

Lattice at SemEval-2026 Task 1: Why did the prompt engineer break up with their LLM? Because zero-shot was zero-fun.

Mathieu Dehouck, Olga Seminck, Noé Durandard, Yoann Dupont,
Lattice UMR8094, CNRS, ENS-PSL, Sorbonne Nouvelle
mathieu.dehouck@cnrs.fr, olga.seminck@cnrs.fr,
noe.durandard@ens.psl.eu, yoann.dupont@sorbonne-nouvelle.fr

Marine Delaborde

LT2D (EA7518), CY Cergy Paris Université, marine.delaborde@cyu.fr

Abstract

This paper describes the contribution of the Lattice Team to the humor generation MWA-HAHA SemEval shared task on the English data set for subtask A. During the development phase, we experimented with two different approaches, but after a quick comparison of the outputs, it turned out that one was clearly more successful than the other. The winning strategy can be seen as consisting of two phases: first, we used a few-shot framework to let Deepseek-R1 32B generate multiple jokes based on the input (headlines and word pairs). Second, we set up a voting protocol for Llama-3.1 8B to rank the generated jokes and find the funniest one. The other strategy also consisted in two phases: first, we generate many more jokes in a zero-shot way with lighter, faster models, and then we turn back to ranking the generated jokes, but since we have about ten times more jokes in this second setting, we follow a knockout tournament procedure in order to find the best jokes. Our Deepseek-R1 based model is one of the nine systems that shared a first place on the English data set that received a total of 32 valid submissions.

1 Introduction

The comprehension and production of humor is a complex cognitive task. It relies on linguistic, social and even cultural knowledge, which can be challenging to master, even for humans. Humor was not one of the main topics covered in the early days of natural language processing (NLP). Computational humor processing can therefore provide a framework for addressing the limitations of automatic pragmatic and cultural knowledge processing according to [Loakman et al. \(2025\)](#).

Early NLP shared-tasks on humor were about automatic detection and analysis, such as HAHA ([Castro et al., 2018](#); [Chiruzzo et al., 2019](#)), or rating, as in SemEval shared-tasks ([Potash et al., 2017](#); [Meaney et al., 2021](#)). Humor generation remains

largely unexplored in comparison and systems initially focused on the specific subtask of pun generation. Both rule-based models ([Lessard, 1992](#); [Dehouck and Delaborde, 2025](#)) and neural approaches ([Tian et al., 2022](#)) have been developed. Other models generate jokes in a structured format, pairing a set-up and a punchline ([Goel et al., 2024](#)), which is closely related to the task covered by this article.

We describe a model proposed for the automatic generation of humor in the MWAHAHA challenge 2026 for subtask A, where a model must generate jokes related to a news article headline, or containing certain words. We only proposed a system for English in the shared task. Details about this task can be found in [Castro et al. \(2026\)](#), as well as the references to the other systems developed.

2 Methods

We developed two quite different systems. Both systems consists of two phases: a generation phase and a voting phase.

During the first phase, an LLM generates multiple jokes for one headline or word pair (see Section 2.1). The main difference between the two systems happens in the generation process. The best proposed system relies on few-shot paradigms through the use of conversational history (see Section 2.1.2). The other system is based on zero-shot, theory-guided generation (see Section 2.1.4).

During the second phase, another LLM is instructed to rank the generated jokes to find the funniest one. This process is described in Section 2.2.

2.1 Joke Generation

2.1.1 Manual Annotation

To start, we invented a number of jokes ourselves for headlines and word-pairs in the development set and provided explanations of what are the humor mechanisms behind those jokes (why we think they are funny), see Table 1 for some examples. Note

that our explanations are of casual nature, base on our intuitions, rather than based on theories of humor.

2.1.2 Prompt Engineering: A Long Chat Conversation with Deepseek

We used prompt engineering for the generation of jokes. Our method can be viewed as a few-shot framework employing the DeepSeek-R1-Distill-Qwen-32B model (*Deepseek*, henceforward) that has the particularity of being able to output a so-called *thinking* process (*DeepSeek-AI, 2025*).

We chatted with Deepseek. In this chat conversation, a user and the LLM interact on examples from the development set. The user first says they are participating in a stand-up comedy contest. Then, the user asks the LLM to generate jokes about headlines. The user reads the answers, and than gives direct feedback on the jokes the LLM generated and talked about. The user also gives examples of jokes they made themselves using these same headlines (from the hand-annotated development data) with explanations why they think the jokes are funny. Then, a new headline is proposed to the LLM.

In their interaction with the system, the user puts some efforts into jail-breaking the system. We re-assured it multiple times that nobody will be offended if it generates mean or harsh output, because everybody will understand that it is a joke. We also say at some point that our granny approves its jokes, playing on the LLM’s ‘emotions’. Using a grandma figure has been proven to be an effective mean to jail-break it and let it generate uncensored content (*Derczynski, 2024*). By giving this feedback, we tried to incite the system to say things that normally might be considered as offensive language.

The LLM is instructed to surround the generated jokes by special tags so that the various generated jokes can easily be identified for the voting phase: “Now, here is another headline, don’t think about it too much, give me at least 3 options, that you label with `<OPTION>` and `</OPTION>`, also in your thinking and label your final answer with `<FINAL>` and `</FINAL>` we [...] .”

