comprising 3953 questions in total, as detailed in Table 3. The subsequent sections of this chapter will introduce these tasks and their corresponding datasets in detail.

2.2 Hungarian WildBench

Task Introduction: The Hungarian WildBench (HuWildBench) task aims to evaluate the performance of LLMs in answering various questions arising from the everyday lives of Hungarians. All questions are sourced from Hungary’s well-known forum website¹ (hereinafter referred to as “g13k” for brevity). These questions cover a wide range of topics, including cultural customs, education, tourism, legal regulations, and business and finance, thus reflect **real-life issues encountered by Hungarians**. Examples of HuWildBench questions are shown in Figure 2 and Table 21. The queries in HuWildBench are user-generated content, therefore their linguistic expressions and question formats tend to be less formal than the structured and polished written language. This poses the realistic challenge for LLMs, as they must adapt to such informal and spontaneous language style. The construction of HuWildBench is detailed in Appendix D.

Metric and judge: We use the WB-Score (Lin et al., 2024) as evaluation metric for HuWildBench, which is calculated in the following manner: We have simultaneously developed the customized checklist for each question during the process of creating each question, to assist the LLM judge in evaluating the responses. Examples of these checklists can be seen in Table 21. Following (Lin et al., 2024), GPT-4o is used as the judge model, which evaluates the quality of each response based on the checklist and provides detailed strengths and weaknesses before assigning a score from 1 to 10. The definition of scores is shown in Table 8 and the judge prompt is detailed in Appendix D.5. Different from (Lin et al., 2024), our final scores are calculated as the average of all test sample scores, with each score multiplied by 10.

2.3 Hungarian SimpleQA

Task Introduction: Hungarian SimpleQA (HuSimpleQA) is designed to assess the ability of LLMs to answer short, fact-seeking questions related to Hungary. Inspired by (Wei et al., 2024)

¹<https://www.gyakorikerdesek.hu/>, which is similar to <https://www.quora.com/> for English-speaking world.

<p>HuWildBench</p> <p>a kérdés az: Mi lesz a jövőben a szocializmus alatt megépül sok panellet? a leírás: Úgy tudom, hogy kb 60 éves életciklusra tervezték őket. Magyarországon (és a környező országokban is) rengeteg ember él bennük. Mi fog történni akkor, ha lakhatatlanná kezdenek válni? Mi lesz azzal a sok emberrel? Meg a panelokkal?</p> <p>The question is: What will happen in the future to the many panels built under socialism? Description: I understand they are designed for a life cycle of about 60 years. There are a lot of people living in them in Hungary (and surrounding countries). What will happen if they start to become uninhabitable? What will happen to all those people? And the panels?</p>	<p>HuMatchingFIB</p> <p>Questions: "Válaszd ki a legördülő listából, hogy melyik fogalom illik a hiányos mondatokba! A faj azon egyedeit, melyek tényleges szaporodási közösséget alkotnak, #0#. A faj #1# mindazoknak a hatásainak az összessége, melyek ténylegesen hatnak az élőlényekre. A populáció méretét jellemző egyik legfontosabb sajátosság a/az #2#. Terület- vagy térfogategységre vonatkoztatott egyedszám a/az #3#. A környezeti tényező azon tartománya, amelyen belül az élőlények életműködéseket mutatnak a #4#. Jellemzően az a környezeti tényező határozza meg a populáció elterjedését, amelyre nézve az adott faj szűk tűrési, ezt nevezzük úgy, hogy #5#." Options: "A.környezet", "B.tűrőképesség", "C.egyedsűrűség", "D.egyedszám", "E.korlátozó tényező", "F.populációnak"</p> <p>Questions: "Select from the dropdown list which concept fits into the incomplete sentences! The individuals of a species that form an actual reproductive community are called #0#. The #1# is the totality of all effects that actually influence living organisms. One of the most important characteristics describing the size of a population is the #2#. The number of individuals per unit area or volume is the #3#. The range of an environmental factor within which living organisms exhibit life processes is the #4#. Typically, the environmental factor that determines the distribution of a population is the one for which the species has a narrow tolerance, and this is called the #5#." Options: "A.environment", "B.tolerance", "C.population density", "D.population size", "E.limiting factor", "F.population"</p>
<p>HuProverbRea-OE</p> <p>Hungarian Phrase: 'Ajándék lónak ne nézd a fogát.' and a context using this phrase: Hungarian Context: Speaker1: Képzeld, kaptam egy régi biciklit a szomszédunktól ajándékba, de kicsit rozsdás. Speaker2: Ne aggódj emiatt! Ajándék lónak ne nézd a fogát. What does the person mean by using this phrase?</p> <p>Hungarian Phrase: 'Don't look at a gift horse's teeth.' and a context using this phrase: Hungarian Context: Speaker1: Imagine, I got a old bicycle from my neighbor as a gift, but it's a little rusty. Speaker2: Don't worry about it! Don't look at a gift horse's teeth. What does the person mean by using this phrase?</p>	<p>HuStandardFIB</p> <p>Questions: "Táld ki a leírások alapján, hogy kiről vagy miről van szó! Írd be a meghatározások után a megfelelő kifejezéseket!" "A. Az #0# udvarában nevelkedett ifjúkorában III. Béla:", "B. A királyi adminisztráció céljából létrehozott intézményrendszer: #1#.", "C. Feltehetően ő volt III. Béla jegyzője: #2#.", "D. 1181-ben tette általánossá III. Béla a hivatali ügyintézésben: #3#.", "E. #4# fogalmazta meg a hivatalos iratokat, okleveleket:", "F. Ebben a városban temették el III. Bét: #5#"</p> <p>Questions: "Based on the descriptions, guess who or what is being referred to! Enter the appropriate terms after the definitions!" "A. In the court of #0#, Béla III spent his youth:", "B. The institutional system created for royal administration: #1#", "C. He was likely the scribe of Béla III: #2#", "D. In 1181, Béla III made this mandatory in official proceedings: #3#", "E. #4# was responsible for drafting official documents and charters:", "F. The city where Béla III was buried: #5#"</p>
<p>HuSimpleQA</p> <p>Question 1: "Melyik évben alakult a Nyíregyháza Spartacus FC?" Answer: "1928" Question 1: "In which year was Nyíregyháza Spartacus FC founded?" Answer: "1928" Question 2: "Melyik magyar film nyerte el a FIPRESCI-díjat az 1983-as Cannes-i Nemzetközi Filmfesztiválon?" Answer: "Szerencsés Dániel" Question 2: "Which Hungarian film won the FIPRESCI Prize at the 1983 Cannes International Film Festival?" Answer: "Lucky Daniel"</p>	

Figure 2: Examples of OpenHuEval. The original text is in black, while the translation into English is in blue. In HuMatchingFIB and HuStandardFIB, the blank is highlighted.

HuSpecificDim	Definition	#Question
Language(L)	Basic knowledge of the Hungarian language and Hungarian proverbs and sayings	1333
History(H)	Historical events and historical development of Hungary	324
Life, Culture, and Custom(LCC)	Religion, rituals, culture, holidays, and the daily life of Hungarians	622
Education and Profession(EP)	Education system in Hungary and related professions	282
Geography and Place(GP)	Geographical knowledge of Hungary, cities, and locations	165
Figure(F)	Famous figures of Hungary	409
Politics, Policy and Law(PPC)	Politics, policies, and laws of Hungary	485
Business and Finance(BF)	Business and finance in Hungary	333

Table 2: The Hungarian-Specific Dimensions (HuSpecificDim).

and (He et al., 2024), we constructed HuSimpleQA based on Hungarian Wikipedia², with the following key characteristics. **Hungarian:** The questions in HuSimpleQA are in Hungarian, and they focus on facts specifically related to Hungary. **Diverse:** The questions in HuSimpleQA cover the eight Hungary-specific dimensions proposed in §2.1. **High-quality:** The construction process of HuSimpleQA (in Appendix E) includes comprehensive and strict quality control procedures, ensuring the quality and accuracy of the questions. **Static:** Similar to SimpleQA, the answers to the questions in HuSimpleQA do not change over time, ensuring that the dataset remains evergreen. **Easy-to-evaluate:** The questions and answers in HuSimpleQA are short and concise, making them ideal for evaluation through LLMs. The examples of HuSimpleQA are shown in Figure 2 and Table 22. The construction of HuSimpleQA is detailed in Appendix E.

Metric and Judge: Following (Wei et al., 2024), we use GPT-4o as a judge to categorize the responses of the LLM to HuSimpleQA into three

classes: CORRECT, INCORRECT or NOT_ATTEMPTED. The judge prompt can be found in Appendix E.7, Figure 27.

2.4 Hungarian Proverb Reasoning

Task Introduction: Hungarian Proverb Reasoning (HuProverbRea), which consists of the collection of Hungarian proverbs, idioms, abbreviations, is a task that requires the LLM to **understand and reason the meaning of Hungarian proverbs in a specific context**. As shown by the examples in Figure 2, LLM is provided with a context in which a Hungarian proverb is used, accompanied by a question: "What does the speaker mean by the saying?". Then, the LLM is tasked with discerning the speaker's true intention, either by selecting the correct interpretation from two provided options (2CQ setting), or by directly articulating the speaker's intended meaning (OE setting). The construction of HuProverbRea is detailed in Appendix F.

Metric and judge: For the 2CQ setting, we simply measure the correct ratio of candidate LLMs. For the OE setting, we adopt GPT-4o as judge to decide if the answer is acceptable. We provide the original proverb, its context and the English expla-

²<https://hu.wikipedia.org/>

Task	HuSpecificDim	Judge	Question type	#Question
HuWildBench	$\mathcal{LCC}, \mathcal{EP}, \mathcal{PPL}, \mathcal{BF}$	llm, checklist	OE	1154
HuSimpleQA	$\mathcal{L}, \mathcal{H}, \mathcal{LCC}, \mathcal{EP}, \mathcal{GP}, \mathcal{F}, \mathcal{PPL}, \mathcal{B}$	llm	OE	1293
HuProverbRea	\mathcal{L}	rule, llm	2CQ/OE	1135
HuMatchingFIB	\mathcal{L}, \mathcal{H}	rule	Matching Filling-in-Blank	278
HuStandardFIB	\mathcal{L}, \mathcal{H}	rule, similarity matching	Standard Filling-in-Blank	93

Table 3: Tasks of OpenHuEval

nation of the proverb as references when judging OE responses. Detailed prompt templates are listed in Appendix F.

2.5 Hungarian Matching and Standard Filling-in-Blank

Task Introduction: Hungarian Matching Fill-in-the-Blank (HuMatchingFIB) is a task where key terms in a text are removed, and a pool of candidate words or phrases is provided. This pool includes both correct answers and distractors. The task requires the LLM to choose the most suitable words from the pool to fill in the blanks, thus restoring the full meaning of the text. The example is shown in Figure 2 and Figure 33. HuMatchingFIB effectively tests the LLM’s abilities in understanding information, reasoning within context, and differentiating correct answers from distractors.

In contrast, Hungarian Standard Fill-in-the-Blank (HuStandardFIB) also involves filling in blanks but does not offer a candidate pool. Instead, the model must rely on its internal knowledge and the provided context to complete the text. The examples are shown in Figure 34. Consequently, HuStandardFIB evaluates the LLM’s overall ability to recall knowledge and reason within context.

The constructions of HuMatchingFIB and HuStandardFIB are detailed in Appendix G.

Metric and Judge: In our inference prompts, we explicitly instruct the LLM to generate responses in a specified format. Additionally, we have established a set of well-defined rules to evaluate the correctness of the LLM’s answers for each blank. The detailed format requirements and judgment criteria can be found in the Appendix G.

For both HuMatchingFIB and HuStandardFIB, we evaluate performance at two levels: blank-level and question-level accuracy. Specifically, **Acc_b** (Blank-level Accuracy) measures the proportion of blanks that the model answers correctly across all questions. On the other hand, **Acc_q** (Question-level Accuracy) evaluates the proportion of questions that the model answers entirely correctly. A question is only considered correct if **all** its associated blanks are answered accurately.

3 Experiments and Analysis

3.1 Experimental setup

We evaluated the currently mainstream LLMs, including GPT-4o (Hurst et al., 2024), GPT-4o mini³, Deepseek-V3 (Liu et al., 2024a), Qwen2.5-Instruct (Yang et al., 2024), and Llama-3.1-Instruct (Dubey et al., 2024), as well as the latest Large Reasoning Models (LRMs) such as OpenAI o1-mini (Jaech et al., 2024), QwQ-32B-Preview (Team, 2024b) (abbreviated as QwQ in following text), and Deepseek-R1 (Team, 2024a). Detailed specifications of these models are provided in Table 4.

We used OpenCompass⁴ in all our experiments. For traditional instruction-based LLMs, we adopted OpenCompass’s default settings for the maximum output length. For Large Reasoning Models, we set the output length to 8192 to ensure sufficient space for reasoning process and to produce a complete final answer, avoiding premature output truncation. For OpenAI models (GPT series and o1-mini), we used their official API with settings following OpenCompass’s default configuration. For Deepseek-V3 and Deepseek-R1, due to the high usage volume of Deepseek’s official API causing instability, we used equivalent API services provided by Alibaba Cloud⁵ and Silicon Valley Flow⁶. The settings followed OpenCompass’s configurations, with the temperature set to 0.7. For other models in Table 4, we performed inference locally with NVIDIA A100 GPUs, using LMDeploy⁷ as the inference backend. The settings followed OpenCompass’s default configuration (Temperature = 1e-6, top_k = 1).

3.2 Overall performance

The overall performance of all LLMs on OpenHuEval is presented in Table 5. It can be observed that across a total of five tasks, Deepseek-R1 ranks first

³We used the gpt-4o-2024-11-20 version for GPT-4o and the gpt-4o-mini-2024-07-18 version for GPT-4o-mini.

⁴<https://github.com/open-compass/opencompass>

⁵<https://cn.aliyun.com/>

⁶<https://siliconflow.cn/>

⁷<https://github.com/InternLM/lmdeploy>

Model	Size	Reasoning Model	Open-source	Inference Method
GPT-4o	-	N	N	Official API
GPT-4o-mini	-	N	N	Official API
Deepseek-V3	-	N	Y	Alibaba Cloud and SiliconFlow API
Qwen2.5-Instruct	7B,72B	N	Y	Local GPU
Llama-3.1-Instruct	8B,70B	N	Y	Local GPU
o1-mini	-	Y	N	Official API
QwQ	32B	Y	Y	Local GPU
Deepseek-R1	-	Y	Y	Alibaba Cloud and SiliconFlow API

Table 4: LLMs evaluated in our experiments.

in three tasks and achieves top-tier performance in the other two tasks. GPT-4o ranks first in two tasks and second in the remaining three tasks. These results demonstrate the exceptional performance of the two models in Hungary-specific tasks.

Open-source models vs Closed-source models:

Among open-source models, Deepseek-R1 stands out, while Deepseek-V3 also demonstrates strong overall performance, ranking highly across all tasks. Llama-3.1-Instruct-70B achieved impressive scores of 93.83% in the HuProverbRea-2CQ task and 36% in the HuSimpleQA task, ranking second only to the closed-source model GPT-4o. This highlights the growing potential of open-source models, led by Deepseek-R1, which are increasingly showing capabilities comparable to closed-source models in Hungarian language tasks.

Traditional LLMs vs. Large Reasoning Models: We compared Traditional LLMs and LRMs within the same series. Across five tasks, Deepseek-R1 consistently outperforms Deepseek-V3 in four of them. Specifically, in the HuMatchingFIB task, Deepseek-R1 achieves relative improvements of 12% at the blank level and 7.19% at the question level compared to Deepseek-V3. Similarly, for the HuStandardFIB task, it achieves gains of 10.32% (blank level) and 7.52% (question level). Although Deepseek-R1 performs slightly worse than Deepseek-V3 on the HuProverbRea task, the performance gap is less than 1%. Considering that both Deepseek-R1 and Deepseek-V3 are based on the same pretrained model, the significantly stronger performance of Deepseek-R1 on the OpenHuEval benchmark demonstrates the effectiveness of LRMs architectures in Hungarian language tasks and domain-specific scenarios. This result underscores the potential of LRMs as a key avenue of exploration in advancing Artificial General Intelligence (AGI).

Model size: From the results, models with larger parameter sizes perform better on OpenHuEval.

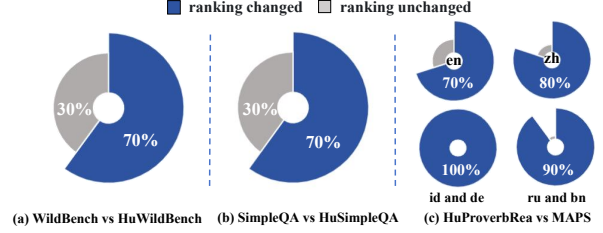


Figure 3: Comparison of model performance on OpenHuEval and similar datasets, highlighting that most LLMs experience rank changes.

For example, GPT-4o, Llama-3.1-Instruct-70B, and Qwen2.5-Instruct-72B outperform their smaller counterparts in the same series (GPT-4o-mini, Llama-3.1-Instruct-7B, and Qwen2.5-Instruct-7B) across all tasks.

3.3 Comparison with Existing Benchmarks

We compare the LLM’s performance rankings on the datasets in OpenHuEval and the existing similar datasets:

HuWildBench vs WildBench (Lin et al., 2024):

As shown in Figure 3(a), 70% LLMs experienced ranking changes. Detailed results can be found in Table 9.

HuSimpleQA vs SimpleQA (Wei et al., 2024):

As shown in Figure 3(b), 70% LLMs experienced ranking changes. Detailed results can be found in Table 12.

HuProverbRea vs MAPS: HuProverbRea was constructed with reference to MAPS (Liu et al., 2024c), which is the proverb reasoning dataset comprising six subsets, each corresponding to a different language: English (en), German (de), Russian (ru), Bengali (bn), Mandarin Chinese (zh), and Indonesian (id). We compared the model performance rankings on HuProverbRea and each subset of MAPS. As shown in Figure 3(c), the percentage of LLMs with ranking changes were: 70% for en, 80% for zh, 90% for ru & bn, and 100% for id & de. Detailed results can be found in Table 15 and Table 16.

These results underscore the importance of evaluating LLMs on Hungarian proverbs and Hungarian-specific questions, highlighting the need for targeted optimization of models to better handle language-specific proverbs and cultural nuances across diverse languages.

Model	HuWildBench	HuSimpleQA	HuProverbRea		HuMatchingFIB		HuStandardFIB	
	WBScore	Acc	Acc. (OE)	Acc. (2CQ)	B acc.	Q acc.	B acc.	Q acc.
GPT-4o	81.09	50.3	89.16	95.51	77.78	43.88	57.36	15.05
GPT-4o-mini	74.19	25.56	84.67	92.16	55.68	19.78	35.08	7.53
QwQ	58.02	9.09	67.49	84.23	38.65	12.23	6.05	0
Deepseek-R1	82.96	34.58	82.29	91.72	80.87	47.12	61.76	17.2
Deepseek-V3	78.42	32.71	83.26	92.51	68.87	39.93	51.44	9.68
Llama-3.1-Instruct-70B	61.78	35.99	80.18	93.83	59.56	24.46	40.99	6.45
Llama-3.1-Instruct-8B	53.62	15.2	63.35	73.48	5.74	0.72	16.64	1.08
o1-mini	76.43	15.8	77.44	87.67	60.83	17.63	45.25	13.98
Qwen2.5-Instruct-72B	74.05	14.9	77.8	90.22	63.8	24.1	32.32	8.6
Qwen2.5-Instruct-7B	42.01	5.22	50.48	67.05	31.88	1.08	7.43	0

Table 5: Overall performance of 10 LLMs on OpenHuEval. The first, second, and third place in each metric are marked with red, green, and blue text, respectively. In the FIB task evaluation metric, **B** represents the blank level, and **Q** represents the question level.

4 Framework for Analyzing the Thinking Process of LRM

When responding to the user’s query, the LRM’s response typically consists of two parts: the **thinking process** and the **answer**.⁸ We developed the framework for the in-depth analysis of the LRM’s thinking process on OpenHuEval. For LRM, we chose Deepseek-R1 and QwQ, as these are the only two models with accessible reasoning processes.

4.1 Task Selection

Among the OpenHuEval tasks, we selected HuSimpleQA and HuMatchingFIB as our subjects of study. Unlike recent work (Wang et al., 2025), which focuses solely on math reasoning datasets, the two tasks we selected each have distinctive characteristics: HuSimpleQA assesses the LLM’s ability to recall and retrieve Hungarian-specific knowledge, as well as its awareness of its own knowledge boundaries. HuMatchingFIB involves questions where multiple competitive blanks exist within the same problem, requiring the model to carefully choose which answers to fill in. Therefore, analyzing these tasks allows us to explore the reasoning mechanisms of LRMs in both broader and deeper contexts, providing the research community with more valuable conclusions and insights.

4.2 Analysis on HuSimpleQA

Method: Following (Wang et al., 2025), the LRM’s thinking process can be broken down into “thoughts”. A “thought” refers to the intermediate cognitive step output by a LRM during its thinking process. Throughout the thinking process, the LRM transitions between multiple thoughts, which are typically separated by reflective phrases such as

⁸In DeepSeek-V1, the thinking process and answer are enclosed within `<think>` `</think>` and `<answer>` `</answer>` tags, respectively.

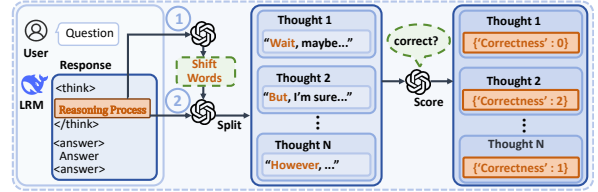


Figure 4: Method for analyzing the LRM’s thinking process on HuSimpleQA.

