Homophonic Pun Generation in Code Mixed Hindi-English

1st Workshop on Computational Humor
COLING 2025

Yash Sarrof

UNIVERSITÄT DES SAARLANDES

Example of Code Mixing

Red - English ; **Blue - Hindi**

Create a new alarm for 9AM on monday 18th June	Monday 18th June ko subah 9 bajhe ke liye ek naye alarm ko create karen
Remove Jim from my reminder to party next wednesday	Agle wednesday ko party ke liye Jim ko mere reminder se hata den

Picture credit - Agarwal, Anmol, et al. "CST5: Data Augmentation for Code-Switched Semantic Parsing." *Proceedings of the 1st Workshop on Taming Large Language Models: Controllability in the era of Interactive Assistants!*. 2023.

UNIVERSITÄT DES SAARLANDES

Pun Generation proficiency in LLMs : Funny but not Creative

Why was the math book sad?
Because it had too many problems. (121)

Why don't scientists trust atoms?
Because they make up everything. (119)

~ 90 % (of 1000) of the generated samples were the same 25 jokes.

Picture Credit : Jentzsch, Sophie, and Kristian Kersting. "ChatGPT is fun, but it is not funny! Humor is still challenging Large Language Models." Proceedings of the 13th Workshop on Computational Approaches to Subjectivity, Sentiment, & Social Media Analysis. 2023.

UNIVERSITÄT DES SAARLANDES

Puns in Hinglish ?

Puns in Hinglish ?

- Goldmine for homophonic puns

UNIVERSITÄT
DES
SAARLANDES

Puns in Hinglish ?

- Goldmine for homophonic puns
- Past work
  - N-grams, smoothing, word embeddings [1]
  - Encoder-decoder architectures [2]

UNIVERSITÄT DES SAARLANDES

Puns in Hinglish ?

- Goldmine for homophonic puns
- Past work
  - N-grams, smoothing, word embeddings [1]
  - Encoder-decoder architectures [2]
- But how do LLMs do ?

UNIVERSITÄT DES SAARLANDES

Puns in Hinglish ?

- Goldmine for homophonic puns
- Past work
  - N-grams, smoothing, word embeddings [1]
  - Encoder-decoder architectures [2]
- But how do LLMs do ?

[1] Aggarwal, Srishti, Kritik Mathur, and Radhika Mamidi. "Automatic Target Recovery for Hindi-English Code Mixed Puns." *arXiv preprint arXiv:1806.04535* (2018).
[2] Agarwal, Kaustubh, and Rhythm Narula. "Humor generation and detection in code-mixed Hindi-English." *Proceedings of the Student Research Workshop Associated with RANLP 2021*. 2021.

Puns in Hinglish ?

- Goldmine for homophonic puns
- Past work
  - N-grams, smoothing, word embeddings [1]
  - Encoder-decoder architectures [2]
- But how do LLMs do ?

Vanilla Prompting Example

Input: Generate a Hindi English mixed pun.
Output: Why did the Hindi teacher bring a ladder to class?
Because he wanted to teach बंदर [bəndə̤r] (bandar), how to climb the बांस [bã:s] (bamboo).

[1] Aggarwal, Srishti, Kritik Mathur, and Radhika Mamidi. "Automatic Target Recovery for Hindi-English Code Mixed Puns." *arXiv preprint arXiv:1806.04535* (2018).
[2] Agarwal, Kaustubh, and Rhythm Narula. "Humor generation and detection in code-mixed Hindi-English." *Proceedings of the Student Research Workshop Associated with RANLP 2021*. 2021.

Puns in Hinglish ?

- Goldmine for homophonic puns
- Past work
  - N-grams, smoothing, word embeddings [1]
  - Encoder-decoder architectures [2]
- But how do LLMs do ?

Vanilla Prompting Example

Input: Generate a Hindi English mixed pun.
Output: Why did the Hindi teacher bring a ladder to class?
Because he wanted to teach बंदर [bənd̪ər] (bandar), how to climb the बांस [bã:s] (bamboo).

Grammatical but Incoherent, and definitely not funny !