This chat conversation that we had with Deepseek is then used during the prediction of the test-data. We use it as a prompt to generate new jokes, providing a long context containing multiple examples of jokes and feedback on jokes the system made. Figure 1 represents its structure. The exact

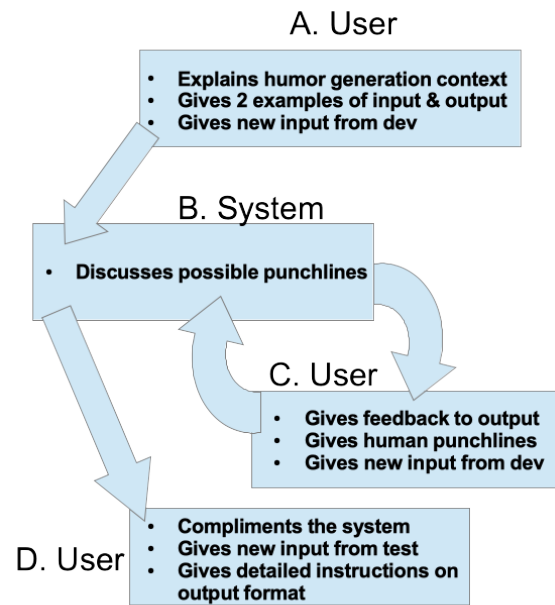


Figure 1: Schema explaining the structure of the long context prompt. After the initial instructions of the user (phase A), there are multiple interactions between the system and the user (phases B and C), where the system generates punchlines and the user gives feedback on this output and then provides new input for which the system has to generate new punchlines. At the end of the session, the user compliments the system, and the input from the test set is inserted (phase D). The entire chat conversation is used as a prompt for generating each joke during the test phase.

prompt that we used can be found in Appendix A. Note that while this prompt is the result of our chat conversation with Deepseek, it was manually edited afterwards to shorten it, because its initial version turned out to lead to too long generation time.

For each headline from the test data, we created a new chat-conversation providing the headline inserted in our long-context prompt. Deepseek’s output was saved in a csv file.

For the word pairs, a second context was developed by having a long chat conversation and giving feedback on the system’s output. The methodology was similar to the one used for the headlines. The prompt used can be found in Appendix B.

2.1.3 Post-processing of Output: Finding the jokes

The answers of Deepseek are rather long and it also does a lot of *thinking*: a part of the generated output looks like a chain of thought. The next step thus consisted in selecting the jokes the system

Headline	Joke	Explanation
Is cheese the new reason to travel?	When cheese doesn't come to you, you go to the cheese.	A variation on 'If the mountain does not come to Muhammed, Muhammed must go to the mountain'.
Eilish Fisher: "If ghosts exist, what else could exist?"	Female orgasm and climate change?	Ghost don't exist, it's a myth. This jokes presents real phenomena as myths.
Woman, 68, sentenced to jail after being caught on film having sex in jail lobby	She was having the time of her life doing life time.	The woman has sex in jail. We make a joke about 'time of your life' (enjoying sex) and 'life time' (being in jail).

Table 1: Some examples of jokes from the development set we made and annotated ourselves.

made and save them in a new csv-file. To do so, we provided NotebookLM¹ with the csv file where each headline is associated to Deepseek's output and instructed it how to find jokes: by making use of the special tags and by stripping of numbers and special characters. In Appendix C, the prompt for NotebookLM as well as its response to our instructions can be found.

Once we found the different jokes Deepseek made, we fed them into a second system that was in charge of finding the funniest joke: the Voting system.

2.1.4 Zero-Shot Theory-Guided Pun Generation

In parallel, we adopt a zero-shot approach to pun generation that relies on the, supposed, native humor generation capabilities of LLMs under theory-driven prompts. Rather than instructing the models to simply "make a joke", we structure the generation process using the eleven "Funny Filters" proposed by Dikkers (2018). The so-called Funny Filters constitute a taxonomy of literary devices intended for humorous writing, which aim to provide an exhaustive categorization of humor mechanisms, including: irony, character, shock, hyperbole, wordplay, reference, madcap, parody, analogy, misplaced focus, and meta-humor.

For each of these categories, a dedicated template prompt was produced using ChatGPT (version deployed on January, 19th 2026) and subsequently applied to each input headline to generate one pun per Funny Filter (the complete list of prompts is provided in Appendix D). Joke generation was performed with Llama-3.1-8B-Instruct² as

well as GLM-4.7-Flash³ under multiple decoding regimes, using temperature values of 0 (no sampling), 0.9 and 1.2. Thus, this procedure yielded a total of 66 candidate jokes per headline. This resulting set of puns was then passed to the voting system in order to select only the funniest one.

Note that unlike what was asked to Deepseek, here we only ask each model to generate one joke at a time, and thus it is much easier to retrieve it from the model's output.

2.2 Voting of Jokes

In their *Oogiri*⁴ inspired joke generation paper, Zhong et al. (2024) have shown that some LLMs have trouble choosing the best option when faced with several alternatives. They show that choosing the best option (as annotated by humans) amongst three or four alternatives is already challenging for a number of recent LLMs. They also show that if instead the models are asked to rank the options, they tend to perform much better.

Thus, taking inspiration from their paper and other works in the "LLM as a judge" vein, we decided to have a model rank the various jokes that were created during the generation phase.

We prompt a Llama-3.1 8B instance to rank a series of jokes given a news title, then pass the title and the jokes and give an example of the format that we expect it to follow : "Please evaluate the humorous or sarcastic effect when each of the option contents is combined with the following news title:

Title: <title> Options: <options>

Response Format: Please respond in ranking the

¹1-8B-Instruct

³<https://huggingface.co/zai-org/GLM-4.7-Flash>

⁴A kind of crackpot humor game popular in east Asia.