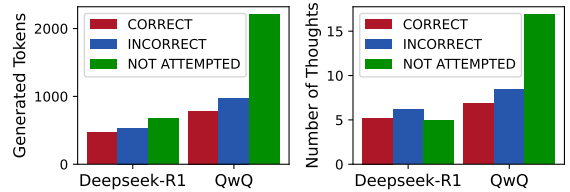


Figure 5: Average number of tokens and thoughts generated per response on HuSimpleQA.

“Alternative”, “Várni”(wait). The examples of the thoughts and the transitions can be found in Figure 40 and Figure 41. Thoughts can be further classified as “correct” or “incorrect”: reasoning along correct thoughts leads to CORRECT responses, while incorrect thoughts result in INCORRECT.

The LRM’s responses to HuSimpleQA have been judged as CORRECT, INCORRECT, or NOT_ATTEMPTED. Then we used GPT-4o to split the thinking process into thoughts (see the prompt in Figure 37 and Figure 38). We evaluated the correctness of each thought (see the prompt in Figure 39), with examples provided Figure 40 and Figure 41 in Appendix H.

Efficiency of thinking process: We measure the length of the process (in terms of token count) and the number of thoughts under the three evaluation outcomes of the responses, as shown in Figure 5. The results indicate that both the reasoning length (in tokens) and the thought count were generally shorter for Deepseek-R1 compared to QwQ. Con-

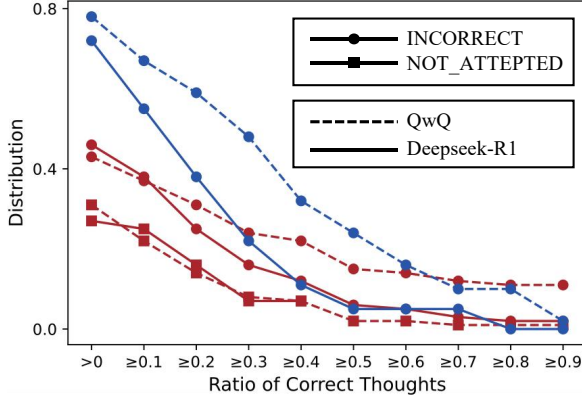


Figure 6: The distribution of thought correctness ratio in INCORRECT and NOT_ATTEMPTED responses. Red lines for HuSimpleQA and blue lines for MATH500-Hard from (Wang et al., 2025).

sidering that Deepseek-R1 performs better than QwQ on the HuSimpleQA task, it suggests that Deepseek-R1 achieves its superior performance with relatively lower reasoning overhead.

Confidence in thinking process: For Deepseek-R1, the reasoning length and thought count showed no significant differences across the three types of evaluation outcomes. In contrast, for QwQ, the length and the number of thoughts were significantly higher in the NOT_ATTEMPTED cases compared to the other two types. This observation suggests that, compared to Deepseek-R1, QwQ is less “confident”, which tends to repeatedly attempt generating answers when faced with uncertainty and is more inclined to abstain from answering altogether.

Correct thoughts in INCORRECT responses: In HuSimpleQA, we analyzed the ratio of INCORRECT and NOT_ATTEMPTED responses from Deepseek-R1 and QwQ that contained correct thoughts. We compared these results with those from MATH500-Hard as reported in (Wang et al., 2025), as shown in Figure 6. In the math reasoning task (MATH500-Hard), the significant portion of LRM’s ultimately incorrect responses still included correct thoughts: 72% for Deepseek-R1 and 78% for QwQ contained at least one correct thought. However, in HuSimpleQA, the ratio of INCORRECT responses containing correct thoughts dropped to 46% for Deepseek-R1 and 42% for QwQ. These findings suggest that the reasoning processes of LRMs differ significantly between memory-intensive tasks (HuSimpleQA) and reasoning-intensive tasks (MATH500-Hard), highlighting the need for targeted analysis and optimization.

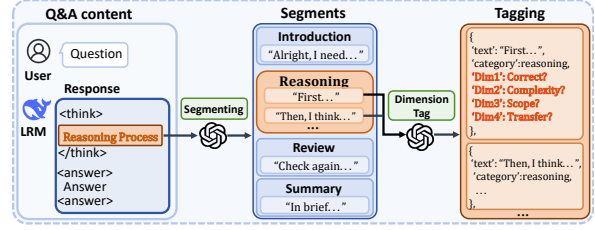


Figure 7: Method for analyzing the LRM’s thinking process on HuMatchingFIB.

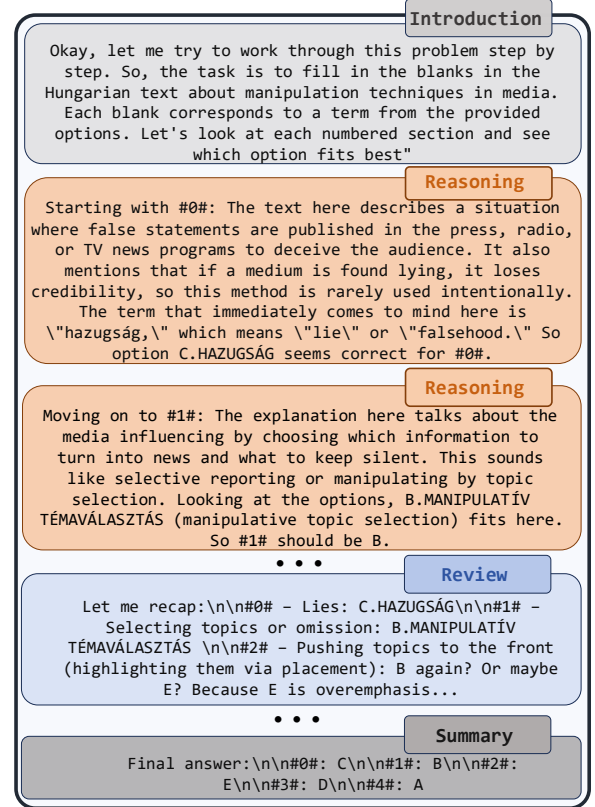


Figure 8: Example for splitting DeepSeek-R1’s thinking process into segments and categorizing these segments on HuMatchingFIB.

4.3 Analysis on HuMatchingFIB

Method: Unlike HuSimpleQA where each query contains only one question, HuMatchingFIB involves **multiple competitive blanks within the same question** that need to be filled (see the example in Figure 2 and Figure 33). Based on extensive case studies, we have developed the analytical method to facilitate further in-depth analysis, as illustrated in Figure 7.

Initially, we split the thinking process of LRMs into several segments, categorizing each as introduction, reasoning, review, or summary. Typically, each thinking process begins with an introduction segment, includes several

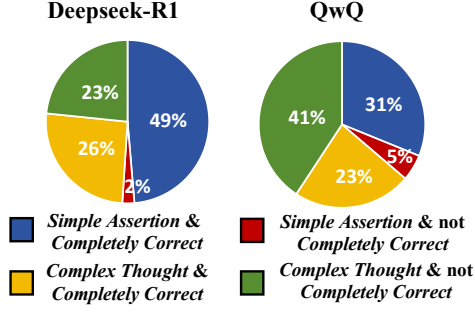


Figure 9: Tagging results of LRM’s reasoning segments in the correctness and complexity dimensions.

reasoning segments and some review segments in the middle, and concludes with a summary segment. In each reasoning segment, the LRM addresses one or more blanks through reasoning. During the review segment, there is usually a reflection on the completed blanks, with potential revisions to the answers for earlier blanks. Examples of these segments can be found in Figure 8, and detailed definitions are provided in Table 18.

We further tagged the reasoning segments according to the following four dimensions: **(Dim1) correctness**: Are the answers in this reasoning segment correct? **(Dim2) complexity**: In this reasoning segment, does the LRM simply assert the answer, or does it involve more complex reasoning? **(Dim3) scope**: Does this reasoning segment focuses on a single blank, modifies previous blanks, or addresses multiple blanks? **(Dim4) language transfer**: Does the LRM switch languages within this reasoning segment? The details of the tagging can be found in Appendix H.2.2. Examples of the tagging results can be found in Figure 47 and Figure 48.

Simple Assertion or Complex Thought? As shown in Table 19, the reasoning segments can be categorized into two classes. The first is referred to as *Simple Assertion*, where LRM directly provides the answer to the blank. The second type is termed *Complex Thought*, where the segment involves repeated thinking, logical reasoning, hypothesis validation, or other complex processes. Examples can be found in Figure 47 and 48.

We analyzed the correctness of the reasoning segments for both *Simple Assertion* and *Complex Thought* (tagging results of dim1), as shown in Figure 9. By comparing the statistical distributions of the Deepseek-R1 and QwQ models, we noted the following: Firstly, the proportion of *Simple Assertion* segments that are not *Completely Cor-*

rect is quite low (3% for Deepseek-R1 and 6% for QwQ), indicating that both models achieve high accuracy when resolving blanks through *Simple Assertion*. This suggests that the models’ calibration is relatively reliable, implying that the models “know what they know”. Secondly, the proportion of *Simple Assertion* segments that are *Completely Correct* is significantly higher for Deepseek-R1 at 49%, compared to only 31% for QwQ. This difference reflects the performance disparity between the two models in the thinking processes. Thirdly, the correctness for *Complex Thought* is notably lower than for *Simple Assertion*, and both models show a higher proportion of cases where no conclusion is reached for the *Complex Thought* tag. This indicates that when faced with unfamiliar problems, the LRM could actively extend the reasoning and analysis process.

Explicit Translation Insertion (ETI): We observed that in some reasoning segments, when faced with the problem in Hungarian, the LRM first translates the key phrase of the original question into English and then proceeds with analysis and reasoning based on this translation. For example, “... *Erőteljes #3# és a költői #4# gazdag használata jellemzi. This translates to* “It is characterized by strong #3# and rich use of poetic #4#.” ...”. We refer to this phenomenon as Explicit Translation Insertion (ETI). Statistical analysis shows that ETI occurs in 42.5% of DeepSeek-R1’s reasoning segments, while for QwQ, the proportion is 31.9%. This demonstrates that the LRM can adaptively translate input from the language in which it is relatively less proficient (such as Hungarian) into the more proficient language (such as English) during its thinking process. By scaling up the length of reasoning during the test stage, the LLM can partially overcome the cross-language performance discrepancies typical of traditional instruction-based LLM.

5 Conclusion

In this paper, we constructed the first benchmark for LLMs focusing on the Hungarian language and its specifics. The results highlight the significant need for evaluation and model optimization tailored to Hungarian language and specifics. We also developed the framework for analyzing the thinking processes of the cutting-edge LRMs. Our work not only advances LLM technology in Hungarian but also provides valuable insights for studying languages of other countries and regions.

6 Limitation

Given the rapid advancements in English evaluation datasets, this research represents only an initial step towards bridging the gap between Hungarian and English evaluation resources. Currently, evaluation datasets for less-resourced languages like Hungarian still lag behind their English counterparts in terms of depth and breadth. Moving forward, we plan to closely monitor developments in English evaluation methodologies, continually refining and enhancing evaluation techniques and datasets for low-resource languages to reduce this disparity.

Moreover, as the field of LLMs evolves rapidly, many promising models, especially those tailored for low-resource languages, remain under-evaluated. Our future goal is to establish a vibrant OpenHuEval community that will regularly update evaluation results for the latest models. This will ensure comprehensive and up-to-date assessments, fostering the optimization and development of models in the low-resource language domain.

7 Ethical Consideration

This work involved human annotation. For all annotators, we explicitly informed them about the use of the data and required them to ensure that the questions included in OpenHuEval do not involve any social bias, ethical issues or privacy concerns during the annotation process.

Acknowledgments

This research was supported by National Key R&D Program of China (2022ZD0160201).

References

- David Ifeoluwa Adelani, Hannah Liu, Xiaoyu Shen, Nikita Vassilyev, Jesujoba O Alabi, Yanke Mao, Hao-nan Gao, and Annie En-Shiun Lee. 2023. Sib-200: A simple, inclusive, and big evaluation dataset for topic classification in 200+ languages and dialects. *arXiv preprint arXiv:2309.07445*.
- Kabir Ahuja, Harshita Diddee, Rishav Hada, Millicent Ochieng, Krithika Ramesh, Prachi Jain, Akshay Nambi, Tanuja Ganu, Sameer Segal, Maxamed Axmed, et al. 2023. Mega: Multilingual evaluation of generative ai. *arXiv preprint arXiv:2303.12528*.
- Shane Arora, Marzena Karpinska, Hung-Ting Chen, Ipsita Bhattacharjee, Mohit Iyyer, and Eunsol Choi. 2024. Calmqa: Exploring culturally specific long-form question answering across 23 languages. *arXiv preprint arXiv:2406.17761*.
- Lucas Bandarkar, Davis Liang, Benjamin Muller, Mikel Artetxe, Satya Narayan Shukla, Donald Husa, Naman Goyal, Abhinandan Krishnan, Luke Zettlemoyer, and Madian Khabsa. 2023. The belebele benchmark: a parallel reading comprehension dataset in 122 language variants. *arXiv preprint arXiv:2308.16884*.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Yi Fung, Ruining Zhao, Jae Doo, Chenkai Sun, and Heng Ji. 2024. Massively multi-cultural knowledge acquisition & lm benchmarking. *arXiv preprint arXiv:2402.09369*.
- Yancheng He, Shilong Li, Jiaheng Liu, Yingshui Tan, Weixun Wang, Hui Huang, Xingyuan Bu, Hangyu Guo, Chengwei Hu, Boren Zheng, et al. 2024. Chinese simpleqa: A chinese factuality evaluation for large language models. *arXiv preprint arXiv:2411.07140*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2020. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*.
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. Xtreme: A massively multilingual multi-task benchmark for evaluating cross-lingual generalisation. In *International Conference on Machine Learning*, pages 4411–4421. PMLR.
- Kaiyu Huang, Fengran Mo, Hongliang Li, You Li, Yuanchi Zhang, Weijian Yi, Yulong Mao, Jinchen Liu, Yuzhuang Xu, Jinan Xu, et al. 2024. A survey on large language models with multilingualism: Recent advances and new frontiers. *arXiv preprint arXiv:2405.10936*.
- Xu Huang, Wenhao Zhu, Hanxu Hu, Conghui He, Lei Li, Shujian Huang, and Fei Yuan. 2025. [Benchmax: A comprehensive multilingual evaluation suite for large language models](#). *Preprint*, arXiv:2502.07346.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Os-trow, Akila Welihinda, Alan Hayes, Alec Radford, et al. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. 2024. Openai o1 system card. *arXiv preprint arXiv:2412.16720*.

- Fajri Koto, Nurul Aisyah, Haonan Li, and Timothy Baldwin. 2023. [Large language models only pass primary school exams in Indonesia: A comprehensive test on IndoMMLU](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12359–12374, Singapore. Association for Computational Linguistics.
- Fajri Koto, Haonan Li, Sara Shatnawi, Jad Doughman, Abdelrahman Boda Sadallah, Aisha Alraesi, Khalid Almubarak, Zaid Alyafeai, Neha Sengupta, Shady Shehata, et al. 2024. Arabicmmlu: Assessing massive multitask language understanding in arabic. *arXiv preprint arXiv:2402.12840*.
- Viet Dac Lai, Chien Van Nguyen, Nghia Trung Ngo, Thuat Nguyen, Franck Dernoncourt, Ryan A Rossi, and Thien Huu Nguyen. 2023. Okapi: Instruction-tuned large language models in multiple languages with reinforcement learning from human feedback. *arXiv preprint arXiv:2307.16039*.
- Haitao Li, Qian Dong, Junjie Chen, Huixue Su, Yujia Zhou, Qingyao Ai, Ziyi Ye, and Yiqun Liu. 2024a. Llm-as-judges: a comprehensive survey on llm-based evaluation methods. *arXiv preprint arXiv:2412.05579*.
- Haonan Li, Yixuan Zhang, Fajri Koto, Yifei Yang, Hai Zhao, Yeyun Gong, Nan Duan, and Timothy Baldwin. 2023. Cmmu: Measuring massive multitask language understanding in chinese. *arXiv preprint arXiv:2306.09212*.
- Jialin Li, Junli Wang, Junjie Hu, and Ming Jiang. 2024b. How well do llms identify cultural unity in diversity? *arXiv preprint arXiv:2408.05102*.
- Noémi Ligeti-Nagy, Gergő Ferenczi, Enikő Héja, László János Laki, Noémi Vadász, Zijian Győző Yang, and Tamás Váradi. 2024. Hulu: Hungarian language understanding benchmark kit. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 8360–8371.
- Bill Yuchen Lin, Yuntian Deng, Khyathi Chandu, Faeze Brahman, Abhilasha Ravichander, Valentina Pyatkin, Nouha Dziri, Ronan Le Bras, and Yejin Choi. 2024. Wildbench: Benchmarking llms with challenging tasks from real users in the wild. *arXiv preprint arXiv:2406.04770*.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2021. Truthfulqa: Measuring how models mimic human falsehoods. *arXiv preprint arXiv:2109.07958*.
- Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. 2024a. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*.
- Chen Cecilia Liu, Iryna Gurevych, and Anna Korhonen. 2024b. Culturally aware and adapted nlp: A taxonomy and a survey of the state of the art. *arXiv preprint arXiv:2406.03930*.
- Chen Cecilia Liu, Fajri Koto, Timothy Baldwin, and Iryna Gurevych. 2024c. [Are multilingual llms culturally-diverse reasoners? an investigation into multicultural proverbs and sayings](#). *Preprint*, arXiv:2309.08591.
- Junho Myung, Nayeon Lee, Yi Zhou, Jiho Jin, Rifki Putri, Dimosthenis Antypas, Hsuvas Borkakoty, Eunsu Kim, Carla Perez-Almendros, Abinew Ali Ayele, et al. 2024. Blend: A benchmark for llms on everyday knowledge in diverse cultures and languages. *Advances in Neural Information Processing Systems*, 37:78104–78146.
- Tarek Naous, Michael J Ryan, Alan Ritter, and Wei Xu. 2023. Having beer after prayer? measuring cultural bias in large language models. *arXiv preprint arXiv:2305.14456*.
- Attila Novák, Borbála Novák, Tamás Zombori, Gergő Szabó, Zsolt Szántó, and Richárd Farkas. 2023. A question answering benchmark database for hungarian. In *Proceedings of the 17th Linguistic Annotation Workshop (LAW-XVII)*, pages 188–198.
- Mátyás Osváth, Zijian Győző Yang, and Karolina Kósa. 2023. Analyzing narratives of patient experiences: A bert topic modeling approach. *Acta Polytechnica Hungarica*, 20(7):153–171.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. Squad: 100,000+ questions for machine comprehension of text. *arXiv preprint arXiv:1606.05250*.
- Yusuke Sakai, Hidetaka Kamigaito, and Taro Watanabe. 2024. mcsqa: Multilingual commonsense reasoning dataset with unified creation strategy by language models and humans. *arXiv preprint arXiv:2406.04215*.
- Guijin Son, Hanwool Lee, Sungdong Kim, Seungone Kim, Niklas Muennighoff, Taekyoon Choi, Cheonbok Park, Kang Min Yoo, and Stella Biderman. 2024. Kmmlu: Measuring massive multitask language understanding in korean. *arXiv preprint arXiv:2402.11548*.
- Jiaxing Sun, Weiquan Huang, Jiang Wu, Chenya Gu, Wei Li, Songyang Zhang, Hang Yan, and Conghui He. 2024. Benchmarking chinese commonsense reasoning of llms: From chinese-specifics to reasoning-memorization correlations. *arXiv preprint arXiv:2403.14112*.
- DeepSeek Team. 2024a. Deepseek-r1-lite-preview is now live: unleashing supercharged reasoning power.
- Qwen Team. 2024b. Qwq: Reflect deeply on the boundaries of the unknown. *Hugging Face*.
- Yue Wang, Qiuzhi Liu, Jiahao Xu, Tian Liang, Xingyu Chen, Zhiwei He, Linfeng Song, Dian Yu, Juntao Li, Zhuosheng Zhang, et al. 2025. Thoughts are all over the place: On the underthinking of o1-like llms. *arXiv preprint arXiv:2501.18585*.

- Jason Wei, Nguyen Karina, Hyung Won Chung, Yunxin Joy Jiao, Spencer Papay, Amelia Glaese, John Schulman, and William Fedus. 2024. Measuring short-form factuality in large language models. *arXiv preprint arXiv:2411.04368*.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, et al. 2024. Qwen2. 5 technical report. *arXiv preprint arXiv:2412.15115*.
- Da Yin, Hritik Bansal, Masoud Monajatipoor, Lillian Harold Li, and Kai-Wei Chang. 2022. Geomlana: Geo-diverse commonsense probing on multilingual pre-trained language models. *arXiv preprint arXiv:2205.12247*.
- Arda Yüksel, Abdullatif Köksal, Lütfi Kerem Şenel, Anna Korhonen, and Hinrich Schütze. 2024. Turkishmmlu: Measuring massive multitask language understanding in turkish. *arXiv preprint arXiv:2407.12402*.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. [HellaSwag: Can a machine really finish your sentence?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4791–4800, Florence, Italy. Association for Computational Linguistics.
- Yidan Zhang, Boyi Deng, Yu Wan, Baosong Yang, Haoran Wei, Fei Huang, Bowen Yu, Junyang Lin, and Jingren Zhou. 2024. P-mmeval: A parallel multilingual multitask benchmark for consistent evaluation of llms. *arXiv preprint arXiv:2411.09116*.