[1] Aggarwal, Srishti, Kritik Mathur, and Radhika Mamidi. "Automatic Target Recovery for Hindi-English Code Mixed Puns." *arXiv preprint arXiv:1806.04535* (2018).
[2] Agarwal, Kaustubh, and Rhythm Narula. "Humor generation and detection in code-mixed Hindi-English." *Proceedings of the Student Research Workshop Associated with RANLP 2021*. 2021.

Self Refine Framework

- Iteratively (3 times)
  - Assess with criterion
  - Feed assessment back
  - Refine the response

Criteria for SelfRefine

Pun Present: Does the text have a pun?
Algorithm followed: Was the algorithm described, if any followed?
Coherence: Is the text coherent?
Funny: Is the text funny?

Madaan, Aman, et al. "Self-refine: Iterative refinement with self-feedback." *Advances in Neural Information Processing Systems* 36 (2024).

Chain of Thought + Algorithmic prompting

Instruction: Construct a code-mixed Hindi-English pun based on the steps below.

Steps:
Step 1. Create a triplet of:
a. English word **EN**
b. Hindi word **HI**, homophone to **EN**
c. **HI** translated to English, labeled **HIToEN**
Step 2. If **HIToEN** == **EN**, redo step 1. Otherwise, proceed to step 3
Step 3. Construct short sentences (less than 10 words) with **EN** as the object of the sentence.
Step 4. Replace **EN** with **HI**.
Step 5. Replace the noun phrase at the start of the sentence with a contextualized phrase that is closely related to the **HIToEN** word.

Example:
Step 1: EN: **dude**, HI: दूध [d̪ud̪ʱə] (milk), HIToEN: **milk**
Step 2: Since **milk** != **dude**, proceed to step 3.
Step 3: Construct sentence with **EN** as the object: "Jack asked - What's up dude?"
Step 4: Replace **EN** with **HI**: "Jack asked - What's up दूध [d̪ud̪ʱə] (dude)?"
Step 5: Replace the noun phrase with a contextualized phrase: "The cow asked - What's up दूध [d̪ud̪ʱə] (dude)?"

He, He, Nanyun Peng, and Percy Liang. "Pun Generation with Surprise."
Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers). 2019.

Chain of Thought + Algorithmic prompting

- Algorithm based on He et al 2019

Instruction: Construct a code-mixed Hindi-English pun based on the steps below.

Steps:
Step 1. Create a triplet of:
a. English word **EN**
b. Hindi word **HI**, homophone to **EN**
c. **HI** translated to English, labeled **HIToEN**
Step 2. If **HIToEN** == **EN**, redo step 1. Otherwise, proceed to step 3
Step 3. Construct short sentences (less than 10 words) with **EN** as the object of the sentence.
Step 4. Replace **EN** with **HI**.
Step 5. Replace the noun phrase at the start of the sentence with a contextualized phrase that is closely related to the **HIToEN** word.

Example:
Step 1: EN: **dude**, HI: दूध [d̪ud̪ʱə] (milk), HIToEN: **milk**
Step 2: Since **milk** != **dude**, proceed to step 3.
Step 3: Construct sentence with **EN** as the object: "Jack asked - What's up dude?"
Step 4: Replace **EN** with **HI**: "Jack asked - What's up दूध [d̪ud̪ʱə] (dude)?"
Step 5: Replace the noun phrase with a contextualized phrase: "The cow asked - What's up दूध [d̪ud̪ʱə] (dude)?"

He, He, Nanyun Peng, and Percy Liang. "Pun Generation with Surprise."
Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers). 2019.

Chain of Thought + Algorithmic prompting

- Algorithm based on He et al 2019
- Modifications made to suit the current context

Instruction: Construct a code-mixed Hindi-English pun based on the steps below.