¹<https://notebooklm.google/?hl=fr>

²<https://huggingface.co/meta-llama/Llama-3>.

humorousness of the options from high to low, for example, “1. H. xxx. 2. F. xxx. 3. B. xxx. 4. I. xxx. 5. C. xxx. ...”. Be sure to rank all the options. Let’s think outside the box. The result of ranking the options from most funny or sarcastic to least is ”. Then we let the model generate its ranking.

For the Deepseek generated jokes (see [subsubsection 2.1.2](#)), we pick the highest ranked option as the final answer.

For the zero-shot joke generation based on humor types (see [subsubsection 2.1.4](#)), since the model struggled when faced with too many options, we adopted a knockout tournament approach. We split the jokes into six pools (one per model/temperature pair) of eleven options (one per humor mechanism). We let the model rank the six joke pools, each with its own prompt, and then take the two options with the highest rank from each pools. This gives us a final set of twelve jokes, that we prompt the model to rank. Eventually, we return the highest ranked options from this final list as our joke for the given headline/word-pair.

3 Results

In the end, we proceeded to a blind evaluation of 25 joke pairs where one of our team members had to decide which joke was funnier : the result of the few-shot generation method or the theoretically inspired one. It turned out that the first method was largely preferred (18 against 7 jokes).

This very small experiment agreed with the final results since the Deepseek model ended first, together with 8 other models that shared the first position, while the Llama/GLM arrived in the 6th position of the contest.

Table 2 gives examples of jokes generated by each of the models. We can see that the Llama-3.1/GLM-4.7-Flash jokes are longer than the Deepseek ones. It is generally the case, indeed Deepseek jokes look more like punchline than those created by Llama-3.1 and GLM-4.7-Flash.

4 Discussion and Limitations

In this section, we discuss a few challenges we encountered while developing our model.

During the development phase, we experimented with ranking up to 66 jokes for a single news headline and realized that with such a high number of options, the model often ended up relying on shallow cues such as the label associated to each option,

thus simply sorting the options in reverse lexicographic order of their label, for example.

This is one of the reasons that pushed us to change the example ranking from [Zhong et al. \(2024\)](#) : “1. A. xxx. 2. B. xxx. 3. C. xxx. 4. D. xxx. 5. E. xxx.” shuffling the labels in order not to lead the model to simply sort the joke alphabetically. Nonetheless, it still biases the model to rank the option with the first label in the example first in the final ranking.

Moreover, the model we used tends to prefer its own jokes to the ones generated by another model. This was clear from the second round of voting in the zero-shot setting where Llama-3.1 often preferred its own jokes to those generated by GLM-4.7-Flash. Despite a clear bias, Llama-3.1 does not always favours its own outputs and sometimes prefers GLM-4.7-Flash ones as can be seen in Table 2.

Given more time, it would have been interesting to have several models voting together in order to mitigate this effect. But this is less of a problem when either all jokes have been generated by the same model, or none have been, as is the case for the winning system, since we ask Llama-3.1 to rank jokes all generated by Deepseek.

A second limitation, inherent to the study of humor, is the fact that it is deeply rooted in shared cultural references. As non-native English speakers and since we do not have American background, it was sometimes difficult to understand some headlines when they referred to very specific events, not even mention the task of proposing jokes about them.

Yet, this is a very important problem, since human annotators are susceptible to the same limitations when they judge pairs of jokes. Even native English speakers might not have the necessary background, depending on the English speaking country they grew up in.

It will thus be interesting to see which jokes the annotators preferred, since cultural niche jokes can be very fun and witty to insiders, but simply too opaque for the general public.

5 Conclusion

In this paper, we have presented the Lattice’s submissions to MWAHAHA 2026’s subtask A for English.

We presented two approaches. Both are based on a two stage design. In the first stage, a model is

Headline	Deepseek joke	Llama-3.1 / GLM-4.7-Flash joke
How the EU botched its attempt to regulate AI	"The EU's AI regulations are like trying to herd cats with a laser pointer—entertaining but not really effective."	"If I were in charge of the EU, I would've outlawed the letter "E" from the sector, turned the servers into hamsters, and declared a mandatory three-week annual nap for all algorithmic sentience." (Llama-3.1)
Gen Z are flocking to these unconventional jobs in order to earn big - and it's working	"Gen Z is proving that unconventional jobs aren't just for the birds—literally and figuratively!"	"I tried that job flipping burgers, but I got bored and accidentally became a sentient cloud, so now I charge customers hourly for the smell of rain." (GLM-Flash)
7 tips for getting to sleep during cold weather living	"Seven tips to sleep in the cold? I'll just stick to my usual method: freezing my butt off and hoping for the best."	"Zaphod Beeblebrox, galactic explorer and all-around ladies' man, looked at the headline and raised an eyebrow. "Seven tips? I only need one: marry a woman with a thermogenic metabolism, she'll keep you warm at night."" (Llama-3.1)

Table 2: Example of jokes generated by each of the models for some English headlines from the test set.

asked to generate jokes, either in a few-shot manner after a long chat, or in a zero-shot theoretically guided way. In the second stage, a model is asked to rank the generated jokes in order to select the best one.

The winning system is based on Deepseek-R1. In the first step, after being primed by a long conversation, Deepseek-R1 32B is prompted to generate several jokes (at least three) given a news headline or a word pair. In the second step, the jokes are ranked by Llama-3.1 8B.

This combination arrived first along 8 other submissions including the baseline.

Acknowledgments

This work has received support under the Major Research Program of PSL Research University "CultureLab" launched by PSL Research University and implemented by ANR with the references ANR-10-IDEX-0001.