A Related Works

A.1 Hungarian benchmarks for LLMs from Translation

Directly translating an existing English evaluation dataset into non-English is an effective and straightforward method for constructing non-English evaluation datasets. Many existing multilingual evaluation datasets are constructed using this approach, such as those described in (Ahuja et al., 2023; Huang et al., 2024). Among these, quite a few evaluation datasets include Hungarian as one of the languages, such as (Hu et al., 2020) and (Adelani et al., 2023). (Lai et al., 2023) utilized GPT-3.5 to translate datasets such as ARC (Clark et al., 2018), MMLU (Hendrycks et al., 2020), TruthfulQA (Lin et al., 2021), and HellaSwag (Zellers et al., 2019) into Hungarian. Belebele (Bandarkar et al., 2023) includes Hungarian as part of its multilingual parallel corpus reading comprehension evaluation dataset. BenchMaX (Huang et al., 2025) is a multilingual parallel corpus evaluation dataset focused on comprehensively assessing LLMs’ capabilities across various generative tasks.

These translation-based parallel corpus evaluation datasets provide a comprehensive assessment of LLMs’ language-agnostic capabilities in Hungarian, including areas like world knowledge, mathematical reasoning, logical reasoning, and code generation. However, they overlook the unique characteristics and linguistic features of Hungarian, such as language nuances, culture, history, and regional context, which are critical for Hungarian users.

A.2 Language specific benchmarks for LLMs

There has been extensive research focusing on evaluating the unique features and capabilities of languages (Liu et al., 2024b). Some studies have constructed benchmarks similar to MMLU (Hendrycks et al., 2020) by collecting exam questions specific to various countries, such as CMMLU (Li et al., 2023), IndoMMLU (Koto et al., 2023), ArabicMMLU (Koto et al., 2024), KMMLU (Son et al., 2024), and TurkishMMLU (Yüksel et al., 2024). Other studies have built evaluation datasets by crawling material and user queries from internet forums, filtering for queries related to linguistic and cultural features, such as the benchmark in (Naous et al., 2023) and CaLMQA (Arora et al., 2024). Some works manually construct evaluation datasets that emphasize linguistic and cultural char-

acteristics, such as BLEND (Myung et al., 2024), CHARM (Sun et al., 2024), and HuLU (Ligeti-Nagy et al., 2024). Others adopt a “LLM generation combined with human expert modification and review” method to construct culturally characteristic evaluation datasets, such as MAPS (Liu et al., 2024c), mCSQA (Sakai et al., 2024), and ChineseSimpleQA (He et al., 2024). Although Hungarian is rarely covered in these works, they provide crucial inspiration and approaches for our work.

Among these works, MILQA (Novák et al., 2023) is the Hungarian question answering benchmark created mainly following the SQuAD (Rajpurkar et al., 2016). HuLU (Ligeti-Nagy et al., 2024) is a comprehensive Hungarian evaluation benchmark kit⁹ that includes a total of seven tasks and corresponding datasets. Of the seven tasks, four are constructed by translating existing English evaluation datasets, and three are manually created based on native Hungarian corpora. However, HuLU only supports multiple-choice and true/false questions, which limits its ability to assess broader LLM capabilities such as language generation, open-domain Q&A, reasoning, and instruction-following.

B Hungarian-Specific Dimensions

OpenHuEval encompasses eight Hungarian-Specific Dimensions (HuSpecificDim), as shown in Table 2: *Language (L)*, *History (H)*, *Life, Culture, and Customs (LCC)*, *Education and Profession (EP)*, *Geography and Place (GP)*, *Figure (F)*, *Politics, Policy, and Law (PPL)*, and *Business and Finance (BF)*.

To select these Hungarian-Specific Dimensions, we first conducted extensive literature research to summarize the knowledge dimensions in existing cultural benchmarks, as shown in Table 6.

Based on above and by considering the knowledge needs of typical Hungarian users and the content on authoritative Hungarian websites, we established the Hu-specific dimensions for OpenHuEval. We believe these dimensions cover most user needs related to Hungarian specifics.

C Inference Detail

C.1 Default Maximum Output Length for LLM in Opencompass

For traditional instruction-based LLMs, we adopted OpenCompass’s default settings for the

⁹<https://hulu.nytud.hu/>

Benchmark	Dimension
CHARM (Sun et al., 2024)	History, Traditional Culture and Arts, Daily Life and Customs, Entertainment, Public Figures, Geography, Chinese Language
CUNIT (Li et al., 2024b)	Clothing, Food
BLEnD (Myung et al., 2024)	Food, Sports, Family, Education, Holidays/Celebrations/Leisure, Work-Life
CultureAtlas (Fung et al., 2024)	Culture, Holidays, Dining Etiquette, Education, Honorifics, etc.
GEOMLAMA (Yin et al., 2022)	Habits and Personal Choices, Cultures and Customs, Policies and Regulations, Geography

Table 6: The knowledge dimensions present in existing cultural benchmarks.

Model	Max out length
Qwen2.5-Instruct-7B	4096
Qwen2.5-Instruct-72B	4096
GPT4-o	2048
GPT-4o-mini	2048
Llama-3.1-Instruct-8B	1024
Llama-3.1-Instruct-70B	4096
Deepseek-V3	2048

Table 7: Default maximum output length for LLM in OpenCompass.

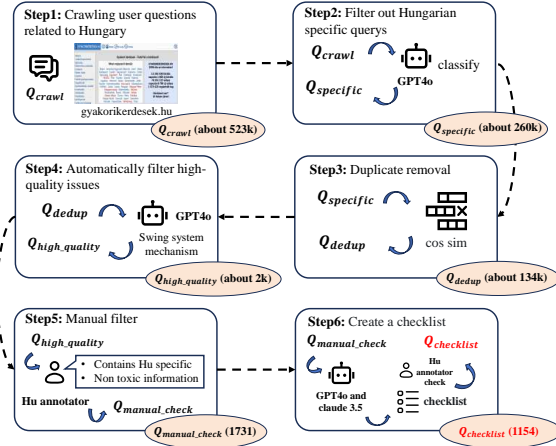


Figure 10: Constuction of HuWildBench.

maximum output length, as shown in Table 7.

D HuWildBench

D.1 Overall Construction Pipeline of HuWildBench

The construction of HuWildBench contains the following steps(see Figure 10):

(1) Crawling: All user queries on the g13k website¹⁰ are systematically categorized into a multi-level tag system, which consists of 27 primary tags and 231 secondary tags. We manually reviewed all the secondary tags and selected 37 of them that contain a higher number of questions related

to Hungary, such as népszokások (folk customs), egészségügyi-ellátások (healthcare services), and rezsi (overheads). We then crawled user queries under these 37 secondary tags, with a query date range from January 1, 2019, to August 31, 2024, resulting in the dataset Q_{crawl} (approximately 523K queries).

(2) Filtering for Hungary-specific content: Although these 37 secondary tags are closely related to Hungary, many of the questions still do not focus on Hungary-specific topics. Therefore, we used GPT-4o to classify the questions in Q_{crawl} (detailed prompt in Figure 13), resulting in a subset of approximately 260K questions, $Q_{specific}$.

(3) Deduplication: To ensure the diversity of questions, we performed deduplication on the Hungary-related questions within each secondary tag. The detailed process is outlined in Appendix D.2. After deduplication, the number of user questions in Q_{dedup} was approximately 134K.

(4) Automatic high-quality question filtering: To ensure that only high-quality questions are extracted from the question pool, we designed a comparative-based high-quality question filtering strategy, as detailed in Appendix D.3. After filtering, the resulting set $Q_{high_quality}$ contained around 2K questions.

(5) Manual filtering: We hired a group of Hungarian native speakers to further manually review the questions in $Q_{high_quality}$. Only questions that met the following two criteria were retained: First, the question should be Hungary-specific and closely related to Hungary. Second, the question must be harmless, meaning it does not contain inappropriate content such as pornography, violence, politics, or taboo topics specific to Hungary. The final set Q_{manual_check} consists of 1731 questions.

(6) Checklist construction: Based on WildBench (Lin et al., 2024), we constructed a checklist for each question. The purpose of the checklist is to assist the LLM judge in evaluating the answers. Each item in the checklist queries a specific

¹⁰<https://www.gyakorikerdesek.hu/>

aspect of the answer to a question. An example of the checklist can be found in Table 21, and the detailed construction method is provided in Appendix D.4. To ensure the relevance of the checklist items to the questions, we hired a Hungarian native speaker to review the checklist for quality, filtering out non-compliant items and performing deduplication. The filtering criterion was whether the item was suitable as an evaluation dimension for the model’s response. To ensure the reliability of the LLM-as-judge, we filtered out user questions with fewer than 8 checklist items. The final set $Q_{checklist}$ contains 1154 questions. In the end, we obtained 1154 user questions along with their corresponding checklists.

D.2 Deduplication of similar questions

Since there are similar questions in the results obtained in the previous step, we design a method to remove similar ones. Specifically, we first use the SentenceTransformer (Osváth et al., 2023)¹¹ model to extract the embedding vector of each question. Then, we calculate the cosine similarity between the embedding of each two questions, and choose a threshold between [0.15-0.25] according to the number of questions under each secondary tag. The larger the number of problems, the larger the threshold. Finally, one of the questions whose similarity is less than the threshold is removed, ensuring that the similarities between all questions are greater than the threshold.

D.3 Automatic high-quality question filtering

In order to automate the filtering of high-quality sample pots, we constructed the prompt that allows the GPT-4o to select the two best Hungarian questions out of the five based on the criteria of linguistic complexity, Hungarian relevance, common-sense accuracy, context-dependence, answer diversity, ambiguity, reasoning requirements, socio-ethical considerations, format diversity, and breadth of knowledge and outputs their indexes in JSON format to output their indexes, as shown in Figure 14. Specifically, we first set the criteria for high-quality questions in Prompt. Then we ask GPT-4o to compare the input questions based on the criteria. In order to mitigate the occurrence of some high-quality questions being eliminated prematurely (or vice versa) when all the questions in the same batch are of high quality, we follow the

following 3 rules when filtering the high-quality questions: (1) filter 2 high-quality questions from 5 questions at a time, instead of filtering 1 high-quality question directly from 2 questions. (2) use the Swiss system mechanism instead of the knock-out mechanism. In each screening round, each question can win in the current round as long as it ensures that it wins in two comparisons, and it will not be eliminated directly because of a failure in one comparison. (3) Our question screening strategy eliminates 65% of the questions in each round, in order to ensure that each secondary label has a sufficient number of high-quality questions. We conducted different elimination rounds for questions under different labels, and finally got about 2K questions.

To ensure the reasonableness and robustness of the standard, we reviewed the results obtained using it. We selected 500 sets of questions filtered by GPT-4o according to this standard (where the model had to select two high-quality questions from each set of five). Human experts annotated these sets, and the overlap rate between the selections of the human experts and the model was 81% (a set was considered overlapping if the two questions chosen by the model matched those chosen by the experts). This demonstrates the effectiveness of the standard for selecting high-quality questions. The model’s recall rate for the questions selected by the human experts was 90%.

The final constructed HuWildBench is shown in Table 21.

D.4 Checklist construction

In the process of building the Checklist, we mainly use large language models to generate it. In order to ensure the diversity of the Checklist and make the judge model can better evaluate the quality of the answers, here we use two non-open source LLM GPT-4o and Claude-3.5, each model generates a list of length 3-5. then we merge the two Checklists into one final Checklist. Checklists are then merged into a final Checklist. ultimately, each problem has a length of 6-10 and a Checklist. The details of our designed Prompt are shown in Figure 15 and the final constructed partial Checklist is shown in Table 21.

D.5 LLM-as-judge

Following WB-Score (Lin et al., 2024), we use GPT-4o as the judge model to evaluate the score of responses generated by LLM. The prompt used in

¹¹<https://github.com/UKPLab/sentence-transformers>

Score	Definition
Score 1-2	The response is very poor and does not make sense at all.
Score 3-4	The response is poor and does not help the user solve the problem meaningfully.
Score 5-6	The response is fair but has issues (e.g., factual errors, hallucinations, missing key information).
Score 7-8	The response is good but could be improved.
Score 9-10	The response is perfect and provides helpful information to solve the problem.

Table 8: Definition of WB-Scores.

Rank	WildBench	HuWildBench
1	Deepseek-V3	Deepseek-R1(↑3)
2	GPT-4o	GPT-4o(-)
3	o1-mini	Deepseek-V3(↓2)
4	Deepseek-R1	o1-mini(↓1)
5	Qwen2.5-Instruct-72B	GPT-4o-mini(↑1)
6	GPT-4o-mini	Qwen2.5-Instruct-72B(↓1)
7	Llama-3.1-Instruct-70B	Llama-3.1-Instruct-70B(-)
8	QwQ	QwQ(-)
9	Qwen2.5-Instruct-7B	Llama-3.1-Instruct-8B(↑1)
10	Llama-3.1-Instruct-8B	Qwen2.5-Instruct-7B(↓1)

Table 9: Performance rankings on WildBench and HuWildBench.

judge is shown in Figure 16.

D.6 Comparing the Performance Ranking on WildBench and HuWildBench

We compare the performance rankings of LLM on HuWildBench and WildBench (Lin et al., 2024), as shown in Table 9.

E HuSimpleQA

E.1 Construction pipeline of HuSimpleQA

The construction of HuSimpleQA consists of the following steps, as shown in Figure 11:

(1) Obtaining corpora rich in Hungary-specific facts and knowledge: First, we chose the Hungarian Wikipedia website¹² as the source of corpus material for question construction. We crawled all the entry pages and extracted their content. Next, we used GPT-4o to classify whether the entries were Hungary-specific, with the prompts detailed in Appendix E.2, Figure 17.¹³ We then used

¹²<https://hu.wikipedia.org/>

¹³We did not classify all the pages but instead randomly

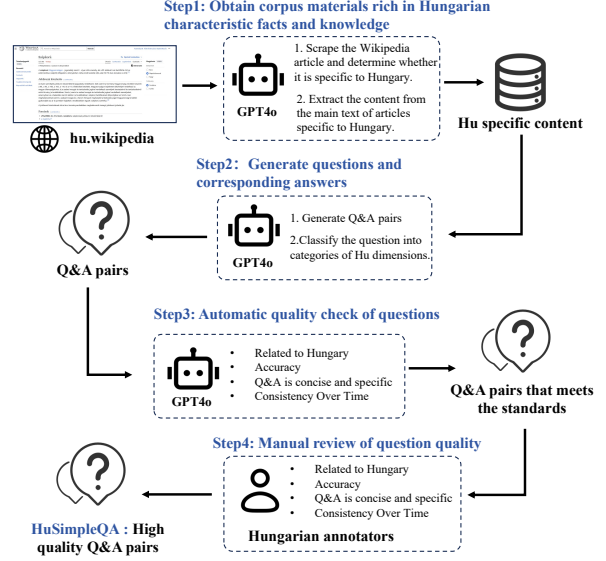


Figure 11: Construction of HuSimpleQA.

GPT to extract key information from the content of these entries suitable for question-answering. The extraction prompts are detailed in Appendix E.2, Figure 18. An example of the extracted key information is shown in Figure 19. As a result, we obtained Hungary-specific key information covering the eight distinct dimensions, totaling 4428 pieces of information.

(2) Generating questions and corresponding answers: We used the GPT-4o model to generate open-ended questions and corresponding answers based on the Hungary-specific key information obtained in the previous step. The prompt used is detailed in Appendix E.3, Figure 20. In this step, we generated a total of 9424 questions based on 4K entries. We then classified the generated questions according to the eight Hungary-specific dimensions outlined in Section 2.1, using GPT-4o with the prompt detailed in Appendix E.3, Figure 21.

(3) Automatic quality checking of questions: To ensure the quality of the questions, we used GPT-4o to check and filter the generated questions. We set the following four criteria, retaining only those questions that met all four standards (the corresponding prompt is detailed in Appendix E.4):

- **Criterion 1:** Hungary-specific: The content of the question-answer pair must align with the eight Hungary-specific dimensions proposed in this paper.

- **Criterion 2:** Accuracy: The information in selected pages until we reached 8K Hungary-specific entries, at which point we stopped.

the question-answer pair must align with the entry description and facts, and the answer should not be directly inferable from the question itself.

- **Criterion 3:** Concise and specific: The question and answer should be clear and concise, with no redundant information. The question should not contain nested sub-questions. The phrasing should be specific and direct, matching the scope of the answer (e.g., for time and location questions, the exact year/month/day/district/city must be specified).

- **Criterion 4:** Consistency Over Time: The answer should remain consistent over time and not be influenced by future events.

After the automatic checking process, we retained 5503 questions corresponding to 2666 entries.

(4) **Manual review of question quality:** To further ensure the quality of the questions, we hired Hungarian native speakers to manually review the questions. Annotators checked whether the questions met the four criteria mentioned in Step 3. During the annotation process, each question was assigned to two annotators, who received the questions but not the answers. A question was considered valid and retained only if both annotators agreed that it met all four criteria and that the provided answer matched the original reference answer. Detailed procedures are provided in Appendix E.4. After these four steps, we obtained a total of 1293 questions, with their distribution across the eight Hungary-specific dimensions shown in Table 22.

E.2 Obtaining corpora rich in Hungary-specific facts and knowledge

In the process of filtering Wikipedia entries with Hungarian characteristics, we randomly selected entries and provided both the entries and the first two paragraphs of the main content to GPT-4o (prompt shown in Figure 17) to determine if they were related to Hungary. If the entry was deemed relevant, it was categorized based on the eight characteristic dimensions proposed in this paper. At this stage, an “Others” category was added to ensure the focus on the eight thematic categories and to exclude interference from entries that belonged to other themes. The screening process stopped once the total number of Hungarian characteristic entries reached 8,000.

Due to the uneven distribution of entry themes on Wikipedia, with more data in the categories

of figures, geography and place, and history, we filtered the data based on the proportion of themes, ensuring that no single category exceeded 1,000 entries. This resulted in 4428 characteristic entries covering the eight dimensions.

Given the varying lengths of content describing entries on Wikipedia, we aimed to streamline the complexity of constructing subsequent question-answer pairs. To achieve this, we first employed GPT-4o to extract key information from the main text of each entry. This step aims to avoid any deviation from the theme caused by redundant content during the construction of the question-answer pairs (prompt shown in Figure 18). The results of the key information extraction are presented in Figure 19.

E.3 Generating questions and corresponding answers

Based on the key information extracted and the provided entries, we utilized GPT-4o to generate 1-3 Hungarian open-ended question-answer pairs for each entry (prompt details in Figure 20). In total, 9,424 question-answer pairs were generated based on 4,000 entries.

Given that the focus and orientation of the generated question-answer pairs may differ from the original entry categories, this paper employed GPT-4o to reclassify the obtained question-answer pairs, with the corresponding prompt detailed in Figure 21.

E.4 Automatic quality checking of questions

We focused on evaluating the quality of the generated questions from two perspectives: the information contained in the question-answer pairs and the formulation of the questions. The quality assessment was divided into two stages, with each stage generating two evaluation metrics. The first stage focuses on the relevance and correctness of the question information. We provided GPT-4o with the entry, its corresponding key information, and the generated question-answer pairs to verify whether the questions contain Hungarian-specific content and whether the information in the question-answer pairs aligns with the provided background material (prompt shown in Figure 22).

Second, from the perspective of the precision of the question formulation, we only provided GPT-4o with the generated question-answer pairs to simulate real user response scenarios. This step emphasized evaluating whether the questions were based

on objective facts, and whether the descriptions were precise and specific enough to allow independent answering without ambiguity. Additionally, we required that the answers remain unaffected by future events, ensuring consistency across any time period and guaranteeing the long-term validity of the dataset (prompt details in Figure 23).

Based on the results of the above automated quality assessment, we retained only those question-answer pairs that passed all four evaluation criteria, resulting in a final set of 5,503 questions.

E.5 Manual review of question quality

To further ensure the quality of the constructed question-answer pairs, we engaged native Hungarian speakers to review these questions. Each question was independently reviewed by two annotators who could only see the questions and not the reference answers. The annotation process consisted of three main steps.

First, the annotators were required to determine whether the given questions aligned with the eight Hungarian-specific dimensions proposed in this paper. Next, they evaluated whether the questions met the four assessment criteria outlined in Step 3, ensuring that the questions were objectively framed, precisely described, had unique answers, contained correct information, and maintained consistent answers over time. Finally, if a question satisfied all the above criteria, the annotators provided the correct answer. During this process, annotators were permitted to search for relevant information online and provided reference sources for their answers.

To address potential issues such as overly obscure questions or non-fixed answers, we used GPT-4o to verify whether the annotated results matched the generated reference answers. If the annotated answer matched the reference answer, it was labeled as "CORRECT"; otherwise, it was labeled as "INCORRECT" (prompt details in Figure 24). We selected question-answer pairs that both annotators deemed valid, Hungarian-specific, and consistent with the original reference answers as candidates for the HuSimpleQA dataset, resulting in a total of 2134 questions.

Considering that the HuSimpleQA dataset should exhibit diversity and broad coverage, we removed question-answer pairs belonging to the same entry, retaining only one question-answer pair per entry that best met the construction and evaluation criteria. The details of selecting optimal question-answer pairs prompt can be seen in Figure

25.

Through this process, we obtained a total of 1293 pieces of Hungarian open-ended question-answer pairs, with the category distribution shown in Table 22.

E.6 Inference prompt

We constructed prompts in two languages for model inference, as shown in Figure 26, while also instructing the model to provide a confidence score (ranging from 1 to 100) to measure the model's confidence in its generated answers.

E.7 LLM-as-judge

Following the approach of SimpleQA (Wei et al., 2024), we employed GPT-4o as the judge to evaluate the LLM's responses. The evaluation criteria for this step were similar to those used in the manual review process of Step 4. In addition to the classification labels CORRECT and INCORRECT, we introduced an additional category, NOT_ATTEMPTED to further assess the model's ability to respond to questions and the breadth of its knowledge coverage (prompt details in Figure 27).

Based on the results from the judge, we evaluate the performance of the LLM on HuSimpleQA using the following five metrics:

- **Correct (CO):** The predicted answer completely encompasses the reference answer without any conflicting or contradictory information.
- **Not Attempted (NA):** The predicted answer does not fully include the reference answer, but there are no contradictions between the two.
- **Incorrect (IN):** The predicted answer contradicts the reference answer, regardless of whether the contradiction is resolved.
- **Correct Given Attempted (CGA):** This metric measures the percentage of accurately answered questions out of all attempted questions.
- **F-score:** This metric calculates the harmonic mean between the proportion of correct answers and the proportion of correct answers among attempted questions.

The formulas for CGA and F-score are as follows:

$$CGA = \frac{c}{c + i} \quad (1)$$

$$F\text{-Score} = \frac{2}{\frac{c+i}{c} + \frac{c+i+n}{c}} = \frac{2c}{2c + 2i + n} \quad (2)$$

Here, c represents the number of CORRECTly answered questions, i represents the number of

Model	Judge Correct Ratio
Deepseek-R1	0.98
Deepseek-V3	1.00
Qwen2.5-Instruct-7B	0.98
Qwen2.5-Instruct-72B	0.98
Llama-3.1-Instruct-8B	1.00
Llama-3.1-Instruct-70B	1.00
GPT-4o	0.99
GPT-4o-mini	1.00
o1-mini	1.00
QwQ	0.98

Table 10: Accuracy of LLM-as-Judge for HuSimpleQA.

INCORRECTLY answered questions, and n represents the number of NOT_ATTEMPTED questions. To verify the effectiveness of LLM-as-judge, we randomly selected 100 samples from each of the 10 evaluated models for inspection, with accuracy exceeding 98% for all models, as shown in Table 10.

E.8 Comparing the Performance Ranking on HuSimpleQA and SimpleQA

We compare the performance rankings of LLM on HuSimpleQA and SimpleQA (Wei et al., 2024), as shown in Table 12.

E.9 Results of Translating HuSimpleQA to English

As shown in Table 5, the current leading LLMs perform poorly on HuSimpleQA. To further analyze and confirm whether these issues stem from the models’ lack of knowledge of the Hungarian language or Hungarian-specific cultural aspects, we translated HuSimpleQA into English and reevaluated the models. The results, presented in Table 13, show that after translation, 2 out of 10 models improved, while 8 experienced a decline in performance, indicating that translation generally worsens performance. This suggests that LLMs have more difficulty understanding Hungarian cultural nuances than the language itself. These findings are consistent with the conclusions in

E.10 Human Performance on SimpleQA and HuSimpleQA

We hire human annotators to complete the tasks in the SimpleQA and HuSimpleQA datasets in order to assess human performance on these datasets. HuSimpleQA was annotated by Hungarian native speakers, while SimpleQA was handled by English native speakers. We used two annotators per dataset, averaging their performance metrics to

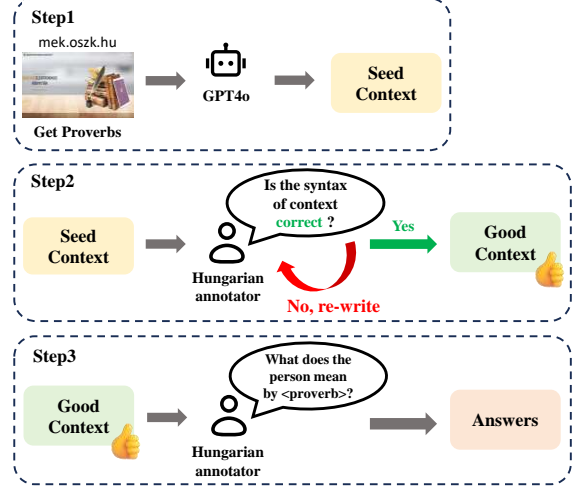


Figure 12: Construction of HuProverbRea.

determine human performance. Annotators were instructed NOT to use external tools like search engines or GPT. For HuSimpleQA and SimpleQA, they were told to answer "I don't know" if uncertain about a question.

The results for HuSimpleQA and SimpleQA are shown in Table 14. The average absolute performance (CGA) of the ten models is 26.65 vs 19.34, highlighting significant room for improvement in LLM performance on both datasets. Regarding average relative performance to humans, the scores are 0.40 vs 0.57, indicating that LLMs perform much worse on HuSimpleQA compared to SimpleQA. Despite better absolute performance on HuSimpleQA, LLMs are significantly weaker in simple factual ability in Hungarian than in English.

F HuProverbRea

F.1 Construction pipeline of HuProverb

The proverbs in HuProverbRea are from 2 separate sources. The first part, 733 traditional Hungarian proverbs, are collected from the website¹⁴, where each proverb is assigned an English or Hungarian explanation. The other 402 proverbs, focusing on abbreviations and Internet slang, are manually collected and explained by native speakers of Hungarian. Inspired by MAPS (Liu et al., 2024c), we adopt a human-in-loop pipeline to generate and refine the context for each Hungarian specific usage, as shown in Figure 12. For each proverb, we first let GPT4 generate a seed context where the proverb is used. Then, we assign it to a Hungarian native

¹⁴<https://mek.oszk.hu/>

model	correct	incorrect	not_attempted	correct given attempted	F-Score
GPT-4o	50.3	40.61	9.09	55.33	52.69
GPT-4o-mini	24.52	74.07	1.42	24.87	24.69
Deepseek-V3	32.71	64.08	3.2	33.8	33.24
QwQ	9.17	52.68	38.15	14.82	11.33
Deepseek-R1	34.58	62.15	3.28	35.75	35.15
Qwen2.5-7B-Instruct	5.29	84.13	10.58	5.92	5.59
Qwen2.5-72B-Instruct	15.05	78.61	6.33	16.07	15.54
Llama-3.1-8B-Instruct	14.9	80.25	4.84	15.66	15.27
Llama-3.1-70B-Instruct	36.36	61.03	2.61	37.34	36.84
o1-mini	16.24	44.19	39.57	26.88	20.25

Table 11: Complete results of HuSimpleQA

Rank	SimpleQA	HuSimpleQA
1	GPT-4o	GPT-4o(-)
2	Deepseek-R1	Llama-3.1-Instruct-70B(↑1)
3	Llama-3.1-Instruct-70B	Deepseek-R1(↓1)
4	Deepseek-V3	Deepseek-V3(-)
5	QwQ	GPT-4o-mini(↑3)
6	Llama-3.1-Instruct-8B	o1-mini(↑3)
7	Qwen2.5-Instruct-72B	Llama-3.1-Instruct-8B(↓1)
8	GPT-4o-mini	Qwen2.5-Instruct-72B(↓1)
9	o1-mini	QwQ(↓4)
10	Qwen2.5-Instruct-7B	Qwen2.5-Instruct-7B(-)

Table 12: Performance rankings on SimpleQA (Wei et al., 2024) and HuSimpleQA

Model	Original HuSimpleQA	Translated to EN
GPT-4o	50.3	35.96 (↓14.34)
GPT-4o-mini	24.52	19.72 (↓4.80)
QwQ	9.17	12.84 (↑3.67)
Deepseek-R1	34.58	32.02 (↓2.56)
Deepseek-V3	32.71	29.08 (↓3.63)
Llama-3.1-70B	36.36	23.82 (↓12.54)
Llama-3.1-8B	14.9	14.08 (↓0.82)
o1-mini	16.24	14.31 (↓1.93)
Qwen2.5-72B	15.05	8.89 (↓6.16)
Qwen2.5-7B	5.29	15.47 (↑10.18)

Table 13: Results of translating HuSimpleQA to English.

speaker to check whether this context is grammatically correct and the use of slang is appropriate. If not, the annotator is required to manually write down a new context for the saying, which will be sent back to another annotator for inspection again. We continue the above procedures until all contexts pass the quality check. It’s worth noting that each option of the 2CQ setting is manually constructed by a human annotator, and only when it passes the double check of two other annotators could it be considered usable. We choose not to involve LLM in this part because designing correct/incorrect options requires deep understanding of sayings, LLM may generate ambiguous options if it does not understand the proverb used in the context, and such pre-provided ambiguous options may negatively influence the creativity of the annotators. Finally, we obtain 1,135 Hungarian proverbs, each equipped

with a context, an English explanation, and two candidate options for question “*What does the speaker mean by the saying?*”.

F.2 More examples of HuProverbRea

The example of HuProverbRea is shown in Figure 28 and Figure 29. The prompt for judging HuProverbRea is shown in Figure 32. The prompt for model inference on HuProverbRea is shown in Figure 31 and 30.

F.3 Comparing the Performance Ranking on HuProverbRea and MAPS

We compare the performance rankings of LLM on HuProverbRea and MAPS (Liu et al., 2024c), as shown in Table 15 and 16.

F.4 Human Performance on HuProverbRea and MAP(en)

We hire human annotators to complete the tasks in the HuProverbRea and MAP(en) datasets in order to assess human performance on these datasets. HuProverbRea was annotated by Hungarian native speakers, while MAP(en) was handled by English native speakers. We used two annotators per dataset, averaging their performance metrics to determine human performance. Annotators were instructed NOT to use external tools like search engines or GPT. The accuracy for HuProverbrea is 80.44%. The accuracy for MAP(en) is 95.31%.

Furthermore, we calculated the LLMs’ relative performance to human, as shown in Table 17.

The average absolute performance of the ten models is 0.8684 vs 0.9350, indicating that LLMs perform worse on HuProverbRea compared to English proverb reasoning, suggesting room for improvement in Hungarian proverb tasks.

For average relative performance to humans, the scores are 1.08 vs 0.98. This indicates that while LLMs slightly underperform compared to humans

Model	HuSimpleQA-abs	HuSimpleQA-rel	SimpleQA-abs	SimpleQA-rel
GPT-4o	55.33	0.85	38.0	1.12
GPT-4o-mini	24.87	0.38	8.7	0.26
Deepseek-V3	33.8	0.52	24.9	0.73
QwQ	14.82	0.23	21.0	0.62
Deepseek_R1	35.75	0.55	30.1	0.89
Qwen2.5-7B-Instruct	5.92	0.09	5.3	0.16
Qwen2.5-72B-Instruct	16.07	0.25	12.2	0.36
Llama-3.1-8B-Instruct	15.66	0.24	12.6	0.37
Llama-3.1-70B-Instruct	37.34	0.58	29.3	0.86
o1-mini	26.88	0.41	11.3	0.33
AVG	26.65	0.40	19.34	0.57

Table 14: The results of HuSimpleQA and SimpleQA, abs is the absolute performance of CGA metric, rel is the performance compared to human.

Rank	HuProverbRea	MAPS (en)	MAPS (bn)	MAPS (id)
1	GPT-4o	GPT-4o (-)	GPT-4o (-)	GPT-4o-mini (↑3)
2	Llama-3.1-Instruct-70B	Llama-3.1-Instruct-70B (-)	Deepseek-V3 (↑1)	Qwen2.5-Instruct-72B (↑4)
3	Deepseek-V3	Qwen2.5-Instruct-72B (↑3)	Deepseek-R1 (↑2)	Llama-3.1-Instruct-70B (↓1)
4	GPT-4o-mini	GPT-4o-mini (-)	Qwen2.5-Instruct-72B (↑2)	GPT-4o (↓3)
5	Deepseek-R1	Deepseek-V3 (↓2)	o1-mini (↑2)	Deepseek-V3 (↓2)
6	Qwen2.5-Instruct-72B	Deepseek-R1 (↓1)	GPT-4o-mini (↓2)	o1-mini (↑1)
7	o1-mini	Qwen2.5-Instruct-7B (↑3)	Llama-3.1-Instruct-70B (↓5)	Deepseek-R1 (↓2)
8	QwQ	Llama-3.1-Instruct-8B (↑1)	Llama-3.1-Instruct-8B (↑1)	Llama-3.1-Instruct-8B (↑1)
9	Llama-3.1-Instruct-8B	o1-mini (↓2)	Qwen2.5-Instruct-7B (↑1)	Qwen2.5-Instruct-7B (↑1)
10	Qwen2.5-Instruct-7B	QwQ (↓2)	QwQ (↓2)	QwQ (↓2)

Table 15: Performance rankings on HuProverbRea and MAPS (Liu et al., 2024c) (part1/2)

on the MAP(en) dataset, they outperform humans on the HuProverbRea dataset, highlighting their strengths in Hungarian language-related knowledge and reasoning.

G HuMatchingFIB and HuStandardFIB

G.1 Construction of HuMatchingFIB and HuStandardFIB

The questions for both HuMatchingFIB and HuStandardFIB are sourced from the Hungarian National Public Education Portal(NKP)¹⁵, a comprehensive platform for cultural funding and support in Hungary. This portal connects artists, cultural organizations, and the public with resources and opportunities to promote Hungarian culture both domestically and internationally. Notably, this website is a government initiative, reflecting the collaborative efforts between the Hungarian government and the European Union, particularly through projects or programs supported by the European Social Fund. After extracting the original questions from the NKP website, we engaged native Hungarian speakers to annotate the data. The annotation process involved manually extracting questions and their corresponding answers¹⁶, classifying the

¹⁵<https://www.nkp.hu/>

¹⁶The questions for HuMatchingFIB and HuStandardFIB on the NKP website are not in plain text but are instead presented in interactive modules, and the answers can only be obtained through additional interactive operations. As a result, the commonly used data cleaning and extraction methods for LLM pre-training datasets are unable to accurately

questions into appropriate categories, and filtering out questions that required additional modalities such as images, tables, audio, or video. This ensured that only purely language-based questions were retained. Through this process, we obtained 278 questions for the HuMatchingFIB task and 93 questions for the HuStandardFIB task, as shown in Table x.

G.2 More examples of HuMatchingFIB and HuStandardFIB

Examples of questions from HuMatchingFIB and HuStandardFIB are provided in Figure 33 and Figure 34. The prompt for model inference on HuMatchingFIB and HuStandardFIB is shown in Figure 36 and Figure 35.

G.3 Metric and Judge of HuMatchingFIB and HuStandardFIB

HuMatchingFIB employs a rule-based evaluation approach, where the assessment is conducted at two levels: the blank level and the question level (as a single question may contain multiple blanks). The evaluation process is analogous to that of multiple-choice questions, and accuracy (acc) is used as

extract these questions and their corresponding answers. Consequently, it can be concluded that the likelihood of these questions being incorporated into the LLM pre-training data in their proper format is minimal, thereby significantly reducing the potential risk of data contamination. This ensures the reasonableness and effectiveness of the test sets for HuMatchingFIB and HuStandardFIB.

Rank	HuProverbRea	MAPS (de)	MAPS (ru)	MAPS (zh)
1	GPT-4o	Llama-3.1-Instruct-70B (↑1)	Deepseek-V3 (↑2)	GPT-4o (-)
2	Llama-3.1-Instruct-70B	GPT-4o-mini (↑2)	o1-mini (↑5)	Qwen2.5-Instruct-72B (↑4)
3	Deepseek-V3	Qwen2.5-Instruct-72B (↑3)	GPT-4o (↓2)	Deepseek-V3 (-)
4	GPT-4o-mini	GPT-4o (↓3)	Qwen2.5-Instruct-72B (↑2)	o1-mini (↑3)
5	Deepseek-R1	Llama-3.1-Instruct-8B (↑4)	GPT-4o-mini (↓1)	Llama-3.1-Instruct-70B (↓3)
6	Qwen2.5-Instruct-72B	o1-mini (↑1)	Deepseek-R1 (↓1)	GPT-4o-mini (↓2)
7	o1-mini	Deepseek-V3 (↓4)	Qwen2.5-Instruct-7B (↑3)	Qwen2.5-Instruct-7B (↑3)
8	QwQ	Qwen2.5-Instruct-7B (↑2)	Llama-3.1-Instruct-70B (↓6)	Llama-3.1-Instruct-8B (↑1)
9	Llama-3.1-Instruct-8B	Deepseek-R1 (↓4)	Llama-3.1-Instruct-8B (-)	Deepseek-R1 (↓4)
10	Qwen2.5-Instruct-7B	QwQ (↓2)	QwQ (↓2)	QwQ (↓2)

Table 16: Performance rankings on HuProverbRea and MAPS (Liu et al., 2024c) (part2/2)

Model	HuProverbRea-abs	HuProverbRea-rel	MAPs(en)-abs	MAPs(en)-rel
Deepseek-V3	92.51	1.15	96.95	1.02
Deepseek_R1	91.72	1.14	98.48	1.03
GPT-4o	95.51	1.19	97.97	1.03
GPT-4o-mini	92.16	1.15	97.97	1.03
Llama-3.1-70B-Instruct	93.83	1.17	98.48	1.03
Llama-3.1-8B-Instruct	73.48	0.91	92.39	0.97
o1-mini	87.67	1.09	92.89	0.97
Qwen2.5-72B-Instruct	90.22	1.12	98.22	1.03
Qwen2.5-7B-Instruct	67.05	0.83	94.16	0.99
QwQ	84.23	1.05	67.51	0.71
AVG	86.84	1.08	93.50	0.98

Table 17: The results of HuProverbRea and MAPs(en), abs is absolute performance, rel is the performance compared to human.

the metric to determine performance. The corresponding formula for blank level accuracy is as follows, where c represents the number of correctly predicted blanks in one question, t represents the number of blanks in one question.

$$\text{Acc}_{\text{blank level}} = \frac{\sum \text{blank}_c}{\sum \text{blank}_t} \quad (3)$$

HuStandardFIB questions are designed with open-ended reference answers to accommodate variations in part of speech and semantics. We employ a many-to-one fuzzy matching mechanism. Fuzzy matching is a technique that calculates the similarity between strings, allowing for flexibility in matching by considering variations such as typos, synonyms, or different word orders. In this context, the model’s answer is compared against a set of possible reference answers (where multiple correct answers may exist for a single question or blank). If the similarity score between the model’s answer and any of the reference answers exceeds a predefined threshold, the answer is considered correct. This approach is particularly suitable for evaluating open-ended questions where exact matches are often infeasible due to the variability in acceptable responses. The annotator information involved in all tasks of OpenHuEval can be found in Appendix I.

H Analyzing LRM’s thinking process on OpenHuEval

H.1 Analyzing LRM’s thinking process on HuSimpleQA

We use GPT-4o to break down the answers into thoughts. This is done in two steps: the first step is to identify expressions that may be a shift in thought (see the prompt in Figure 37), and the second step is to confirm whether it is indeed a shift in thought (set the prompt in Figure 38). Then, We utilized the LLM to evaluate whether each idea would lead to the correct answer, the prompt is shown in Figure 39. We consider a confident score of 2 as the correct thought.

Examples of the thoughts and corresponding correctness can be found in Figure 40 and Figure 41.

H.2 Analyzing LRM’s thinking process on HuMatchingFIB

H.2.1 Splitting LRM’s thinking process into segments and classifying these segments

We first divided the thinking process of LRMs into multiple segments, with each segment categorized as introduction, reasoning, review, or summary. The definitions of these four segment types are provided in Table 18. Typically, each thinking process begins with an introduction segment, includes several reasoning segments and some review segments in the middle, and concludes with a summary segment. The segmentation

and classification were performed using GPT-4o, with the prompt template detailed in Figures 42, 43 and 44. Examples of the segmentation and classification can be found in Figure 47 and Figure 48.

H.2.2 Tagging the reasoning segments along the dimensions

We further tagged the reasoning segments according to the following four dimensions: **(Dim1) correctness**: Are the answers in this reasoning segment correct? **(Dim2) complexity**: In this reasoning segment, does the LRM simply assert the answer, or does it involve more complex reasoning? **(Dim3) scope**: Does this reasoning segment focuses on a single blank, modifies previous blanks, or addresses multiple blanks? **(Dim4) language transfer**: Does the LRM switch languages within this reasoning segment? The details of the tagging can be found in Table 19. Examples of the tagging results can be found in Figure 47 and Figure 48.

I Information of the Annotators

We submitted the annotation task online to a professional data annotation company, which organized annotators to complete the annotation work. In the construction phase of OpenHuEval, the annotations were carried out by professional annotators who are native Hungarian speakers. Table 20 shows the number of annotators and the total time spent on each task. All annotators involved in this project hold a bachelor's or master's degree, with academic backgrounds in fields such as Social Sciences, Translating and Interpreting, English Studies, and IT Engineering. They all possess the ability to distinguish subtle aspects of the Hungarian language and handle Hungarian-specific knowledge effectively.

Segment types	Definition
introduction	typically located at the beginning of the reasoning process; usually consists of the LRM’s brief restatement of the question and the descriptive account of the work it is about to undertake; does not include the actual start of the analysis of the question.
reasoning	typically constitutes the main body of the reasoning process; includes the detailed thinking and reasoning steps undertaken by the LRM to solve the fill-in-the-blank questions.
review	usually occurs after the reasoning process is essentially complete but before the final output. This section typically includes a review of the entire reasoning process and may contain keywords or phrases such as “Overall, ...” or “double check...” .
summary	Summarizes the overall content or provides final conclusions, often using phrases like ‘in conclusion’ or ‘overall’.

Table 18: Definitions of the four categories of segments in the thinking process of LRM on HuMatchingFIB.