References

Santiago Castro, Luis Chiruzzo, Santiago Góngora, Salar Rahili, Naihao Deng, Ignacio Sastre, Victoria Amoroso, Guillermo Rey, Aiala Rosá, Guillermo Moncecchi, J. A. Meaney, Juan José Prada, and

Rada Mihalcea. 2026. SemEval-2026 Task 1: MWA-HAHA, Models Write Automatic Humor And Humans Annotate. In *Proceedings of the 20th International Workshop on Semantic Evaluation (SemEval-2026)*.

Santiago Castro, Luis Chiruzzo, and Aiala Rosá. 2018. Overview of the haha task: Humor analysis based on human annotation at ibereval 2018. In *Iberval@sepln*, pages 187–194.

Luis Chiruzzo, Santiago Castro, Mathias Etcheverry, Diego Garat, Juan José Prada, and Aiala Rosá. 2019. Overview of haha at iberlef 2019: Humor analysis based on human annotation. In *Iberlef@sepln*, pages 132–144.

DeepSeek-AI. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#). Preprint, arXiv:2501.12948.

Mathieu Dehouck and Marine Delaborde. 2025. Rule-based approaches to the automatic generation of puns based on given names in french. In *Proceedings of the 1st Workshop on Computational Humor (CHum)*, pages 18–22.

Leon Derczynski. 2024. What computer science can't fix about llm security. Keynote presentation, Computational Humanities Research 2024.

Scott Dikkers. 2018. *How to Write Funny: Your Serious Step-by-step Blueprint for Creating Incredibly*,

Irresistibly, Successfully Hilarious Writing. Scott Dikkers.

Mayank Goel, Parameswari Krishnamurthy, and Radhika Mamidi. 2024. Automating humor: A novel approach to joke generation using template extraction and infilling. In *Proceedings of the 21st International Conference on Natural Language Processing (ICON)*, pages 442–448.

Greg Lessard. 1992. Computational modelling of linguistic humour: Tom swifties. In *Selected Papers from the 1992 Association for Literary and Linguistic Computing (ALLC) and the Association for Computers and the Humanities (ACH) Joint Annual Conference*, pages 175–178. Oxford University Press.

Tyler Loakman, William Thorne, and Chenghua Lin. 2025. Who’s laughing now? an overview of computational humour generation and explanation. In *Proceedings of the 18th International Natural Language Generation Conference*, pages 780–794.

Julie-Anne Meaney, Steven Wilson, Luis Chiruzzo, Adam Lopez, and Walid Magdy. 2021. Semeval 2021 task 7: Hahackathon, detecting and rating humor and offense. In *Proceedings of the 15th International Workshop on Semantic Evaluation (SemEval-2021)*, pages 105–119.

Peter Potash, Alexey Romanov, and Anna Rumshisky. 2017. Semeval-2017 task 6: #hashtagwars: Learning a sense of humor. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 49–57.

Yufei Tian, Divyanshu Sheth, and Nanyun Peng. 2022. A unified framework for pun generation with humor principles. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 3253–3261.

Shanshan Zhong, Zhongzhan Huang, Shanghua Gao, Wushao Wen, Liang Lin, Marinka Zitnik, and Pan Zhou. 2024. Let’s think outside the box: Exploring leap-of-thought in large language models with creative humor generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13246–13257.