Dimension	Description	Tags
Dim 1: Correctness	Are the answers in this reasoning segment correct? Class 4 is used when no conclusion is reached.	Class 1: Completely Incorrect Class 2: Partially Correct Class 3: Completely Correct Class 4: Non Conclusion
Dim 2: Complexity	In this reasoning segment, does the LRM simply assert the answer, or does it involve more complex reasoning?	Class 1: Simple Assertion Class 2: Complex Thought
Dim 3: Scope	Does this reasoning segment focuses on a single blank, modifies previous blanks, or addresses multiple blanks?	Class 1: Only Current Blank Class 2: Modify Previous Blanks Class 3: Current Blank and Consecutive Blank
Dim 4: Language Transfer	Does the LRM switch languages within this reasoning segment (e.g., Hungarian to English)?	Class 1: Contains Language Transfer Class 2: No Language Transfer

Table 19: Tagging dimensions of the reasoning segments in LRM’s thinking process on HuMatchingFIB

Task	# Anotater	Total working hours
HuSimpleQA	14	161.9
HuWildBench	5	55.2
HuProverbRea	15	118.2
HuMatchingFIB and HuStandardFIB	8	84.5

Table 20: Information of the Annotators

Hungarian-specific dimensions: CCC	Count: 365
Question Example: a kérdés az: Mi lesz a jövőben a szocializmus alatt megépül sok panellel? a leírás: Úgy tudom, hogy kb 60 éves életciklusra tervezték őket. Magyarországon (és a környező országokban is) rengeteg ember él bennük. Mi fog történni akkor, ha lakhatatlanná kezdenek válni? Mi lesz azzal a sok emberrel? Meg a panelokkal? Translation to EN: The question is: What will happen in the future to the many panels built under socialism? The description is: I understand they are designed for a life cycle of about 60 years. There are a lot of people living in them in Hungary (and surrounding countries). What will happen if they start to become uninhabitable? What will happen to all those people? And the panels?	Checklist: "Does the response provide an analysis of the current condition and expected lifespan of the panel buildings in Hungary and neighboring countries?", "Does the response address the expected lifespan of panel buildings and their current age?", "Are there any historical or international examples included to illustrate possible outcomes or strategies?", "Does the response consider the economic implications of renovating or replacing panel buildings?", "Does the response include potential government or private sector plans or policies addressing the future of these buildings and their residents?", "Does the answer discuss potential scenarios for when these buildings become uninhabitable?", "Are environmental and urban planning aspects of dealing with aging panel buildings mentioned?", "Is there an explanation of possible solutions or government plans for relocating residents?"
Hungarian-specific dimensions: \mathcal{EP}	Count: 201
Question Example: a kérdés az: A kárpátaljai magyarok Ukrajnában oroszul vagy ukránul tanultak meg a 2000-es évek közepén? a leírás: Mit tanítottak az iskolákban? Mennyire reális az, hogy valakire szinte semmi se ragad a környezetéből? Vannak olyan tömb területek ahol mondjuk egy magyar gyereknek egyáltalán nem kell helyi ukránokkal beszélnie? Egyáltalán a helyi ukránok ukránul beszéltek a 2000-es években? Translation to EN: The question is: Did Hungarians in Transcarpathia learn Russian or Ukrainian in Ukraine in the mid-2000s? The description is: What was taught in schools? How realistic is it that almost nothing sticks to someone from their environment? Are there block areas where, say, a Hungarian child doesn't have to speak to local Ukrainians at all? Did local Ukrainians even speak Ukrainian in the 2000s?	Checklist: "Does the answer provide information on the language predominantly spoken by local Ukrainians in Transcarpathia in the 2000s?", "Does the response discuss the social and linguistic dynamics in areas with significant Hungarian populations, including interactions with local Ukrainians?", "Does the response clearly explain the educational policies and language of instruction in schools for Hungarians in Transcarpathia during the mid-2000s?", "Does the response accurately describe the language of instruction in Transcarpathian Hungarian schools in the mid-2000s?", "Does the response consider the historical and political context of language policies in Ukraine during this period?", "Does the response provide insight into whether local Ukrainians predominantly spoke Ukrainian during the 2000s?", "Does the response offer a balanced view of cultural and linguistic integration in Transcarpathia during the specified period?", "Does the answer address the likelihood of a Hungarian child not acquiring any local language skills from their environment?", "Does the response discuss the existence of predominantly Hungarian areas where interaction with local Ukrainians might be limited?"
Hungarian-specific dimensions: \mathcal{PPL}	Count: 299
Question Example: a kérdés az: Mi történt azzal, aki az 50-es években a felhívás ellenére sem jegyzett "önként" békekölcsönt? Érhette ezért retorzió az embert? a leírás: Persze nyilván volt, amilyen "bolondos" időket jártak nálunk akkortájt. Biztos kikiáltották reakciónak vagy fasisztának, meg a "népi demokrácia" ellenségének. Translation to EN: The question is: What happened to the man who did not "voluntarily" subscribe to a peace charter in the 1950s, despite the call? Could he have been retaliated against for this? The description is: He must have been branded a reactionary or a fascist or an enemy of 'people's democracy'.	Checklist: "Does the answer address the political labels mentioned in the description (e.g., 'reactionary', 'fascist', 'enemy of people's democracy')?", "Does the response differentiate between official consequences and social/societal repercussions for not subscribing to the peace loan?", "Does the response address potential consequences for individuals who did not subscribe to the peace loan, with references to historical examples or documentation?", "Does the response provide a balanced view, considering both potential punitive measures and any instances of leniency or exceptions, if applicable?", "Is there a clear explanation of what 'békekölcsön' (peace loan) was and its significance during that time period?", "Does the response accurately describe the historical context of the 1950s in Hungary?", "Is there an analysis of the societal and governmental attitudes toward dissenters in Hungary during the 1950s, including any possible labels or accusations they might have faced?", "Does the response provide specific examples of potential retaliations against those who didn't subscribe to the peace loan?"
Hungarian-specific dimensions: \mathcal{BF}	Count: 289
Question Example: a kérdés az: Meddig tartható fent Magyarország negatív külkereskedelmi mérlege? a leírás: Nem a háború óta, hanem már 2021 nyarától folyamatosan negatív az ország külkereskedelmi mérlege. Júliusban és augusztusban összesen több, mint 1000 milliárd forintnyi mínusz keletkezett. Persze a többi hónap nem volt ennyire szörnyű, de ez csak erre az évre már több, mint 2000 milliárd forintnyi mínusz. Változatlan devizaimport mellett a mérséklődött energiaárakkal is több, mint 1000 milliárdos negatív mérleg hozható össze 2023-ban. Meddig lehet ezt tovább folytatni? Meddig elég a devizataralék a hiány pótlására? Translation to EN: The question is: How long can Hungary maintain a negative trade balance? The description is: In July and August there was a total deficit of more than HUF 1000 billion. Of course, the other months were not so bad, but for this year alone it is already more than HUF 2000 billion in deficit. Even with unchanged foreign exchange imports and moderating energy prices, a negative balance of more than 1,000 billion in 2023 could be created. How long can this go on? How long will foreign exchange reserves be enough to cover the deficit?	Checklist: "Does the response analyze Hungary's current foreign exchange reserves and their sufficiency in covering the trade deficit?", "Is there an exploration of historical trends and comparisons to similar situations in other countries to provide context?", "Is the impact of energy prices on the trade balance accurately assessed in the response?", "Does the response offer a clear and supported prediction or timeframe for how long Hungary can sustain its negative trade balance?", "Is there an analysis of the factors affecting Hungary's foreign exchange reserves and their ability to cover the deficit?", "Does the answer provide a clear timeline or projection for how long the negative balance can be sustained?", "Are there comparisons made to similar situations in other countries or historical precedents in Hungary?", "Does the response accurately explain the current state of Hungary's foreign trade balance?"

Table 21: Examples of HuWildBench. The blue font is the English translation of the original OpenHuEval examples, used for visualization.

Hungarian-specific dimensions	Count	Question-Answer Pairs	translation to EN
L	11	Question1: Mit jelent a Kara török eredetű régi magyar személynév? Answer1: fekete Question2: Melyik régi magyar név a Pantaleon megfelelője? Answer2: Pentele	Question1: What does the old Hungarian personal name Kara of Turkish origin mean? Answer1: black Question2: Which old Hungarian name is the equivalent of Pantaleon? Answer2: Pentele
H	140	Question1: Melyik király nevezte ki Szapolyai Imrét szepesi örökletes főispánná 1465-ben? Answer1: Mátyás király Question2: Melyik várost foglalta el Báthory Gábor 1610. december 11-én? Answer2: Szeben	Question1: Which king appointed Imre Szapolyai as the hereditary ispán of Szepes in 1465? Answer1: King Matthias Question2: Which city was captured by Gabriel Báthory on December 11, 1610? Answer2: Sibiu
LCC	257	Question1: Melyik magyar film nyerte el a FIPRESCI-díjat az 1983-as Cannes-i Nemzetközi Filmfesztiválon? Answer1: Szerencsés Dániel Question2: Melyik legendára épít az 'Eredet / Origins' táncjáték? Answer2: Csodaszarvas-legendára	Question1: Which Hungarian film won the FIPRESCI Prize at the 1983 Cannes International Film Festival? Answer1: Lucky Daniel Question2: Which legend is the 'Origin / Origins' dance play based on? Answer2: Legend of the Miraculous Deer
EP	81	Question1: Melyik városban alapították a Gandhi Gimnáziumot 1994-ben? Answer1: Pécsen Question2: Melyik évben alapította a Magyar Tudományos Akadémia az Acta Juridica Hungarica folyóiratot? Answer2: 1959	Question1: In which city was the Gandhi High School founded in 1994? Answer1: Pécs Question2: In which year did the Hungarian Academy of Sciences establish the journal Acta Juridica Hungarica? Answer2: 1959
GP	165	Question1: Melyik magyar vármegyében található Nemesmedves? Answer1: Vas vármegyében Question2: Mi a neve Magyarország legmagasabb fekvő csillagvizsgálójának, amely a Pizskés-tetőn található? Answer2: Pizskéstetői Observatórium	Question1: In which Hungarian county is Nemesmedves located? Answer1: Vas county Question2: What is the name of Hungary's highest observatory, located on Pizskés Peak? Answer2: Pizskés Peak Observatory
F	409	Question1: Nádasdy Kálmán hányszor kapott Kossuth-díjat élete során? Answer1: Háromszor Question2: Balogh József melyik magyar városban született 1946. április 15-én? Answer2: Nagykanizsán	Question1: How many times did Kálmán Nádasdy receive the Kossuth Prize during his lifetime? Answer1: Three times Question2: In which Hungarian city was József Balogh born on April 15, 1946? Answer2: Nagykanizsa
PPL	186	Question1: Melyik szervezet jogkörét vette át a Nép gazdasági Tanács 1949. június 11-én? Answer1: Gazdasági Főtanács Question2: Melyik törvénycikk rendelkezett 1878-ban Magyarországon a réz-váltópénz szaporításáról? Answer2: 1878. évi VI. törvénycikk	Question1: Which organization's authority was taken over by the National Economic Council on June 11, 1949? Answer1: Supreme Economic Council Question2: Which statute regulated the increase of copper coinage in Hungary in 1878? Answer2: Act VI of 1878
BF	44	Question1: Milyen néven működött az ÉVITERV 1954-től az 1980-as évek elejéig? Answer1: ÉM Szerelőipari Tervező Vállalat Question2: Melyik cég gyártotta a Puli autótípust a gyártás kezdeti időszakában? Answer2: HÓDGÉP	Question1: Under what name did ÉVITERV operate from 1954 to the early 1980s? Answer1: ÉM Installation Industry Design Company Question2: Which company manufactured the Puli car model in the early production period? Answer2: HÓDGÉP

Table 22: Examples of HuSimpleQA. The rightmost column is the English translation of the original OpenHuEval examples, used for visualization.

""Given the following question, identify whether it has a characteristic related to Hungary. A question is considered to have a Hungarian characteristic if it meets any of the following criteria:

Hungary-Specific Context: The question itself directly references or relates to Hungary. For example, "What is the capital of Hungary?" clearly has a Hungarian characteristic.

Hungary-Specific Answer: The question might not directly reference Hungary, but the answer would vary depending on the country, particularly Hungary. For example, "What is the minimum wage according to labor laws?" The answer would depend on Hungary's laws and practices.

Hungary-Specific Context and Answer: Both the question and the likely answer have strong connections to Hungary. For example, "How do you view Hungary's 2024 foreign policy?" is likely to have both the question and the answer centered on Hungary.

Any Other Model-Identified Hungarian Characteristic: If the model identifies a Hungarian characteristic based on context, culture, or any other relevant factors.

There are several special rules to follow:

The language of the question should not be used as an evidence.

For a question to which the Hungarian answer is not significantly different from the answer of the rest of the world, the question is not considered having a Hungarian characteristic.

If a question only mentions a Hungarian-related term, such as the Hungarian currency, the forint, or a certain place in Hungary, but the question itself is not more related to Hungary's cultural, social, political, economic, military, life, etc., the question is not considered having a Hungarian characteristic.

If the answer to a question is open-ended, for example, "Will you buy a flower for your mom on Mothers' Day?", the question is not considered having a Hungarian characteristic.

Please respond strictly in JSON format. Do not include any additional text outside the JSON structure.

```
{
  "Question": "[The original question]",
  "HasHungarianCharacteristic": "yes/no",
  "Reason": "[Explanation for why this question was classified as having a Hungarian characteristic]"
  "Score": "[This score is used to evaluate how relevant this issue is to Hungary, with 0 being the lowest and 10 being the highest.]
}
```

The question is: <question>.

Note that each question is composed of a question itself and a question description.""

Figure 13: Prompt template for automatic filtering of user questions related to Hungarian specifics in the construction of HuWildBench.

```

"""# Instruction
You are an expert responsible for evaluating the capabilities of a language model in handling
questions related to the Hungarian language and context.
The questions are sourced from the Gyakori kérdések website, and the objective is to assess the
model's performance by selecting the best questions based on a set of criteria.
You will be given five Hungarian questions.
Based on the 9 criteria listed below, select the two best questions.

## Evaluation Criteria
<|begin of evaluation|>
1. Linguistic complexity: Does the question contain complex syntactic structures and rich
vocabulary, testing the model's ability to process complex language?
2. Hungarian-specific relevance: Is the question highly relevant to Hungarian culture, society,
history, or daily life, testing the model's understanding of Hungary-specific context?
3. Requirement for common knowledge and factual accuracy: Does the question require knowledge
of Hungarian common sense or factual information, allowing for the evaluation of the model's
knowledge base and accuracy?
4. Context dependency: Does the question require the model to understand or infer from the
context, testing the model's ability to use prior or surrounding information?
5. Answer diversity: Does the question allow for multiple reasonable answers, testing the model's
creativity and ability to generate diverse responses?
6. Ambiguity: Does the question contain ambiguity or multiple meanings, testing the model's ability
to handle uncertain or vague information?
7. Reasoning requirement: Does the question require logical reasoning or causal inference, testing
the model's ability to analyze and reason through complex information?
8. Social and ethical considerations: Does the question involve social, ethical, or moral issues,
testing the model's ability to generate responses that align with ethical standards?
9. Format diversity: Does the question come in a unique format (e.g., multiple choice, open-ended,
narrative, etc.), testing the model's ability to handle different types of question formats?
10. Breadth of knowledge: Does the question cover a broad range of knowledge areas (e.g., science,
arts, technology), testing the model's general knowledge across various domains.
<|end of evaluation|>

## Questions
<|begin of questions|>
1. <question0>.
2. <question1>.
3. <question2>.
4. <question3>.
5. <question4>.
<|end of questions|>
Note that each question is composed of a question itself and a question description.

## Output format
Your output should be in JSON format as follows:
Please respond strictly in JSON format. Do not include any additional text outside the JSON
structure.
{
  "question_indices": [a list of the indices of the two best question in int type],
}"""

```

Figure 14: Prompt template for automatic filtering of high-quality question in the construction of HuWildBench.

""You are a model designed to assist in evaluating responses to questions. You will receive a question about Hungary, and your task is to provide a list of 3–5 evaluation criteria. Each item in the list should be a distinct angle for assessing whether the response to the question meets the required standard. For example, if the question is: "Is a monthly income of \$1000 sufficient to cover normal living expenses in the capital city of Hungary?", the list could include criteria such as:

- 1.Does the response comprehensively outline all relevant living expenses in Budapest?
- 2.Are the amounts mentioned for each expense aligned with objective facts?
- 3.Does the response provide an overall conclusion on whether \$1000 is enough for living expenses?

Each criterion should assess a different aspect of the response, ensuring no overlap in evaluation angles. Please respond strictly in JSON format. Do not include any additional text outside the JSON structure.

```
{
  "Checklist": "[The evaluation criteria list]"
}
```

The question is: <question>.
Note that each question is composed of a question itself and a question description.""

Figure 15: Prompt template for constructing the checklist in the construction of HuWildBench.

```

"""# Instruction
You are an expert evaluator. Your task is to evaluate the quality of the responses generated by AI models.
We will provide you with the user query and an AI-generated responses.
You should first read the user query and the conversation history carefully for analyzing the task,
and then evaluate the quality of the responses based on and rules provided below.

# Conversation between User and AI

## History
<|begin_of_history|> {history} <|end_of_history|>

## Current User Query
<|begin_of_query|> {user_query} <|end_of_query|>

## AI Response
<|begin_of_response|> {prediction} <|end_of_response|>

# Evaluation

## Checklist
<|begin_of_checklist|> {checklist} <|end_of_checklist|>
Please use this checklist to guide your evaluation, but do not limit your assessment to the checklist.

## Rules
You should compare the above response based on your analysis of the user queries and the conversation history.
You should first write down your analysis and the checklist that you used for the evaluation, and then provide your assessment according to the checklist.
The scores are in the range of 1~10, where 1 means the response is very poor and 10 means the response is perfect.
Here are more detailed criteria for the scores:

– Score 1~2: The response is very poor and does not make sense at all.
– Score 3~4: The response is poor and does help user solve the problem in a meaningful way.
– Score 5~6: The response is fair but has some issues (e.g., factual errors, hallucinations, missing key information).
– Score 7~8: The response is good enough but could be improved in some ways.
– Score 9~10: The response is perfect and provides helpful information that can help user solve the problem.

## Output Format
First, please output your analysis for the model response, and then summarize your assessment to two aspects: "strengths" and "weaknesses"; Finally, please write down your rating for the assessment.
Please provide your evaluation results in the following json format by filling in the placeholders in []:
```{"strengths": "[analysis for the strengths of the response]", "weaknesses": "[analysis for the weaknesses of the response]", "score": "[1~10]"}```

```

Figure 16: Prompt template for LLM as judge on HuWildBench (WBScore).

""""

- Role: Expert in Hungarian Culture and Data Classification
- Background: You are tasked with classifying data that is deeply related to Hungarian-specific content. This data may involve Hungarian history, culture, art, folklore, language, traditions, tourism, and more. Your expertise is critical to ensuring the classification is precise and adheres strictly to Hungarian cultural relevance.
- Goal: Analyze the detailed description in the input data and categorize it into one of the following nine predefined categories:
  1. Language: Content related to Hungarian language, including proverbs, idioms, or linguistic knowledge.
  2. History: strictly for content describing specific historical events or developments in Hungary. Examples include wars, revolutions, significant treaties, or influential periods of political or societal change.
  3. Life, Culture, and Customs: Hungarian religion, etiquette, cultural practices, holidays, and daily life (including tourism).
  4. Education and Profession: Information on Hungary's education system or associated occupations.
  5. Geography and Place: Hungarian geography, city locations, landmarks, and travel-related content.
  6. Figure: Notable Hungarian individuals and their achievements.
  7. Politics, Policy, and Law: Hungarian political systems, policies, or legal regulations.
  8. Business and Finance: Hungarian economy, business practices, or financial systems.
  9. Others: Content not relevant to Hungarian culture or not fitting into the above categories.

Constraints:

- Cultural Accuracy: Your classification must be based on an in-depth understanding of Hungarian culture and the context provided in the input. Avoid assumptions or generic classifications that lack cultural alignment.
- Systematic Approach: Follow a logical and consistent process to ensure every input is matched to the most relevant category. If the content cannot be clearly classified into one category, opt for "Others"
- Specificity: Focus on how the content relates explicitly to Hungary. Avoid overgeneralizing or assigning tags that are only loosely connected to the data.

Please classify the following data according to the above requirements and example:  
<input\_question>

Please respond strictly in JSON format. Do not include any additional text outside the JSON structure.

```
{
 "cn_specific_label": "[predicted label]"
}
```

""""

Figure 17: Prompt template for selecting Hungarian specific Wikipedia entries in the construction of HuSimpleQA.

```

""""
As a general knowledge expert, please judge the knowledge value of the material and
extract key information from the following descriptive materials. The requirements
are as follows:
1. The extracted content is the most critical information of the text description
subject. Please extract the core content of the description text in a targeted manner.
2. Please ensure that the extracted information is accurate and unambiguous.
3. The extracted key information is in <language>.
4. The key information extracted should be related to the title corresponding to the
material.

[Contextual information]
{
 "title": "704-es busz",
 "content": "# 704-es busz\n\nA 704-es jelzésű elővárosi autóbusz Százhalombatta,
DE-Zrt. 2 sz. kapu és Martonvásár, vasútállomás között közlekedik. A járatot a
Volánbusz üzemelteti.\n\n## Megállóhelyei",
}
[Extract key information]
{
 "key_info": "1. **Üzemeltetési útvonal**: A 704-es autóbuszjárat Százhalombatta
DE-Zrt. 2-es kapuja és Martonvásár vasútállomása között közlekedik.\n 2. **Üzemeltető
cég**: A járatot a Volánbusz üzemelteti."
}

[Contextual information]
{
 "title": <title>,
 "content": <content>
}

Please respond strictly in JSON format. Do not include any additional text outside
the JSON structure:
{
 "key_info": "[the key information extracted from the given Contextual material]"
}
""""

```

Figure 18: Prompt template for extracting key information from the wikipedia entries in the construction of HuSimpleQA.



```

{
 "title": "Bodor Anikó",
 "content": "# Bodor Anikó\n\nBodor Anikó (Zenta, 1941. június 15. – Zenta, 2010. július 9.) vajdasági népzenekutató, tanár.\n\n## Életrajz\n\n1960-ban a zentai gimnáziumban érettségizett, majd 1969-ig jogi tanulmányokat folytatott Újvidéken és Zágrábban. 1966-1972 között zenetudományi, művészettörténeti tanulmányokat folytatott Stockholmban és Uppsalában, ahol zenetudományi diplomát szerzett. Tanulmányait 1976-1980 között a belgrádi Zeneakadémia etnomuzikológia szakán folytatta, ahol 1984-ben a népzene tudományok magisztere lett. 1972 és 1973 között az Újvidéki Rádió és Televízió munkatársa, 1975 és 1995 között a zentai alsófokú zeneiskola tanára, 1995-től a Zentai Városi Múzeum munkatársa.\n\nMunkásságát bizonyítja a több mint félszáz tanulmány, nagyobb cikk és népzenei kiadvány (könyv, kotta, lemez). Ezek közül a legnagyobb horderejű az öt könyvre tervezett Vajdasági magyar népdalok című sorozata volt, amelyből eddig négy kötet jelent meg. Szerkesztésében készült el a Daloló vajdasági fiatalok és a Vajdasági élő magyar népzene című népzenei lemezsorozat. A Délvidéki Népzenei Archívum létrehozója és gondozója volt.\n\n## Főbb művei\n\n- Hallottatok-e hírét? (1977)\n- Tiszából a Dunába folyik a víz (1978)\n- sajtó alá rendezte a Gombos és Doroszló népzeneje (1982), Az aldunai székelyek népdalai (1984) és A drávaszögi magyarok dalai (1989) c. könyveket.\n- A szlavóniai szigetmagyarság népdalai I. (Kiss Lajossal, 1990)\n- Vajdasági magyar népdalok I. (1997)\n- Vajdasági magyar népdalok II. (1999)\n- Vajdasági magyar népdalok III. (2003)\n- Vajdasági magyar népdalok IV. (2008)",
 "key_info": "1. **Születési és halálozási adatok**: Bodor Anikó 1941. június 15-én született Zentán és 2010. július 9-én hunyt el Zentán.\n2. **Szakmai tevékenység**: Vajdasági népzenekutató és tanár.\n3. **Tanulmányok**: Jogot tanult Újvidéken és Zágrábban, zenetudományt és művészettörténetet Stockholmban és Uppsalában, etnomuzikológiát a belgrádi Zeneakadémián.\n4. **Munkásság**: Több mint félszáz tanulmány és népzenei kiadvány szerzője, a Vajdasági magyar népdalok című sorozatból négy kötet jelent meg.\n5. **Fontos művek**: 'Hallottatok-e hírét?' (1977), 'Tiszából a Dunába folyik a víz' (1978), 'Vajdasági magyar népdalok' sorozat (1997, 1999, 2003, 2008)."
```

Figure 19: Example of the extracted key information from the Hungarian wikipedia entries. (HuSimpleQA)

```

"""
As a general knowledge expert, please generate 1 to 3 factual open-ended questions with their corresponding answers, based on
the specified knowledge material. Ensure the questions meet the following criteria:
1. Content Relevance:
 • The question content should be related to the title corresponding to the key information. Only objective knowledge should
 be tested, such as the life story of important historical figures, information about important events, leaders of
 important events, or important attributes of certain objects and concepts. Do not test irrelevant information.
 • Minimize questions that are based solely on time and place, and instead, focus on unique and detailed aspects of the
 subject matter.
2. Clarity and Scope:
 • Each question is an independent and unambiguous question and can be answered independently without the help of other
 materials.
 • The question stem must specify the scope of the answer. Avoid broad or open-ended questions. Ensure answer is clear and
 objective, avoiding subjective speculation.
 • For example, instead of asking 'hol találkozott Barack és Michelle Obama' (for which could have multiple answers 'Chicago'
 or 'a Sidley & Austin ügyvédi iroda'), questions had to specify 'melyik városban' or 'melyik cégnél'. Another common
 example is that instead of asking simply 'mikor' or 'melyik időpontban' (meaning "when" or "what time"), the question
 should ask 'melyik évben' or 'melyik napon' (meaning "which year" or "which day").
 • Answers should be brief, without additional explanations or redundancy. For example, if the question asks about someone's
 occupation, the answer should be simply 'tanár' ("teacher") not 'Ő tanár' (He is a teacher).
3. Consistency over Time:
 • Ensure that reference answers do not change over time. Try to avoid generating content that will change due to the
 progress of historical research, entertainment works, construction and updates of transportation roads, etc.
 • For example, instead of broadly asking "ki Meredith párja a Grey's Anatomy-ban", which could change as new seasons are
 produced, questions about TV shows, movies, video games, and sports typically require specifying a point in time (e.g.,
 "ki Meredith párja a Grey's Anatomy 13. évadában").
4. Question type: The questions should be open-ended, with a clear problem description and answer.
5. Moderate difficulty: Ensure the questions have appropriate readability and difficulty, allowing for clear differentiation
 between correct answers while maintaining accuracy.
6. Distinct Knowledge Points:
 • For each material, generate 1 to 3 questions and answers, ensuring the knowledge points being tested are distinct and do
 not overlap. Each question should offer a unique perspective and related answer.
 • All questions should be related to Hungarian-specific knowledge, reflecting aspects of Hungarian history, culture,
 geography, economy, figure, education or other uniquely Hungarian topics.
7. Language: The questions and answers are in Hungarian.

Example1:
[Input title and key information]:
{
 "title": "2004-es Formula-1 magyar nagydíj",
 "key_info": "1. A 2004-es Formula-1 magyar nagydíj a 2004-es világbajnokság tizenharmadik futama volt, amelyet 2004.
 augusztus 15-én rendeztek meg a Hungaroringen. Ez volt a 19. Formula-1-es futam Magyarországon.\n2. Michael Schumacher..."
}
[questions generated based on the information]
{
 "1":{
 "question": "A 2024-es Forma-1-es Magyar Nagydíj hányadik Forma-1-es versenylvolt Magyarországon?",
 "answer": "Michael Schumacher"
 },
 "2":{
 "question": "Milyen büntetést kapott Felipe Massa a 2004-es Formula-1-es Magyar Nagydíjon a motorcsere miatt?",
 "answer": "Tízhelyes rajtbüntetés",
 },
 "3":{
 "question": "A 2004-es Forma-1-es Magyar Nagydíj a 2004-es világbajnokság melyik futama volt?",
 "Answers": "tizenharmadik futam"
 }
}

Example2:
...

Please strictly follow the above requirements to generate the questions and answers in Json format, Do not add extra
irrelevant format or content.
[Input title and key information]:
{
 "title": <title>,
 "key_info": <key_info>
}
[questions generated based on the information]
{
 "1":{
 "question": "str",
 "answer": "str"
 },
 "2":{
 "question": "str",
 "answer": "str"
 },
 ...
}
"""

```

Figure 20: Prompt template for generating Hungarian question-answer pairs in the construction of HuSimpleQA.

```

""""
- Role: Hungarian Featured Content Identification Expert
- Background: Your role is to classify given questions and answers, determining whether they are related to Hungary and identifying their specific category.
- Goals:
 1. Determine whether given question and answer is related to Hungary.
 2. If related, identify the category it belongs to, and assign the appropriate label from the predefined list of categories.
- Classification Categories:
 1. Language: Content related to the Hungarian language, including proverbs, idioms, or linguistic knowledge.
 2. History: Content strictly describing specific historical events or developments in Hungary, such as wars, revolutions, significant treaties, or influential political or societal periods.
 3. Life, Culture, and Customs: Information about Hungarian religion, etiquette, cultural practices, holidays, daily life, and tourism.
 4. Education and Profession: Details about Hungary's education system or associated professions.
 5. Geography and Place: Content about Hungary's geography, cities, landmarks, or travel-related topics.
 6. Figure: Information about notable Hungarian individuals and their achievements.
 7. Politics, Policy, and Law: Information about Hungary's political systems, policies, or legal regulations.
 8. Business and Finance: Content related to Hungary's economy, business practices, or financial systems.
 9. Others: Content unrelated to Hungarian culture or not fitting into the above categories.
- Constraints:
 - Relevance: Only classify the content of question and answer related to Hungary. If the content is in Hungarian but unrelated to Hungary or is generic, classify it as unrelated.
 - Strict adherence to categories: Ensure consistent and accurate classification according to the nine dimensions.
 - Unclear content: For texts that cannot be clearly categorized, assign them to the "Others" category.
 - Each question can only have one category label.

Example1:
Input:
{
 "question": "Milyen posztumusz díjat kapott Fehér Sándor hegedűművész 2013. január 10-én?",
 "answer": "Magyar Civil Becsületrend",
}
Output:
{
 "hu_related": "True",
 "question_specific_label": "Figure"
}

Example2:
...

Please strictly follow the above format classify given questions and answers, do not add extra irrelevant format or content.
Input:
{
 "question": <question>,
 "answer": <answer>
}

Please respond strictly in JSON format. Do not include any additional text outside the JSON structure.
Output:
{
 "hu_related": [If question and answer are related to Hungarian characteristics, enter "True". Otherwise, enter "False"]
 "question_specific_label": [Predicated label should be chosen from the above nine categories. If there is an exception or it cannot be judged, set the string to an empty string.]
}
""""

```

Figure 21: Prompt template for categorizing the generated question-answer pairs in the construction of HuSimpleQA.

```

"""
- Role: Hungarian Content Review Expert
- Background Information:
You need to determine whether an open-ended question and its answer are relevant to Hungarian characteristics and align with the provided background information. The given information includes various aspects about Hungary, including history, culture, language, geography, people, law, economy and more. If the question and answer involve content related to Hungary, you need to ensure the content is accurate and consistent with the background information.

- Task:
1. Determine relevance to Hungarian characteristics:
 - Ensure that the question and answer relate to the given title and key information, particularly with regard to Hungarian history, life, culture, customs, people, geography, politics, economy, education, etc. If the content does not align with Hungarian characteristics, it is considered irrelevant.
 - If the question is in Hungarian but unrelated to Hungarian characteristics, it is also considered irrelevant.
2. Ensure consistency with background information:
 - Verify that the question and answer are not only relevant to Hungary but also fit the background information provided. For example, Hungarian historical events should match the correct time and facts, and cultural references should align with actual Hungarian culture.
 - Ensure that the question description and answer are consistent with the information provided in the materials, without any deviation or omission
3. Appropriate Difficulty:
 - The question should not be overly simple, and the answer should not be directly obvious from the question itself.

Example1:
Input: # This question is irrelevant to Hungary characteristics.
{
 "title": "Ipari Termékosztályozás",
 "key_info": "1. Definíció: Az Ipari Termékosztályozás (ITO) egy hierarchikus statisztikai osztályozás, amely az Eurostat PRODCOM jegyzékének hazai sajátosságokkal kiegészített változata, és ipari termékek és szolgáltatások gazdasági megfigyelésére használják.\n2. Struktúra: Az ITO kód 12 számjegyből áll, amelyek a TEAOR'08, TESZOR, PRODCOM, és KSH által képzett kódok kombinációjából állnak.\n3. Történet: Az ITO 2008. január 1-jén lépett hatályba, elődje a Belföldi Termékosztályozás (BTO) volt, amely 2007. december 31-ig volt érvényben.\n4. Jogszabályok: Az ITO-ra vonatkozó jogszabályok közé tartozik a 6/2018. (III.12.) MvM rendelet, a Bizottság 2017/2119 rendelete, a Bizottság 912/2004/EK rendelete, a Bizottság 1209/2014/EU rendelete, az Európai Parlament és a Tanács 451/2008/EK rendelete, az Európai Parlament és a Tanács 1893/2006/EK rendelete, a 16/2011. (V. 10.) KIM rendelet, a Bizottság 2017/2119 rendelete, és a Tanács 3924/91/EGK rendelete."
 "question": "Melyik évben lépett hatályba az Ipari Termékosztályozás (ITO)?",
 "answers": "2008"
}
Output:
{
 "question_hu_relevant" : "fail",
 "answer_hu_correct": "pass"
}

Example2:
...

Please strictly follow the above format to judge the quality of question and answer, do not add extra irrelevant format or content.
Input:
{
 "title": <title>,
 "key_info": <key_info>,
 "question": <question>,
 "answer": <answer>,
}
Please respond strictly in JSON format. Do not include any additional text outside the JSON structure:
Output:
{
 "question_hu_relevant": [If the question and answer is relevant to Hungarian characteristics, enter "pass". Otherwise, enter "fail"],
 "answer_hu_correct": [If the question and answer is consistent to key information, enter "pass". Otherwise, enter "fail"]
}
"""

```

Figure 22: Prompt template for evaluating the relevance and correctness of question-answer pairs in the construction of HuSimpleQA.



```

"""
- Role: Hungarian Content Review Expert
- Background Information: You are responsible for evaluating whether an open-ended question and its
corresponding answer meet the following standards:
1. Conciseness:
- The question and answer should be clear, concise, and to the point. Avoid unnecessary details
or redundant descriptions.
- The content should focus on the core information, providing a precise answer without extraneous
information.
- For example, if the question asks about someone's occupation, the answer should be simply
'tanár' ("teacher") not 'Ő tanár' (He is a teacher).
2. Single Question:
The question should contain only one query. Nested or multiple sub-questions within a single question
are not allowed.
3. Specificity:
- The question must be precise and targeted. Avoid broad open-ended questions like 'Miért' or
'Milyen hatása volt?' which require analysis or subjective answers.
- Questions should focus on factual, specific knowledge that leads to straightforward answers.
4. Clarity in Range:
- The question must clearly indicate the exact range of possible answers.
- For time-related questions, **do not use vague terms like 'Mikor' ('When'/'what time') in
question**. Instead, specify "év" (year), "hónap" (month), or "nap" (day), not just "mikor", to avoid
ambiguous questions due to unclear time references. Ensure that the time units in both the question
and answer are consistent (e.g., 'melyik nap', 'melyik év', 'melyik hónap')(meaning "which year" or
"which day").
- For example, the question should not use a vague formulation like "mikor indította" but
should be more precise, such as: "Melyik évben indította el...".
- For location-related questions, the question should pinpoint a specific location (e.g.,
'város', 'ország', 'régió')(meaning city, country, or region) rather than a vague 'hol' (where').
5. Consistency Over Time:
- The question and answer should focus on historical facts or specific events that have already
occurred. Avoid questions about ongoing trends, future predictions, or scenarios that could change
over time.
- Answers should remain consistent regardless of future developments.
- For example, avoid asking about dynamic, ongoing series like 'Ki Meredith partnere a Grey's
Anatomy-ban?' Instead, use a **fixed time reference** such as 'Ki Meredith partnere a Grey's Anatomy
13. évadában?'

Example1:
Input:
{
 "question": "Melyik törvénycikk szabályozta 1886-ban Magyarországon a trachoma terjedésének
meggátlását?",
 "answer": "1886. évi V. törvénycikk",
}
Outputs:
{
 "question_hu_qualified" : "pass",
 "question_hu_time_consistency": "pass"
}

Example2:
...

Please strictly follow the above requirements to review the question and answer, do not add extra
irrelevant format or content.
Input:
{
 "question": <question>,
 "answer": <answer>,
}

Please respond strictly in JSON format. Do not include any additional text outside the JSON structure:
Output
{
 "question_hu_qualified": [If the question or answer are concise, specific and only have one
question, enter "pass". Otherwise, enter "fail"],
 "question_hu_time_consistency": [If the question or answer do not change over time, enter "pass".
Otherwise, enter "fail"]
}
"""

```

Figure 23: Prompt template for evaluating the precision and consistency of question-answer pairs in the construction of HuSimpleQA.

```

"""
- Role: Hungarian factual Knowledge open-ended questions Review Expert
- Background: Hungarian General Knowledge open-ended question involves many aspects of Hungarian history, culture, geography, politics, people, etc., and a professional review expert is needed to ensure the accuracy of the answers.
- Goals: According to the given questions and reference answers, accurately judge whether the answers entered by the user are correct, and give the evaluation results: "CORRECT", "INCORRECT".
- Evaluation criteria for the given answer to be correct:
 • Predict answer fully contain the important information in the gold target.
 • Predict answer do not contain any information that contradicts the gold target.
 • Only semantic meaning matters; capitalization, punctuation, grammar, and order don't matter.
 • Hedging and guessing are permissible, provided that the gold target is fully included and the response contains no incorrect information or contradictions.
 • The answer must be consistent with the scope of the question. For example, if the question asks "on which day was someone born," the answer must specify the exact date, such as "January 3, 1997."

- Workflow:
1. Receive questions, reference answers, and user answers.
2. Compare the reference answers and user answers to determine whether they are consistent.
3. Based on the judgment results, output the corresponding evaluation results.
- Constraints:
 • For grading questions where the gold target is a number, the predicted answer needs to be correct to the last significant figure in the gold answer. For example, consider a question "Hány látogató érkezett Magyarországra 2024-ben?" with gold target "440k".
 • Predicted answers "440k", "444k", and "435k" are all CORRECT.
 • Predicted answers "400k" and "413k" are INCORRECT.
 • The gold target may contain more information than the question. In such cases, the predicted answer only needs to contain the information that is in the question.
 • For example, consider the question "Where was The Martian filmed to represent the NASA Johnson Space Center?" with the gold target "Budapesti Bálna (HU)". Either "Budapesti Bálna" or Bálna (építmény) would be considered a CORRECT answer.
 • Do not punish predicted answers if they omit information that would be clearly inferred from the question.
 • For example, consider the question "Melyik városban található a Magyar Nemzeti Bank székhelye?" with the gold target "Budapest, Magyarország". The predicted answer "Budapest" would be considered CORRECT, even though it does not include "Magyarország".
 • Consider the question "Milyen díjat nyert a 'Napsugár' című magyar film a 2018-as Velencei Filmfesztiválon?" with the gold target "Legjobb Forgatókönyv Díja". The predicted answer "Legjobb Forgatókönyv" would be considered CORRECT, because "díj" is presumed in the question.
 • For the question "Mi Dominik Szoboszlai magassága méterben?" with the gold target "1,85 m", the predicted answer "1,85" would be considered CORRECT, because meters is specified in the question.
 • For the question "Mi Magyarország első női miniszterelnökének neve?" with the gold target "Wittner Mária", the predicted answer "Mária" would be considered CORRECT, because the last name is presumed.
 • Do not punish for typos in people's names if it's clearly the same name.
 • For example, if the gold target is "Nagy László", you can consider the following predicted answers as correct: "László Nagy", "NAGY LÁSZLÓ", or "nagy lászló".

Example1:
Input:
{
 "question": "Melyik városban született Tisza Lajos 1832. szeptember 12-én?",
 "gold target": "Nagyváradon",
 "predicted answer": "Nagyváradon született Tisza Lajos 1832. szeptember 12-én."
}
Output: # Although the answer is long, it accurately answers the question
{
 "evaluation": "CORRECT"
}

Example2:
...

Please strictly follow the above example and requirements, evaluate the following answer.
Input:
{
 "question": <question>,
 "gold target": <std_answer>
 "predicted answer": <pred_answer>
}

Please respond strictly in JSON format. Do not include any additional text outside the JSON structure.
Output:
{
 "evaluation": "CORRECT"/"INCORRECT"
}
"""

```

Figure 24: Prompt template for evaluating human-annotated answers in the construction of HuSimpleQA.

```

"""
- Role: Question Screening Expert
- Goals: Screen out the most suitable questions from multiple Hungarian general knowledge questions,
ensuring that the questions and answers meet the following standards:
 1. Relevance to Hungarian Characteristics: Ensure that the question and answer are related to
 Hungarian history, culture, geography, etc.
 2. Appropriate Difficulty: The question should not be overly simple, and the answer should not be
 immediately obvious.
 3. Conciseness: The question and answer should be clear and to the point, avoiding unnecessary
 details.
 4. Single Question: Each question should contain only one query, no sub-questions.
 5. Specificity: The question should be precise and not too broad. Avoid vague, open-ended questions.
 6. Clear Range: For time or location-related questions, avoid vague inquiries like "Mikor" (When),
 as they do not provide a clear timeframe. Instead, ensure the question explicitly asks for a
 specific year, month, day, or a defined period.
 7. Historical Consistency: Focus on fixed, historical facts and events. Avoid questions about
 ongoing trends or future scenarios.
 8. Time/Geography-Specific Queries: If a question includes specific time limitations (such as year,
 month, or specific period) or specific geographic or personal details, the answer should be
 considered fixed and not subject to change over time.
 • This is especially important for questions related to transportation, geography, historical
 landmarks, and iconic structures.
- Constrains:
 1. Selecting the Best-Matching Question and Answer: From a group of questions, select the question
 and answer that best meet the criteria and mark it as 1. All other questions in the group should
 be marked as 0. In a group, there may be at most one question that is selected, but it is also
 possible that none of the questions meet the requirements.
 2. Consistent Evaluation Results: The number of evaluation results must match the number of input
 questions. Ensure that for every question, there is a corresponding evaluation result.
 3. Limit Time-Related Questions for Answer Diversity: Avoid selecting too many questions that focus
 on specific time-related aspects, such as the year an event occurred or a person's birth year.
 Aim to ensure that the questions generate a diverse range of answers.

Example1:
Input:
{
 "question1": "Milyen feltételek szükségesek a bírósági titkári kinevezéshez?",
 "answer1": "Cselekvőképesség, magyar állampolgárság, jogi egyetemi diploma és jogi szakvizsga,
 érvényes pályaalakalmassági vizsgálat, érvényes hatósági erkölcsi bizonyítvány",
 "question2": "Melyik törvény szabályozza az igazságügyi alkalmazottak jogállását Magyarországon?",
 "answer2": "Az 1997. évi LXVIII. törvény (Iasz.)"
}
Output: # Question 1 is too broad and does not meet the requirement of asking and answering concisely
and objectively.
{
 "is_selected": [0,1]
}

Example2:
...