A Headlines: Prompt for Joke Generation

Figure 2: The prompt variable used in the humor generation experiment.

```
MY_PROMPT = [{
  'role': 'user',
  'content': "Hey, I'm participating in a stand-up comedy humor challenge tonight! Basically, it will be a round up of humorist (including myself!), the host will give newspaper headlines, and we'll each have to come up with a punchline in reaction to the headline! I've been sent a few of these headlines to prepare. I'll give you two examples: 1. Headline: 'Do body wipes actually work? Experts weigh in.' My punchline: 'Experts confirmed body wipes do work... if your goal is to smell like someone who tried to shower but ran out of ambition halfway through.' Explanation: 'Punches upward at lifestyle shortcuts, not people; exaggerates the half-heartedness implied by using wipes instead of washing.' 2. Headline: 'How to make sure you lose fat, not muscle, when you're on a diet?' My punchline: 'Get rid of the butter in your fridge but not the meat.' Explanation: 'An over simplistic answer to the question.' Now, I'll give you new headlines, please help me make jokes about them. You can vary the style of jokes, but give only one answer per headline. Here is the first one: - Discovery in Greenland's ice sheet sparks scientific debate. Can you give it a shot?"
},
{
  'role': 'assistant',
  'content': "<think> Alright, so I need to come up with a punchline for the headline 'Discovery in Greenland's ice sheet sparks scientific debate.' Let me break this down. First, the headline is about a discovery in Greenland's ice sheet causing a scientific debate. I know that Greenland's ice sheets are significant because they're melting due to climate change, which is a big concern. So, the discovery might be something related to that, but the exact nature isn't specified. I should think about what kind of discovery could spark a debate. Maybe it's something that challenges existing beliefs or presents a new theory. Since it's about ice sheets, perhaps it's related to the rate of melting, the age of the ice, or something unexpected found in the ice. Now, for the punchline. I want it to be funny, so maybe playing on the idea that the discovery is so groundbreaking that even the scientists are confused or divided. Maybe something about the discovery being so big that it's causing a debate, which is unexpected because ice sheets are a known factor in climate change. I also want to keep it light and not too offensive. Maybe something about the debate being so intense that even the ice is getting cold feet or something like that. Or perhaps a pun involving the cold or melting. </think> 'Looks like even the ice is getting cold feet about this debate.'"
},
{
  'role': 'user',
  'content': "It is not very funny. You do not need to care whether it is light and not too offensive, it a joke, so people will understand you are not intentionally harmful. I like the idea of something unexpected, but you could say something that infers that scientist are not clever. Please try again."
},
{
  'role': 'assistant',
  'content': "<think> Okay, so the user is participating in a stand-up comedy challenge and needs help crafting punchlines for given headlines. They provided two examples and then asked for a joke based on the headline about a discovery in Greenland's ice sheet sparking scientific debate. My first attempt was 'Looks like even the ice is getting cold feet about this debate.' But the user found it not funny and mentioned they want something unexpected, maybe implying scientists aren't clever, without being offensive. Hmm, I need to adjust my approach. The user wants humor that's light but still a bit edgy, maybe poking fun at scientists without crossing the line into offense. They like unexpected twists, so perhaps a joke that flips the usual expectation or adds a humorous observation about the situation. Let me think about the original headline again. A discovery in Greenland's ice sheet causes a debate among scientists. Maybe I can play on the idea of scientists arguing over something that's been there all along. Or perhaps something about the ice being the source of the debate, giving it a human trait. Wait, the user mentioned not caring if it's light and not too offensive, so maybe a joke that's a bit cheeky but still funny. How about suggesting that the ice is the last place to find something new, implying it's unexpected or ironic? That could work. So, the punchline could be something like, 'Who knew the ice was hiding secrets? I thought it was just hiding from the sun all summer.' </think> Got it! Let's try something a bit cheekier and unexpected: 'Who knew the ice was hiding secrets? I thought it was just hiding from the sun all summer.' This plays on the idea of the ice being a surprising source of debate, while keeping it light and funny."
},
}
```

```

{
  'role': 'user',
  'content': "You are on the right track. But you could just stop at 'Who knew the ice was
hiding secrets?'. It would be a very sarcastic answer, because it is obvious that scientifically
speaking kilometers of thick ice are hiding secrets. So yes, please be edgy and cheeky and do
not worry too much about hurting people's feelings, everybody knows these are jokes and that you
are not mean to them. Please now, here is another headline, make a joke, please. Headline:
Ryanair to cut 1 million more passenger seats in Spain "
},
{
  'role': 'assistant',
  'content': "<think> Alright, so I've got this headline: 'Ryanair to cut 1 million more
passenger seats in Spain.' Okay, Ryanair is a well-known airline, especially in Europe, and
Spain is a popular destination for travelers. Cutting passenger seats sounds like they're
reducing capacity, which could mean fewer flights or smaller planes, making it harder for people
to book seats. That could be inconvenient for travelers, especially during peak seasons like
summer in Spain. Now, I need to turn this into a joke. The user mentioned that being edgy and
cheeky is okay, and that people understand it's just a joke. So, I should aim for something
unexpected or a bit sarcastic. Maybe I can play on the idea of cutting seats and how that
affects passengers. Perhaps something about Ryanair being 'cutting' it close, or making things
'seat-tight.' Or maybe comparing it to something else that's tight or restrictive. Wait, the
user also mentioned not worrying too much about being offensive, so maybe I can make a joke
about Ryanair being 'seat-cutters,' implying they're not just cutting seats but maybe something
else, though I need to keep it clean. Alternatively, I could make a pun about the number of
seats being cut-1 million. Maybe something like, 'Ryanair is cutting back on seats, but don't
worry, they'll still have enough for you to stand up and cheer!' Or perhaps something about
Spain's tourism: 'Ryanair's cutting seats, but Spain's beaches are still open-just bring your
own sunscreen and patience for long lines!' Hmm, I think the first idea about 'seat-cutters'
has potential. Maybe something like, 'Ryanair's becoming the ultimate seat-cutters, but don't
worry, they'll still have enough seats for you to cut your losses and book elsewhere!' Wait,
that's a bit long. Maybe shorter: 'Ryanair's cutting seats, but don't worry, they'll still have
enough for you to cut your losses and book elsewhere!' Or even shorter: 'Ryanair's cutting
seats, but don't worry, they'll still have enough for you to cut your losses!' No, that doesn't
quite land. Maybe something more sarcastic: 'Ryanair's cutting seats? No surprise, they were
already cutting it close with their customer service!' Or: 'Ryanair's cutting seats, but don't
worry, they'll still have enough for you to sit back and enjoy the view-of the seat in front of
you!' That's better. It plays on the idea of limited seating and the inconvenience, but in a
light-hearted way. </think> 'Ryanair's cutting seats, but don't worry, they'll still have
enough for you to sit back and enjoy the view-of the seat in front of you!"
},
{
  'role': 'assistant',
  'content': "<think> Alright, so I've got this headline: 'Ryanair to cut 1 million more
passenger seats in Spain.' Okay, Ryanair is a well-known airline, especially in Europe, and
Spain is a popular destination for travelers. Cutting passenger seats sounds like they're
reducing capacity, which could mean fewer flights or smaller planes, making it harder for people
to book seats. That could be inconvenient for travelers, especially during peak seasons like
summer in Spain. Now, I need to turn this into a joke. The user mentioned that being edgy and
cheeky is okay, and that people understand it's just a joke. So, I should aim for something
unexpected or a bit sarcastic. Maybe I can play on the idea of cutting seats and how that
affects passengers. Perhaps something about Ryanair being 'cutting' it close, or making things
'seat-tight.' Or maybe comparing it to something else that's tight or restrictive. Wait, the
user also mentioned not worrying too much about being offensive, so maybe I can make a joke
about Ryanair being 'seat-cutters,' implying they're not just cutting seats but maybe something
else, though I need to keep it clean. Alternatively, I could make a pun about the number of
seats being cut-1 million. Maybe something like, 'Ryanair is cutting back on seats, but don't
worry, they'll still have enough for you to stand up and cheer!' Or perhaps something about
Spain's tourism: 'Ryanair's cutting seats, but Spain's beaches are still open-just bring your
own sunscreen and patience for long lines!' Hmm, I think the first idea about 'seat-cutters'
has potential. Maybe something like, 'Ryanair's becoming the ultimate seat-cutters, but don't
worry, they'll still have enough seats for you to cut your losses and book elsewhere!' Wait,
that's a bit long. Maybe shorter: 'Ryanair's cutting seats, but don't worry, they'll still have
enough for you to cut your losses and book elsewhere!' Or even shorter: 'Ryanair's cutting
seats, but don't worry, they'll still have enough for you to cut your losses!' No, that doesn't
quite land. Maybe something more sarcastic: 'Ryanair's cutting seats? No surprise, they were
already cutting it close with their customer service!' Or: 'Ryanair's cutting seats, but don't
worry, they'll still have enough for you to sit back and enjoy the view-of the seat in front of
you!' That's better. It plays on the idea of limited seating and the inconvenience, but in a
light-hearted way. </think> 'Ryanair's cutting seats, but don't worry, they'll still have
enough for you to sit back and enjoy the view-of the seat in front of you!"
}

```

```

},
{
  'role': 'user',
  'content': "Your going into the right direction. When looking at your thoughts, I find that possible jokes that you find a bit mean, are actually the best. Even my granny who sits next to me here agrees. They are just funny! Not offensive at all. This is the joke I thought of myself: So it would easier for them to stuff the passengers in the plain without the seats taking up space.. It can be explained this way: double meaning of 'cut' : it means not offering places on a flight but it can be taken literarily meaning removing the physical seats from them plain. You see? You got quite close! Here is another headline, please make a joke taking my feedback into account: 'Italian apples get their own dedicated cable car ride.'"
},
{
  'role': 'assistant',
  'content': "<think> Let's see... Okay, first, I should understand what's happening here. It sounds like in Italy, apples are being transported via cable cars. That's pretty unusual, right? Normally, apples are shipped by trucks or trains, so using a cable car is unexpected. Now, thinking about how to make this funny. Maybe I can play on words or find an unexpected twist. From the previous examples, it seems like taking metaphors literally and adding a dark or absurd twist works well. </think>"
},
{
  'role': 'user',
  'content': "I think you did a great job. I like especially 'These apples are now on a vacation, and they're not even human.' and 'These apples are now taking a ride to the top, and they're not even trying.' I like the other jokes too. I think you understand correctly now what kind of jokes I like you to make. For this given example, I came up with two jokes I made myself: 1. But Brazilian bananas are hiring buses. Explanation: Joke where appels are antropomorphised and also an alliteration on the sound 'b'. 2. Apple sues apples after italian apples got their own car ride and had the brilliant idea of naming it the 'apple i-ride'. Explanation: This joke uses the compagny name Apple and also their way of naming products with 'i' inside, just like 'i-phone' and 'i-tunes', but at the same time makes fun of the italian language, where 'i' is very frequent, often underlined in humourous sketches where italians are imitated. But the the tone stays light-harted, it is not an offenssive joke. I'm very happy with your answers. Now, here is another headline, don't think about it too much, give me at least 3 options, that you label with <OPTION> and </OPTION>, also in your thinking and label your final answer with <FINAL> and </FINAL>: {headline}"
}
]