Input:
<input_question>

Please respond strictly in JSON format. Do not include any additional text outside the JSON structure.
{
 "is_selected": [evaluation results]
}
"""

```

Figure 25: Prompt template for selecting optimal question-answer pairs in the construction of HuSimpleQA.

HuSimpleQA Inference Prompt	Translation to EN
<pre> """Kérdés: {question} Kérjük, magyar nyelven adja meg a legjobb választ erre a kérdésre, és 0-tól 100-ig terjedő pontszámmal jelezzze, hogy bízik a válaszában. Kérjük, válaszát a következő JSON formátumban adja meg: {   "answer": "Az Ön válasza itt",   "confidence_score": szám } """ </pre>	<pre> """Question: {question} Please provide your best answer to this question in Hungarian and indicate your confidence in your answer using a score from 0 to 100. Please provide your response in the following JSON format: {   "answer": "Your answer here",   "confidence_score": number } """ </pre>

Figure 26: Prompt template for model inference on HuSimpleQA. The left is the original prompt, the right is the English translation for visualization



```

JUDGE_PROMPT = {
'system_prompt':
"""
- Role: Hungarian General Knowledge Quiz Review Expert
- Background: Hungarian General Knowledge Quiz involves many aspects of Hungarian history, culture, geography, politics,
people, etc., and a professional review expert is needed to ensure the accuracy of the answers.
- Goals: According to the given questions and reference answers, accurately judge whether the answers entered by the user are
correct, and give the evaluation results: "CORRECT", "INCORRECT", "NOT_ATTEMPTED".
- Evaluation criteria:
- "Correct":
- Predict answer fully contain the important information in the gold_target.
- Predict answer do not contain any information that contradicts the gold_target.
- Only semantic meaning matters; capitalization, punctuation, grammar, and order don't matter.
- Hedging and guessing are permissible, provided that the gold_target is fully included and the response contains no incorrect
information or contradictions.
- The answer must be consistent with the scope of the question. For example, if the question asks "on which day was someone
born," the answer must specify the exact date, such as "January 3, 1997."
- "Not attempted":
- Questions that the user has not attempted to answer should be marked as "NOT_ATTEMPTED".
- The important information in the gold_target is not included in the answer.
- No statements in the answer contradict the gold_target.
- Workflow:
1. Receive questions, reference answers, and user answers.
2. Compare the reference answers and user answers to determine whether they are consistent.
3. Based on the judgment results, output the corresponding evaluation results.
- Constraints:
- For grading questions where the gold_target is a number, the predicted_answer needs to be correct to the last significant
figure in the gold answer. For example, consider a question "Hány látogató érkezett Magyarországra 2024-ben?" with gold_target
"440k".
- predicted_answers "440k", "444k", and "435k" are all CORRECT.
- predicted_answers "400k" and "413k" are INCORRECT.
- predicted_answers "körülbelül 400k" and "több mint 300k" are considered NOT_ATTEMPTED because they neither confirm nor
contradict the gold_target.
- The gold_target may contain more information than the question. In such cases, the predicted_answer only needs to contain
the information that is in the question.
- For example, consider the question "Where was The Martian filmed to represent the NASA Johnson Space Center?" with the
gold_target "Budapesti Bálna (HU)". Either "Budapesti Bálna" or Bálna (épitmény) would be considered a CORRECT answer.
- Do not punish predicted_answers if they omit information that would be clearly inferred from the question.
- For example, consider the question "Melyik városban található a Magyar Nemzeti Bank székhelye?" with the gold_target
"Budapest, Magyarország". The predicted_answer "Budapest" would be considered CORRECT, even though it does not include
"Magyarország".
- Consider the question "Milyen díjat nyert a 'Napsugár' című magyar film a 2018-as Velencei Filmfesztiválon?" with the
gold_target "Legjobb Forgatókönyv Díja". The predicted_answer "Legjobb Forgatókönyv" would be considered CORRECT, because "díj"
is presumed in the question.
- For the question "Mi Dominik Szoboszlai magassága méterben?" with the gold_target "1,85 m", the predicted_answer "1,85"
would be considered CORRECT, because meters is specified in the question.
- For the question "Mi Magyarország első női miniszterelnökének neve?" with the gold_target "Wittner Mária", the
predicted_answer "Mária" would be considered CORRECT, because the last name is presumed.
- Do not punish for typos in people's names if it's clearly the same name.
- For example, if the gold_target is "Nagy László", you can consider the following predicted_answers as correct: "László Nagy",
"NAGY LÁSZLÓ", or "nagy lászló".
Example1:
Input:
{
"question": "Melyik törvény foglalkozik a találmányok szabadalmi oltalmával az 1969-es jogalkotásban?",
"gold_target": "1969. évi II. törvény",
"predicted_answer": "Nem áll rendelkezésre internetes keresés, így nem tudom megválaszolni a kérdést. Azonban 1969-ben valóban
elfogadták a szabadalmi védelmi törvényt."
}
Output:
{
"evaluation": "NOT_ATTEMPTED"
}
Example2:
...
""",
'user_prompt':
"""Please strictly follow the above example and requirements, evaluate the following answer. Input:
{{
"question": {question},
"gold_target": {answer},
"predicted_answer": {pred_answer}
}}
Please respond strictly in JSON format. Do not include any additional text outside the JSON structure.
Output:
{{
"evaluation": "Correct"/"Incorrect"/"NOT_ATTEMPTED"
}}
""",
}

```

Figure 27: Prompt template for LLM as judge on HuSimpleQA.

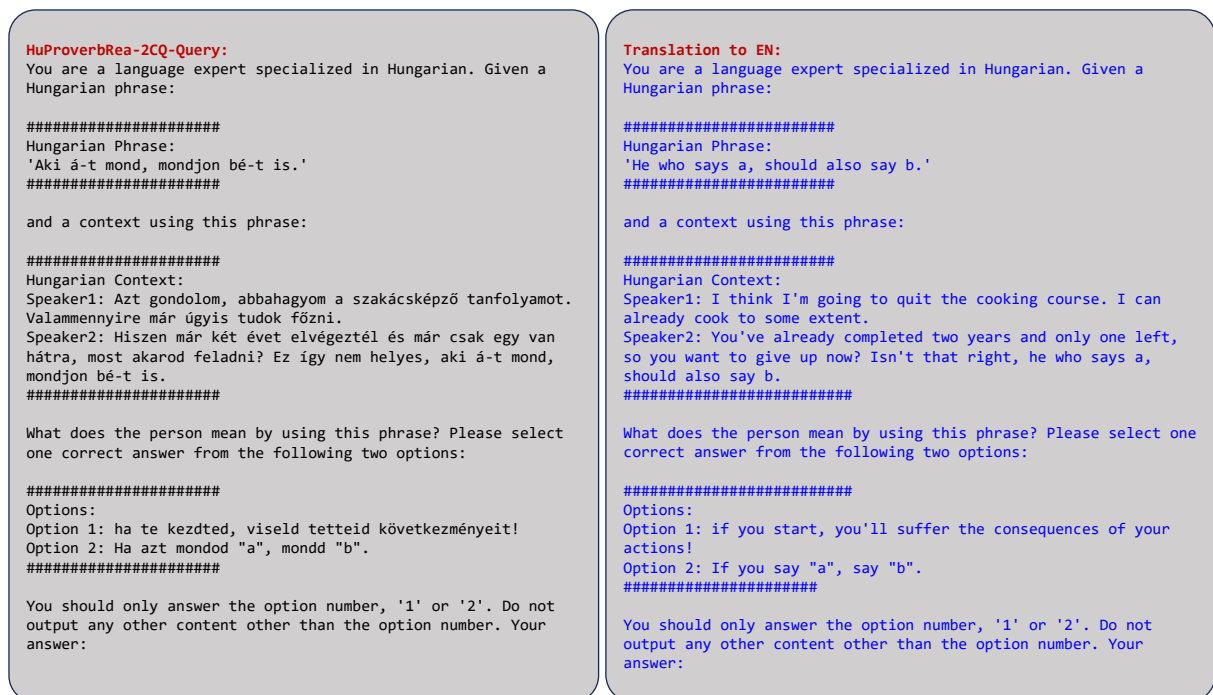


Figure 28: Example of HuProverbRea (2CQ). The left is the original example in OpenHuEval, the right is the English translation for visualization.

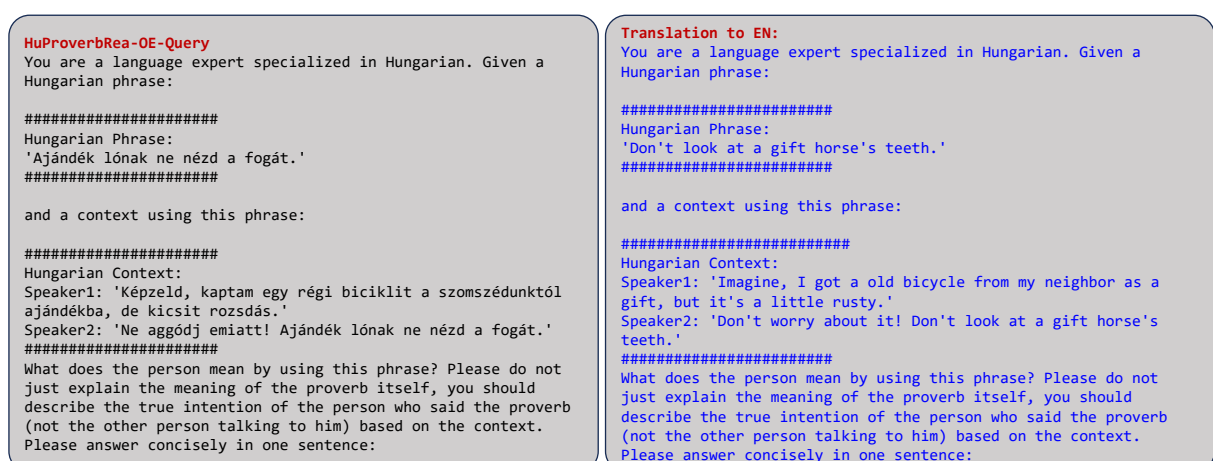


Figure 29: Example of HuProverbRea (OE). The left is the original example in OpenHuEval, the right is the English translation for visualization.

```

"""
You are a language expert specialized in Hungarian. Given a Hungarian phrase:
#####
Hungarian Phrase:

'{hu_text}'
#####
and a context using this phrase:
#####
Hungarian Context:

{context}
#####
What does the person mean by using this phrase? Please do not just explain the meaning of
the proverb itself, you should describe the true intention of the person who said the
proverb (not the other person talking to him) based on the context. Please answer concisely
in one sentence:
"""

```

Figure 30: Prompt template for model inference on HuProverbRea (OE).

```

"""
You are a language expert specialized in Hungarian. Given a Hungarian phrase:
#####
Hungarian Phrase:

'{hu_text}'
#####
and a context using this phrase:
#####
Hungarian Context:

{context}
#####
What does the person mean by using this phrase? Please select one correct answer from the
following two options:
#####
Options:

Option 1: {option1}
Option 2: {option2}
#####
You should only answer the option number, '1' or '2'. Do not output any other content other
than the option number.
Your answer:
"""

```

Figure 31: Prompt template for model inference on HuProverbRea (2CQ).

```

'en_system': "" Please act as an impartial judge specialized in Hungarian language and culture. Given a
Hungarian saying, a context using that saying, and two analyses explaining 'what does the person mean by using
that saying in the context?', please decide whether the given two analyses express the same meaning. If they
reflect the same understanding of the saying's meaning, you should answer YES. If they are based on different
interpretations of the saying, you should answer NO. Do not output anything other than 'YES' or 'NO'. Avoid
any position biases and ensure that the order in which the analyses were presented does not influence your
decision. Do not allow the length of the analyses to influence your judge, focus on their core meanings and
their understandings of the Hungarian saying. ""
'en_user': ""
[The start of Hungarian saying]
{proverb}
[The end of Hungarian saying]

[The start of the context]
{conversation}
[The end of the context]

[The start of the first analysis]
{answer}
[The end of the first analysis]

[The start of the second analysis]
{raw_pred}
[The end of the second analysis]
""
'Your decision:

```

Figure 32: Prompt template for LLM as judge on HuProverbRea.

#### HuMatchingFIB-Hugarian

```

{
 "q_main": "Válassz ki a legördülő listából, hogy
melyik fogalom illik a hiányos mondatokba!\nA faj
azon egyedeit, melyek tényleges szaporodási
közösséget alkotnak, #0# nevezzük.\nA/az #1#
mindazoknak a hatásoknak az összessége, melyek
ténylegesen hatnak az élőlényekre.\nA populáció
méretét jellemző egyik legfontosabb sajátosság a/az
#2#.\nTerület- vagy térfogategységre vonatkoztatott
egyedszám a/az #3#.\nA környezeti tényező azon
tartománya, melyen belül az élőlények
életműködéseket mutatnak a #4#.\nJellemzően az a
környezeti tényező határozza meg a populáció
elterjedését, amelyre nézve az adott faj szűk tűrésű,
ezt nevezzük úgy, hogy #5#.",
 "options": [
 "A.környezet",
 "B.tűrőképesség",
 "C.egyedsűrűség",
 "D.egyedszám",
 "E.korlátozó tényező",
 "F.populációnak"
],
 "std_ans": [
 "#0#F",
 "#1#A",
 "#2#D",
 "#3#C",
 "#4#B",
 "#5#E"
]
}

```

#### Translation to EN

```

{
 "q_main": "Select from the dropdown list which
concept fits into the incomplete sentences!\nThe
individuals of a species that form an actual
reproductive community are called #0#.\nThe #1# is
the totality of all effects that actually influence
living organisms.\nOne of the most important
characteristics describing the size of a population
is the #2#.\nThe number of individuals per unit area
or volume is the #3#.\nThe range of an environmental
factor within which living organisms exhibit life
processes is the #4#.\nTypically, the environmental
factor that determines the distribution of a
population is the one for which the species has a
narrow tolerance, and this is called the #5#.",
 "options": [
 "A.environment",
 "B.tolerance",
 "C.population density",
 "D.population size",
 "E.limiting factor",
 "F.population"
],
 "std_ans": [
 "#0#F",
 "#1#A",
 "#2#D",
 "#3#C",
 "#4#B",
 "#5#E"
]
}

```

Figure 33: Example of HuMatchingFIB. The left is the original example in OpenHuEval, the right is the English translation for visualization.



HuStandardFIB-Hungarian	Translation to EN
<pre> {   "q_main": "Találd ki a leírások alapján, hogy kiről vagy miről van szó! Írd be a meghatározások után a megfelelő kifejezéseket!",   "std_ans": [     "#0#Mánuel;Mánuelcsászár",     "#1#kancellária",     "#2#Anonymus",     "#3#írásbeliség",     "#4#jegyző",     "#5#Székesfehérvár;Fehérvár"   ],   "formatted_q_sub": [     "A.Az #0# udvarában nevelkedett ifjúkorában III. Béla:",     "B.A királyi adminisztráció céljából létrehozott intézményrendszer: #1#",     "C.Feltehetően ő volt III. Béla jegyzője: #2#",     "D.1181-ben tette általánossá III. Béla a hivatali ügyintézésben: #3#",     "E.#4# fogalmazta meg a hivatalos iratokat, okleveleket:",     "F.Ebben a városban temették el III. Bélát: #5#"   ],   "formatted_std_ans": [     "#0#Mánuel;Mánuelcsászár",     "#1#kancellária",     "#2#Anonymus",     "#3#írásbeliség",     "#4#jegyző",     "#5#Székesfehérvár;Fehérvár"   ], } </pre>	<pre> {   "q_main": "Based on the descriptions, guess who or what is being referred to! Enter the appropriate terms after the definitions!",   "std_ans": [     "#0#Manuel;Emperor Manuel",     "#1#chancellery",     "#2#Anonymus",     "#3#written records",     "#4#scribe",     "#5#Székesfehérvár;Fehérvár"   ],   "formatted_q_sub": [     "A. In the court of #0#, Béla III spent his youth:",     "B. The institutional system created for royal administration: #1#",     "C. He was likely the scribe of Béla III: #2#",     "D. In 1181, Béla III made this mandatory in official proceedings: #3#",     "E. #4# was responsible for drafting official documents and charters:",     "F. The city where Béla III was buried: #5#"   ],   "formatted_std_ans": [     "#0#Manuel;Emperor Manuel",     "#1#chancellery",     "#2#Anonymus",     "#3#written records",     "#4#scribe",     "#5#Székesfehérvár;Fehérvár"   ], } </pre>

Figure 34: Example of HuStandardFIB. The left is the original example in OpenHuEval, the right is the English translation for visualization.

```

""The following questions are in Hungarian language on {hu_specific_dim}, please read the questions, and try
to fill in the blanks in the question list. Please organize the answer in a list. An example:
{
 "instruction": "Írd be a megfelelő meghatározás mellé a fogalmat!",
 "questions": ["A.A szerzetesi közösségek szabályzatának elnevezése latinul: #0#", "B.Az első ún. kolduló
rend: #1#", "C.A szerzetesek által kézzel másolt mű: #2#", "D.Papi nőtlenség: #3#", "E.A pápát megválasztó
egyházi méltóságok: #4#", "F.A bencés rend megújítása ebben a kolostorban kezdődött a 10. században: #5#"],
}
The answers are:
{
 "answers": ["#0#regula", "#1#ferencesrend", "#2#kódex", "#3#cölibátus", "#4#bíborosok", "#5#Cluny"]
}
Now try to answer the following questions, your response should be in a JSON format. Contain the "answers"
like the case given above.
The questions are:
{
 "instruction": {instruction},
 "questions": {questions},
}
""

```

Figure 35: Prompt template for model inference on HuStandardFIB.

```

"""You are a native Hungarian teacher. The following question is in Hungarian language on {hu_specific_dim}.
Please read the question, and choose the appropriate option from the provided "options" list to fill in each
blanks in the text based on the context. Read the entire text, then fill in the blanks. Some options can be
selected repeatedly. Please organize the answer in a list. An example:
{
 "question": "Egészítsd ki a Janus Pannonius életére vonatkozó rövid szöveget! Segítségként használd az
internetet! Vigyázz, nem minden szót kell felhasználnod!\nJanus Pannonius nem csupán költőként volt jelentős
személyisége kora Magyarországnak. #0# unokaöccseként a politikából is hamar kivette a részét. #1#
tanulmányai után pécsi #2# lett, majd a királyné mellett #3#. Főkecsként és a #4# báni cím elnyerésével
komoly politikai karriert futott be Mátyás király udvarában. A királlyal megromló kapcsolata miatt részt vett
a #5# elleni összeesküvésben, ezért menekülnie kellett. Ez, és az akkor már súlyosbodó betegsége okozta
halálát #6#.",
 "options": ["A.érsek", "B.szlavón", "C.Vitéz János", "D.püspök", "E.főpohárnok", "F.Ulászló", "G.1474-ben",
"H.főkancellár", "I.Itáliai", "J.Kinizsi Pál", "K.Kálmán", "L.1472-ben", "M.Prágai", "N.Mátyás"],
},
The answer is:
{
 "answer": ["#0#C", "#1#I", "#2#D", "#3#H", "#4#B", "#5#N", "#6#L"]
}
Now try to answer the following question, your response should be in a JSON format. Contain the "answer" like
the case given above.
The question and options are:
{
 "question": {question},
 "options": {options},
}
"""

```

Figure 36: Prompt template for model inference on HuMatchingFIB.

"""

Given a thinking process for answering a question, follow these steps to extract contrastive expressions from the answer text:

1. Identify the Primary Language:

- First, determine the primary language of the answer text. The language could be English, Hungarian, or any other language.

2. Extract Contrastive Words, Phrases, or Expressions:

- Identify all the phrases that express a shift in opinion, explanation, or answer, phrases that signal a contrast or change in direction.
- For English: "However," "but," "On the other hand," "Although," "Nevertheless," "Yet," "Despite," "In contrast," "Instead," "Even though."
- For Hungarian: "azonban," "De," "Másrészt," "Ellentétben," "Pedig," "MÉGIS," "Bár,"

- Requirements:

1. Identify and list all the contrastive words or phrases that indicate a shift in meaning, thought, or direction.
2. These expressions should be **\*\*at the beginning of a sentence\*\*** to signal a shift.
3. Keep the original text's meaning and context intact.
4. Ensure to maintain the original capitalization of the words (e.g., "However" vs. "however").
5. Provide a clear list of all the identified contrast words or phrases.

Input text:

```
{
 "question": <question>,
 "answer": <answer>,
}
```

Please respond strictly in JSON format. Do not include any additional text outside the JSON structure. The output should also include the detected language type.

Output:

```
{
 "language": "<detected_language>",&br/> "shift_expression": [list]
}
```

"""

Figure 37: Prompt template for splitting the LRM's thinking process into thoughts on HuSimpleQA (step 1/2).

```

"""
Given the thinking process, identify all the phrases that express a shift in
opinion, explanation, or answer, i.e., phrases that signal a contrast or
change in direction (commonly known as "contradiction," "contrast," or
"transition" phrases). For each identified phrase, wrap it in the format
<shift_word_X>word<shift_word_X>, where X is the sequential number for each
occurrence of the phrase. Notice that the text provided may be in either
Hungarian or English.
You should provide a list of the sequence numbers corresponding to the
phrases that convey a shift in meaning.
If there is no "<shift_word_X>" tag in text, return an empty list.

Input text:
<input_text>
Please respond strictly in JSON format. Do not include any additional text
outside the JSON structure.
Output:
{
 "shift_word_ids": [list of sequential numbers of shifting phrases]
}
"""

```

Figure 38: Prompt template for splitting the LRM's thinking process into thoughts on HuSimpleQA (step 2/2).

```

"""
Question Q = <question>
Splited Thought S = <split_solutions>
Correct Answer A = <expected_answer>

1. Please analyze the relevance between the splited thought S and the
questin Q, and conduct some verifications to check the correctness of the
solution itself. Please think step by step to give an explanation
EXPLANATION.

2. If you think the splited thought S can lead to the orrect answer A of the
problem P, please stick to the line of thinking without deviation and carry
it through to completion. If you think it cannot yield the correct answer or
you're not sure, don't force yourself to give an answer and generate
None.

3. Please tell me honestly how confident you are that you can solve the
problem P correctly based on the the splited thought S. Out of 2, please
generate your confidence score **CONFIDENT SCORE**.

Please output **EXPLANATION** and **CONFIDENT SCORE** according to the
following JSON format:
{
 "expalanation": "EXPLANATION",
 "confident_score": "CONFIDENT SCORE"
}
"""

```

Figure 39: Prompt template for evaluating the correctness of the thoughts of the LRM's thinking process on HuSimpleQA.

**Question:** Melyik hivatal látja el a Hamisítás Elleni Nemzeti Testület titkársági feladatait?

Which office handles the secretarial tasks of the National Board Against Counterfeiting?

**Correct Answer:** Szellemi Tulajdon Nemzeti Hivatala  
National Office of Intellectual Property

Thought1	Okay, let's tackle this question. The user is asking which office handles the secretarial tasks for the National Anti-Forgery Committee (Hamisítás Elleni Nemzeti Testület) in Hungary. I need to figure out the correct answer and assess my confidence. First, I recall that in Hungary, various government bodies have specific roles. The National Anti-Forgery Committee is likely related to financial or document security. I remember that the Hungarian National Bank ...the secretariat duties could be under a different department.	✓
Thought2	<u>Alternatively</u> , the Government Office (Kormányhivatal) usually handles administrative tasks for various committees.	✗
Thought3	<u>Wait</u> , the National Tax and Customs Administration (Nemzeti Adó- és Vámhivatal, NAV) deals with tax, customs, and financial fraud. Since counterfeiting currency is a big part of their work, maybe NAV supports the committee	✗
.....	.....	
Thought5	<u>But</u> given the context, NAV seems correct	✗
<b>Predicted Answer:</b> A Hamisítás Elleni Nemzeti Testület titkársági feladatait a Nemzeti Adó- és Vámhivatal (NAV) látja el. The secretariat tasks of the National Body Against Counterfeiting are carried out by the National Tax and Customs Administration (NAV).		

Figure 40: Example of the thoughts in Deepseek-R1's thinking process on HuSimpleQA. ✓ and ✗ indicate the correctness of the thoughts. The original text is in black, while the translation into English is in blue.



**Question:** Melyik magyar rendező filmjében nyújtott alakításáért kapott Helen Mirren különdíjat a 2012-es Európai Filmdíjakon?

For her performance in which Hungarian director's film did Helen Mirren receive a special award at the 2012 European Film Awards?

**Correct Answer:** Szabó István (Istvan Szabo)

Thought1	<p>Melyik magyar rendező filmjében nyújtott alakításáért kapott Helen Mirren különdíjat a 2012-es Európai Filmdíjakon? Nagyon nehéz kérdés, nem vagyok teljesen biztos a válaszban, de megpróbálom kitalálni. Először is, Helen Mirren egy angol nagyjátékostermek, de számos filmben játszott, így valószínűleg valami brit vagy Európa kapcsolattal filmben játszott, amelyet magyar rendező forgattat.</p> <p>In which Hungarian director's film did Helen Mirren receive a special award at the 2012 European Film Awards? This is a very difficult question, and I'm not completely sure of the answer, but I'll try to figure it out. First of all, Helen Mirren is a renowned English actress who has appeared in many films, so it's likely that she acted in a British or Europe-related film directed by a Hungarian director.</p>	✓
Thought2	<p><b>De</b> a kérdésről szóló részlet szerint a magyar rendező filmjében nyújtott alakításáért kapta a díjat. Hmmm, ki lehetne a magyar rendező? Tudom, hogy néhány magyar rendező sikert ért el nemzetközi filmekben, de pontosan nem emlékszem rájuk most. Lehet Béla Tarr?</p> <p><b>But</b> according to the detail of the question, she received the award for her performance in a film by a Hungarian director. Hmm, who could the Hungarian director be? I know that some Hungarian directors have achieved success in international films, but I can't exactly recall them right now. Could it be Béla Tarr?</p>	✗
.....	.....	
Thought20	<p><b>De most</b>, a fejemben, nem emlékszem rá. Így, a válaszom a kérdésre, hogy sajnálom, de nem tudom, melyik magyar rendező filmjében nyújtott alakításáért kapott Helen Mirren különdíjat a 2012-es Európai Filmdíjakon.</p> <p><b>But right now</b>, I can't remember it. So, my answer to the question is: Sorry, I don't remember which Hungarian director's film Helen Mirren received the special award for her performance at the 2012 European Film Awards.</p>	✗
	<p><b>Predicted Answer:</b> Sajnálom, de nem emlékszem arra a filmre, ahol Helen Mirren nyújtott alakításáért kapott különdíjat a 2012-es Európai Filmdíjakon, és melyik magyar rendező műve volt.</p> <p>Sorry, but I don't remember the film where Helen Mirren received a special award for her performance at the 2012 European Film Awards, nor which Hungarian director's work it was.</p>	

Figure 41: Example of the thoughts in QWQ's thinking process on HuSimpleQA. ✓ and ✗ indicate the correctness of the thoughts. The original text is in black, while the translation into English is in blue.

You are a semantic paragraph segmentation expert, responsible for dividing the thinking process content I provide to you (generated by a large language model) into paragraphs. The content of the thinking process pertains to the reasoning and solving process of fill-in-the-blank questions in Hungarian. The background of the thinking process content is the reasoning and solving process for fill-in-the-blank questions in the Hungarian version. I need you to segment the original complete thought process content into multiple paragraphs and assign each paragraph a tag strictly limited to the categories: "Introduction", "Reasoning", "Review", and "Summary", based on its content. Below, I will describe the characteristics of these four types in detail and provide examples for reference. In most cases, the thinking process content is presented in English, with a very small portion in Hungarian. You can apply the same logic for segmentation. Please note that no additional content should be added or removed from the original thinking process content; Additionally, there should be no overlap between the divided paragraphs.

#### Segment 1: Introduction

Description: The introduction is typically located at the beginning of the thinking process content. It usually consists of the large language model's brief restatement of the problem and a descriptive account of the work it is about to undertake. It does not include the actual start of the analysis of the problem.

Such statements may generally include the following:

- (1) Alright, I have this history question to complete. It's about the concept of royal power and political systems in Western Europe, specifically in England and France during a certain period. I need to fill in the blanks using the provided options. Let's see, there are nine blanks, and I have nine options to choose from, but some might be used more than once, though the example didn't specify that. I'll approach this step by step.
- (2) I have this task to complete a diagram by dragging expressions to their corresponding numbers. The expressions are...
- (3) I have this task here. I need to find the odd one out from each group of words. Each group has words that belong to one part of speech, except for one word that doesn't fit in that group. I need to identify the odd one out and state its part of speech.
- (4) I'm going to answer this question about the Csörsz-ditch. I need to decide whether each statement is true or false based on the information provided and any knowledge I have about the topic. Let's go through each one step by step.
- (5) I'm going to try to fill in the blanks in this text. It seems like a story about someone exploring unknown places, maybe flying or something like that. I have a list of options to choose from, and I need to pick the right ones to complete the sentence properly. I should pay attention to the context and make sure the words fit grammatically and make sense in the story.

#### Segment 2: Reasoning (Important)

Description: The reasoning process typically constitutes the main body of the thinking process content. It includes the detailed thinking and reasoning steps undertaken by the large language model to solve the fill-in-the-blank questions. You should collect, as thoroughly and sequentially as possible, the content that you identify as part of the "reasoning".

The use of '\n\n' paragraph separators may serve as a suitable paragraph division choice, but please note that answer-related statements may also utilize '\n\n' for line breaks or section divisions. Exercise judgment to distinguish between these usages. Paragraphs in the Reasoning section should neither be excessively brief nor unduly lengthy.

#### Segment 3: Review (Important)

Description: The review usually occurs after the reasoning process is essentially complete but before the summary. This section typically includes a review of the entire reasoning process and may contain keywords or phrases such as "Overall, ..." or "double check..."

Please note that not all thinking process content necessarily includes a review content; in some cases, the reasoning process may be directly followed by the summary. In such instances, you can refer to the example response format provided.

Figure 42: Prompt template for splitting LRM's thinking process into segments and categorizing these segments on HuMatchingFIB. (part 1/3).

#### Segment 4: Summary

Description: The summary is generally the model's ultimate output, i.e., the part where the model provides the summary after completing all the reasoning in the thinking process. It may also be presented at the very beginning of the thinking process, in which case it might be directly displayed in a JSON format, requiring your judgment. It typically includes some indicative phrases, such as "...final answer..." or "...final choices...". However, please note that content containing "...summarize..." may not necessarily be the summary; it could be part of the intermediate reasoning process. Be sure to distinguish such content and exclude it from the summary.

The statements in the summary may generally include the following:

(1) So, my final answer is:

```
{
 "answer": ["#0#B", "#1#F", "#2#H", "#3#D", "#4#G", "#5#H", "#6#C", "#7#C"]
}
```

(2) I'll present this in the required JSON format.

**\*\*Final Answer\*\***

```
\\[\\boxed{ \\{ "answer": ["#0#A", "#1#G", "#2#C", "#3#I", "#4#B", "#5#C", "#6#H",
"#7#E", "#8#F"] \\} } \\]
```

(3) So, the final answer should be:

```
{
 "answer": ["#0#L", "#1#H", "#2#A", "#3#I", "#4#E", "#5#C", "#6#K", "#7#M", "#8#B",
"#9#M", "#10#D", "#11#F"]
}
```

The input content you receive is after [input thinking process content] and you should response strictly in the provided format. The specific content should be added after the [Your segmentation results] field and must in JSON format:

[input thinking process content]

COTs content

[Your segmentation results]

```
{
 'segment':[
 {
 'text': content you regard as "Introduction",
 'category': "Introduction"
 },
 {
 'text': content you regard as "Reasoning",
 'category': "Reasoning"
 },
 ...,
 {
 'text': content you regard as "Review",
 'category': "Review"
 },
 {
 'text': content you regard as "Summary",
 'category': "Summary"
 }
]
}
```

Figure 43: Prompt template for splitting LRM's thinking process into segments and categorizing these segments on HuMatchingFIB. (part 2/3).

Some specific examples are as follows:

```
[input thinking process content]
...
[Your segmentation results]
{
 'segment':[
 {
 'text': "<think>Okay, let me try to figure out how to answer this. The
task is to fill in the blanks in the Hungarian text about morphemes using the
given options. Let's look at each blank step by step.\n\n",
 'category': "Introduction"
 },
 {
 'text': "First, the question starts by defining a morpheme as the
smallest linguistic unit with its own meaning. So blank #0# should be
\"unit\", which in Hungarian is \"egység\". Checking the options, H is
\"egység\".",
 'category': "Reasoning"
 },
 {
 'text': "Then #1# is \"meaning\", which is \"jelentése\" (option B).\n\n",
 'category': "Reasoning"
 },
 ...
 {
 'text': "Putting it all together:\n\n#0: H (egység)\n#1: B
(jelentése)\n#2: A (toldalék)\n#3: C (egyszerű)\n#4: D (összetett)\n#5: F
(képző)\n#6: G (jel)\n#7: E (rag)</think>{\n \"answer\": [\"#0#H\",
\"#1#B\", \"#2#A\", \"#3#C\", \"#4#D\", \"#5#F\", \"#6#G\", \"#7#E\"]\n}",
 'category': "Summary"
 }
]
}

[input thinking process content]
...
[Your segmentation results]
...
```

Now, the target content you need to split is as follows. Please provide your standardized answer after [Your segmentation results] in JSON format:

```
[input thinking process content]
{Raw_COT}
[Your segmentation results]
```

Figure 44: Prompt template for splitting LRM's thinking process into segments and categorizing these segments on HuMatchingFIB. (part 3/3).

**Role**

You are a text judgement and reasoning expert.

**Task Description**

Your task is to perform multidimensional classification of the output from a reasoning model. The model's output has been segmented into multiple segments (Introduction, Reasoning, Review, Final\_answer), among which there are Reasoning-type segments. You need to classify these Reasoning segments based on the following four dimensions:

**Dimensions and Classification Rules:**

**Dimension 1: Correctness of the Result**

Based on the options, the standard answer (std\_ans) and the model's answer (model\_ans), determine whether the result in each Reasoning segment is correct. The classification is as follows:

Class 1: Completely Incorrect

All blank-filling results in the segment do not match the standard answer.

Class 2: Partially Correct

Some blank-filling results in the segment match the standard answer, while others do not.

Class 3: Completely Correct

All blank-filling results in the segment match the standard answer.

Class 4: Non Conclusion

No conclusion has been provided yet.

**Constraints:**

If the segment involves multiple blanks, compare each result with the standard answer.

If the segment does not explicitly mention the blank-filling results, infer based on the context.

**Dimension 2: Reasoning Complexity**

Determine whether the reasoning process in each segment is a simple assertion or involves complex thinking. The classification is as follows:

Class 1: Simple Assertion

The segment directly provides the answer without detailed reasoning.

Class 2: Complex Thought

The segment includes repeated thinking, logical reasoning, hypothesis validation, or other complex processes.

**Constraints:**

If the segment contains keywords such as: "Wait, perhaps...", "I need to consider...", "Alternatively...", "Hmm, maybe...", "Let me think..." classify it as "Complex Thought."

If the segment only directly provides the answer (e.g., "#1# is H.508"), classify it as "Simple Assertion."

**Dimension 3: Reasoning Scope**

Determine whether the reasoning in each segment involves modifying any previously solved blanks. The classification is as follows:

Class 1: Only Current Blank

The segment only provides an answer for the unresolved blank and does not modify previously solved blanks.

Class 2: Modify Previous Blanks

The segment not only provides an answer for the unresolved blank but also modifies or corrects previously solved blanks.

Class 3: Current Blank and Consecutive Blank

The segment provides an answer for the current unresolved blank and also addresses consecutive blanks, either by solving them or making adjustments.

**Constraints:**

If the segment explicitly mentions modifying previously solved blanks (e.g., "Wait, I need to change #2# to..."), classify it as "Modify Previous Blanks."

If the segment only focuses on the current blank, classify it as "Only Current Blank."

If the segment addresses both the current blank and consecutive blanks, classify it as "Current Blank and Consecutive Blank."

Figure 45: Prompt template for tagging the reasoning segments of LRM's thinking process along the four dimensions on HuMatchingFIB (part 1/2).



**Dimension 4: Language Transfer**  
Determine whether each Reasoning segment includes the process of translating Hungarian into English. The classification is as follows:  
**Class 1: Contains Language Transfer**  
The segment includes a translation process similar to:  
“Erőteljes #3# és a költői #4# gazdag használata jellemzi.”  
This translates to “It is characterized by strong #3# and rich use of poetic #4#.”  
**Class 2: No Language Transfer**  
The segment does not include the above translation process.  
**Constraints:**  
If the segment contains an explicit translation process (e.g., “This translates to...”), classify it as “Contains Language Transfer.”  
If the segment only uses Hungarian or English without translation, classify it as “No Language Transfer.”

**Example Input:**

```
{
 "options": [...],
 "std_ans": [...],
 "model_ans": [...],
 "segments": [
 {
 'text': ...
 'category': "Introduction"
 },
 {
 'text': ...
 'category': "Reasoning"
 },
 ...
]
}
```

**Example Output:**

```
{
 "segments": [
 {
 'text': ...
 'category': "Introduction"
 },
 {
 'text': ...
 'category': "Reasoning",
 'Dimension1': "Non Conclusion",
 'Dimension2': "Complex Thought",
 'Dimension3': "Only Current Blank",
 'Dimension4': "No Language Transfer",
 },
 ...
 {
 'text': ...
 'category': "Reasoning",
 'Dimension1': "Completely Correct",
 'Dimension2': "Complex Thought",
 'Dimension3': "Current Blank and Consecutive Blank",
 'Dimension4': "No Language Transfer",
 },
 ...
]
}
```

**Notes**

1. Ensure that every Reasoning segment is classified, and the classification results are logical.
2. If a dimension cannot be clearly classified for a segment, infer the most appropriate category based on the context.
3. The output must be in JSON format and include classification results for all four dimensions.

Figure 46: Prompt template for tagging the reasoning segments of LRM’s thinking process along the four dimensions on HuMatchingFIB (part 2/2).

```
[
 {
 "text": "<think>Okay, let me try to work through this problem step by step.
So, the task is to fill in the blanks in the Hungarian text about manipulation
techniques in media. Each blank corresponds to a term from the provided options.
Let's look at each numbered section and see which option fits best.",
 "category": "Introduction"
 },
 ...
 {
 "text": "#4 is \"Egy gondolat bánt engemet...\", described as a \"zaklatott
menetű költemény\" (disturbed flow poem). The options include rapszódia (E), which
is a rhapsody, often passionate or irregular in form. That seems to fit, so #4 is
E.",
 "category": "Reasoning",
 "Dimension1": "Completely Correct",
 "Dimension2": "Simple Assertion",
 "Dimension3": "Only Current Blank",
 "Dimension4": "Contains Language Transfer"
 },
 {
 "text": "#5 is \"Szeptember végén\", a love poem to his wife. The options
left include A (dal), C (elégia), etc. But \"szerelmi vallomás\" (love confession)
might be a song or lyric poem. Dal (A) is possible, but I'm not certain. Wait, the
example says \"szerelmi vallomása\", which could be a lyric poem. Dal (song) is a
common form for love poems. So #5 is A.",
 "category": "Reasoning",
 "Dimension1": "Completely Correct",
 "Dimension2": "Complex Thought",
 "Dimension3": "Only Current Blank",
 "Dimension4": "Contains Language Transfer"
 },
 ...
 {
 "text": "Let me confirm:\n\n- #0: János vitéz → G (elbeszélő költemény)\n-
#1: Helység kalapácsa → H (vígeposz)\n- #2: Borozó → A (dal)\n- #3: Felhők ciklus
verse → C (elégia)\n- #4: Egy gondolat... → E (rapszódia)\n- #5: Szeptember végén →
A (dal)\n- #6: Tigris és hiéna → B (regény)\n- #7: A hóhér kötele → D (dráma)",
 "category": "Review"
 },
 ...
 {
 "text": "{\n \"answer\": [\n \"#0#G\", \n \"#1#H\", \n
\"#2#A\", \n \"#3#C\", \n \"#4#E\", \n \"#5#A\", \n
\"#6#B\", \n \"#7#D\" \n]\n}",
 "category": "Summary"
 }
]
```

Figure 47: Example of splitting Deepseek-R1's thinking process into segments, categorizing these segments, and tagging the reasoning segments on HuMatchingFIB. The categorizing results are in green and the tagging results are in red.

```
[
 {
 "text": "Alright, I have this task to fill in the blanks in a text about bank cards using the provided options. The question is in Hungarian, and I need to make sure that the words I choose fit correctly in the context. Let's see, the text is about bank cards being a substitute for cash, and how they can be used for purchases, withdrawing money, and checking balances. It also mentions that the card belongs to the issuing bank, which can reclaim it in case of non-payment or cancel its use.",
 "category": "Introduction",
 },
 ...
 {
 "text": "Moving on to the third blank: \"Segítségével a vásárláson kívül még pénzt lehet felvenni a bankjegykiadó automatából, vagy tájékozódni lehet a #2# egyenlegéről.\" This translates to \"With its help, besides shopping, you can also withdraw money from an ATM or check your... balance.\" So, it's about checking a balance related to something. Looking at the options, \"bankszámla\" means \"bank account,\" which makes sense because you check your account balance. \"Megtakarításunk\" means \"our savings,\" which could also be related, but \"bankszámla\" is more precise. So, I'll go with \"A.\tbankszámla.\",
 "category": "Reasoning",
 "Dimension1": "Completely Correct",
 "Dimension2": "Complex Thought",
 "Dimension3": "Only Current Blank",
 "Dimension4": "Contains Language Transfer"
 },
 ...
 {
 "text" : "Fifth and final blank: \"...illetve a #4# letiltani.\" This means \"and to cancel its use.\" So, it's about the bank being able to cancel the card's usage. Looking at the options, \"használatát\" means \"its usage,\" which fits perfectly here. So, \"E.\thasználatát\" is the right choice.",
 "category": "Reasoning",
 "Dimension1": "Completely Correct",
 "Dimension2": "Simple Assertion",
 "Dimension3": "Only Current Blank",
 "Dimension4": "Contains Language Transfer"
 },
 ...
 {
 "text": "Let me double-check all the choices:\n\n- #0# F.\tfizetési (payment)\n\n- #1# B.\tkészpénzzel (with cash)\n\n- #2# A.\tbankszámla (bank account)\n\n- #3# G.\tkibocsátó (issuing)\n\n- #4# E.\thasználatát (its usage)\n\nSeems logical and consistent with the context provided. I don't think I need to make any changes.",
 "category": "Review"
 },
 {
 "text": "***Final Answer**\n\n\\[\\boxed{ \\text{#0#F, #1#B, #2#A, #3#G, #4#E} } \\]",
 "category": "Summary"
 }
]
```

Figure 48: Example of splitting QwQ's thinking process into segments, categorizing these segments, and tagging the reasoning segments on HuMatchingFIB. The categorizing results are in green and the tagging results are in red.