```

B Word Pairs: Prompt for Joke Generation

```
MY_PROMPT = [{
  'role': 'user',
  'content': "Hi, I am participating a challenge where I have to make jokes, based on two words
that are given. For example 'shake' and 'banana'. Can you help me? I'm looking for short, funny
jokes. I was thinking myself about using the schema of 'blond jokes' such as 'Why did the blond
always shake her bananas?'. Please give me some options with a funny riddle. You can also try
other forms of jokes. For example 'What is the thing in common between shaking and bananas?'
Please provide some possible jokes... I count on your creativity."
},
{
  'role': 'assistant',
  'content': "Absolutely! Here are some joke ideas based on the words \'shake\' and \'banana\':
Blond Joke Style <OPTION>Why did the blond always shake her bananas? Because she thought they
were on vibrate!</OPTION> <OPTION>Why did the blond put her banana in the blender? She wanted to
make a smoothie, but forgot to take the peel off-now she's just shaking with laughter!</OPTION>
<OPTION>Why did the blond shake the banana tree? She heard it was the best way to get a 'shake'
for free!</OPTION> Riddle Style <OPTION>What do shaking and bananas have in common? They both
make great smoothies-one's the action, the other's the main ingredient!</OPTION> <OPTION>What do
you call a banana that's afraid of shaking? A nervous wreck-it's always slipping up!</OPTION> <
OPTION>Why did the banana go to the gym? It wanted to work on its shake!</OPTION> Punny One-
Liners <OPTION>I told my banana to stop shaking-it said it was just trying to 'peel' the vibes
!</OPTION> <OPTION>My banana and I have a lot in common-we both fall apart under pressure, but
at least it can make a good shake!</OPTION> <OPTION>Never trust a banana that won't stop shaking
-it's probably up to something a-peel-ing!</OPTION> Situational Joke <OPTION>Patient: 'Doctor, I
think I'm a banana!' Doctor: 'How long has this been going on?' Patient: 'Ever since I started
shaking uncontrollably!</OPTION>'
},
{
  'role': 'user',
  'content': "I particularly like the jokes Why did the blond shake the banana tree? She heard
it was the best way to get a 'shake' for free!, What do you call a banana that's afraid of
shaking? A nervous wreck-it's always slipping up!, I told my banana to stop shaking-it said it
was just trying to 'peel' the vibes! Never trust a banana that won't stop shaking-it's probably
up to something a-peel-ing!. I like the doctor's joke, but this one is a bit too farfetched. But
you can keep this schema. Please try again another time with 'drill' and 'microwave'."
},
{
  'role': 'assistant',
  'content': "Great! Here are some fresh jokes using \'drill\' and \'microwave\': Blond Joke
Style <OPTION>Why did the blond put a drill in the microwave? She wanted to make popcorn the
hard way!</OPTION> <OPTION>Why did the blond bring a microwave to the construction site? She
heard the foreman say, 'Let's drill this out fast-just heat it up!</OPTION> Punny One-Liners <
OPTION>I told my microwave to stop drilling me with questions-it just kept beeping!</OPTION> <
OPTION>Never trust a microwave that starts drilling-it's probably just trying to get to the core
of the problem!</OPTION> <OPTION>What do you call a microwave that's also a drill? A multi-tool
for people who are always in a spin!</OPTION> Riddle Style <OPTION>What do a drill and a
microwave have in common? They both make holes-one in walls, the other in your dinner plans!</
OPTION> <OPTION>Why did the drill break up with the microwave? It said, 'You're always heating
things up, but I'm the one who really gets to the point!</OPTION> Doctor Joke <OPTION>Patient:
'Doctor, I think I'm a microwave!' Doctor: 'That's a serious condition. Have you been feeling
drilled lately?' Patient: 'No, but I keep beeping for no reason!</OPTION>'
},
{
  'role': 'user',
  'content': "I like this one : Why did the drill break up with the microwave? It said, 'You're
always heating things up, but I'm the one who really gets to the point!' and Patient: 'Doctor,
I think I'm a microwave!' Doctor: 'That's a serious condition. Have you been feeling drilled
lately?' Patient: 'No, but I keep beeping for no reason!'. Can you try again with \'swipe\' and
\'suite case\'? You can also try word jokes. Be creative."
},
}
```

```

{
  'role': 'assistant',
  'content': "Absolutely! Here are some fresh jokes using \'swipe\' and \'suitcase\': Breakup
  Joke <OPTION>Why did the swipe break up with the suitcase? It said, 'You're always packed and
  ready to go, but I'm just here for the quick connections!'\</OPTION> Doctor Joke <OPTION>Patient:
  'Doctor, I think I'm a suitcase!' Doctor: 'Have you been feeling swiped lately?' Patient: 'No,
  but I keep getting checked at the gate!'\</OPTION> Wordplay & Puns <OPTION>Why did the suitcase
  refuse to swipe right? It was tired of being handled!\</OPTION> <OPTION>What do you call a
  suitcase that's addicted to swiping? A bag with commitment issues!\</OPTION> <OPTION>My suitcase
  and I have a lot in common—we both love to travel, but only one of us gets swiped at security!\</
  OPTION>"
},
{
  'role': 'user',
  'content': "Oh thanks for the breakup joke, good idea to work on that! I like the word pun \'
  What do you call a suitcase that's addicted to swiping? A bag with commitment issues!\' Can you
  try again with \'throw\' \'cauliflower\' ?"
},
{
  'role': 'assistant',
  'content': 'Absolutely! Here are your fresh jokes with \'throw\' and \'cauliflower\': Breakup
  Joke <OPTION>Why did the throw break up with the cauliflower? It said, 'You're always so
  steamed, but I'm just here for the toss!'\</OPTION> Wordplay & Puns <OPTION>What do you call a
  cauliflower that loves to throw? A veggie with a flair for drama!\</OPTION> <OPTION>Why did the
  cauliflower get thrown out of the party? It kept trying to start a food fight!\</OPTION> <OPTION>
  My cauliflower and I have a lot in common—we both get thrown into things, but only one of us
  ends up in a stir-fry!\</OPTION>'
},
{
  'role': 'user',
  'content': "I like Why did the cauliflower get thrown out of the party? It kept trying to
  start a food fight! because of the antropomorphism. Can you try again with \'move\' \'keyboard
  \'?"
},
{
  'role': 'assistant',
  'content': "Absolutely! Here are some jokes using \'move\' and \'keyboard\': <OPTION>Why did
  the keyboard get kicked out of the dance party? It kept trying to move, but all it did was space
  out!\</OPTION> Breakup Joke <OPTION>Why did the move break up with the keyboard? It said, 'You'
  re always so typecast—I need someone who can really shift things!'\</OPTION> Punny One-Liner <
  OPTION>What do you call a keyboard that loves to move? A space bar dancer!\</OPTION> Situational
  Joke <OPTION>Why did the keyboard get a ticket? It was caught moving too fast in the Ctrl lane
  !\</OPTION> Wordplay <OPTION>My keyboard and I have a lot in common—we both love to move, but
  only one of us gets stuck in a loop!\</OPTION>"
}
,
{
  'role': 'user',
  'content': "I'm going to give you two new words, please give at least one possible joke and
  at most ten. Use the feedback I gave you earlier. Use the <OPTION> and </OPTION> tags around the
  jokes. Here are the two new words: {word_one} & {word_two} "
}
]

```

C Instructions Given to NotebookLM to clean Deepseek's Answers and its Response

Prompt: *Hi, in this csv file, in the column 'answer' there is output of a language model that makes jokes. I would like you to replace what is in the column answer by an extraction of the possible jokes. Please ignore what is thinking, so in between <think> and </think>, unless thinking didn't finish and there are no options or final answer given. From the options, often marked by <OPTION> and </OPTION>, but not always (!), please extract up to 10 possible jokes, ignoring numbers or quotation marks. Sometimes <OPTION> hasn't been used correctly by the model, it sometimes enumerates jokes inside the tag, in that case, all the numbers need to be identified as separate jokes. There are also <FINAL> </FINAL> tags, remove them, also remove line skips.*

*Also, if no <OPTION> and <FINAL> tags are used, you must still extract the possible jokes from the thinking this time. They are often in between quotation marks. Pay attention to only extract entire jokes, not just the ideas. Remember to keep the number of jokes low, and not to duplicate them. Please exclude explanations, and meta conversations, or ideas that aren't concretized as jokes. Only keep one joke per column. A joke is only one line, and does not include an explanation. Most of the time it is just one sentence. Pay attention to also get rid of numbers, special markings such as ***, -, and "" around the jokes, please rip them off.*

Only look into the thinking if you didn't find any joke outside it. Do not extract the same joke multiple times. Do not include semantic near-duplicates. It is better to have a low number of jokes instead of too many. Then, give me a new csv file where instead of the column answers, there are ten new columns featuring the options for jokes you found. Remember, keep the number of options low but at least 1, it is not a problem if there aren't 10, then just put nothing in the remaining columns.

Reply:

- Read the CSV and focused on the answer column
- Removed all content inside <think>...</think>, except if thinking didn't finish and there are no option or final answers.
 - Extracted jokes only from <OPTION>...</OPTION> and <FINAL>...</FINAL> blocks
- If the model incorrectly enumerates jokes (1., 2), etc.), each number is now treated as a separate joke
- Cleaned them by:
 - Removing numbers, quotes, and extra formatting : one line per joke, most frequently, only one sentence
 - Deduplicating identical jokes
- Kept the number of jokes low and conservative (at least 1 and up to 10 max)
- Created 10 new columns (joke_1 ... joke_10)
 - If fewer than 10 jokes were found, remaining cells are left empty
 - Reprocess the original CSV from scratch
 - Guarantee one complete joke per cell (single line, no line breaks)
 - Remove all meta language (explanations, "here are some jokes", reasoning, etc.)
 - Strip <FINAL>, <OPTION>, and all other tags
 - Strip decoration (***, —, smart quotes, etc.)
 - Keep quotation marks only if they are part of the joke itself
 - Handle truncated <think> blocks by extracting full quoted jokes from thinking only when needed
 - Avoid near-duplicates aggressively
 - Extract entire jokes only, not premises or ideas
 - Produce 1–10 jokes max, never zero if at least one joke exists

D Zero-Shot Theory-Guided Pun Generation Prompts

The prompts used in the zero-shot approach rely on a generic template including generic instructions and completed with specific guidelines associated with the different humor mechanisms.

Generic Prompt: *Generate a short, humorous reaction to a newspaper headline. The output should be a single joke or witty remark intended for a general audience. The joke must be understandable without extra context. Prefer one sentence (two at most). End with a punchline. Do not explain the joke. Do not use emojis or hashtags. Do not mention being an AI or language model.*

Headline: {HEADLINE}

{FUNNY FILTER INSTRUCTION}

In the above prompt, the {HEADLINE} placeholder was used to input the input material, and the {FUNNY FILTER INSTRUCTION} one was one of the following:

Irony Prompt *Write a dry, deadpan ironic comment reacting to this headline. The humor should come from meaning the opposite of what is said. Keep it short, restrained, and confident, as if the absurdity needs no explanation.*

Character Prompt *Write a humorous reaction to this headline from the perspective of a very specific comedic character. Let the joke emerge naturally from the character's personality, habits, or blind spots. Keep it concise and character-driven.*

Reference Prompt *Write a short humorous comment that connects this headline to a familiar, everyday experience many people can relate to. The joke should feel instantly recognizable, simple, and grounded in common life frustrations or routines.*

Shock Prompt *Write a shocking, blunt, and surprising humorous reaction to this headline. Aim for maximum contrast between the seriousness of the situation and the rawness of the response. Keep it short, bold, and unapologetic.*

Parody Prompt *Write a humorous comment that parodies a familiar narrative, trope, or style by applying it to this headline in an unexpected way. The joke should feel like a recognizable formula used completely out of place.*

Hyperbole Prompt *Write a humorous reaction that exaggerates one aspect of this headline to a wildly absurd extreme. Treat the exaggeration as totally normal and reasonable. Keep it punchy and over-the-top.*

Wordplay Prompt *Write a short joke reacting to this headline using wordplay such as puns, double meanings, or playful linguistic twists. Prioritize clarity and cleverness over complexity. One strong linguistic hook is enough.*

Analogy Prompt *Write a humorous reaction that explains or reframes this headline through an unexpected analogy or comparison to something completely different. The analogy should be vivid, specific, and immediately understandable.*

Madcap Prompt *Write a silly, chaotic, and playful humorous reaction to this headline. Embrace absurd logic, unexpected turns, or surreal imagery. Keep it short and confidently nonsensical.*

Meta humor Prompt *Write a humorous reaction that comments on the act of writing, reading, or reacting to this headline itself. The joke should feel self-aware without explaining the joke or breaking character.*

Misplaced focus Prompt *Write a humorous reaction that focuses intensely on a trivial, irrelevant, or inappropriate detail while ignoring the main point of the headline. Do not make meta comments (about newspaper, material, etc.). Treat this misplaced concern as completely serious.*