Steps:
Step 1. Create a triplet of:
a. English word **EN**
b. Hindi word **HI**, homophone to **EN**
c. **HI** translated to English, labeled **HIToEN**
Step 2. If **HIToEN** == **EN**, redo step 1. Otherwise, proceed to step 3
Step 3. Construct short sentences (less than 10 words) with **EN** as the object of the sentence.
Step 4. Replace **EN** with **HI**.
Step 5. Replace the noun phrase at the start of the sentence with a contextualized phrase that is closely related to the **HIToEN** word.

Example:
Step 1: EN: **dude**, HI: दूध [d̪ud̪ʱɐ] (milk), HIToEN: **milk**
Step 2: Since **milk** != **dude**, proceed to step 3.
Step 3: Construct sentence with **EN** as the object: "Jack asked - What's up dude?"
Step 4: Replace **EN** with **HI**: "Jack asked - What's up दूध [d̪ud̪ʱɐ] (dude)?"
Step 5: Replace the noun phrase with a contextualized phrase: "The cow asked - What's up दूध [d̪ud̪ʱɐ] (dude)?"

He, He, Nanyun Peng, and Percy Liang. "Pun Generation with Surprise."
Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers). 2019.

Chain of Thought + Algorithmic prompting

- Algorithm based on He et al 2019
- Modifications made to suit the current context
- Zero shot, 1 shot, Few shot
- Negative as well as positive examples

Instruction: Construct a code-mixed Hindi-English pun based on the steps below.

Steps:
Step 1. Create a triplet of:
a. English word **EN**
b. Hindi word **HI**, homophone to **EN**
c. **HI** translated to English, labeled **HIToEN**
Step 2. If **HIToEN** == **EN**, redo step 1. Otherwise, proceed to step 3
Step 3. Construct short sentences (less than 10 words) with **EN** as the object of the sentence.
Step 4. Replace **EN** with **HI**.
Step 5. Replace the noun phrase at the start of the sentence with a contextualized phrase that is closely related to the **HIToEN** word.

Example:
Step 1: EN: **dude**, HI: दूध [d̪ud̪ʱə] (milk), HIToEN: **milk**
Step 2: Since **milk** != **dude**, proceed to step 3.
Step 3: Construct sentence with **EN** as the object: "Jack asked - What's up dude?"
Step 4: Replace **EN** with **HI**: "Jack asked - What's up दूध [d̪ud̪ʱə] (dude)?"
Step 5: Replace the noun phrase with a contextualized phrase: "The cow asked - What's up दूध [d̪ud̪ʱə] (dude)?"

He, He, Nanyun Peng, and Percy Liang. "Pun Generation with Surprise."
Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers). 2019.

Chain of Thought + Algorithmic prompting

- Algorithm based on He et al 2019
- Modifications made to suit the current context
- Zero shot, 1 shot, Few shot
  - Negative as well as positive examples
- Generation temperature varied

Instruction: Construct a code-mixed Hindi-English pun based on the steps below.

Steps:
Step 1. Create a triplet of:
a. English word **EN**
b. Hindi word **HI**, homophone to **EN**
c. **HI** translated to English, labeled **HIToEN**
Step 2. If **HIToEN** == **EN**, redo step 1. Otherwise, proceed to step 3
Step 3. Construct short sentences (less than 10 words) with **EN** as the object of the sentence.
Step 4. Replace **EN** with **HI**.
Step 5. Replace the noun phrase at the start of the sentence with a contextualized phrase that is closely related to the **HIToEN** word.

Example:
Step 1: EN: **dude**, HI: दूध [d̪ud̪ʱɐ] (milk), HIToEN: **milk**
Step 2: Since **milk** != **dude**, proceed to step 3.
Step 3: Construct sentence with **EN** as the object: "Jack asked - What's up dude?"
Step 4: Replace **EN** with **HI**: "Jack asked - What's up दूध [d̪ud̪ʱɐ] (dude)?"
Step 5: Replace the noun phrase with a contextualized phrase: "The cow asked - What's up दूध [d̪ud̪ʱɐ] (dude)?"

He, He, Nanyun Peng, and Percy Liang. "Pun Generation with Surprise."
Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers). 2019.

Chain of Thought + Algorithmic prompting

- Algorithm based on He et al 2019
- Modifications made to suit the current context
- Zero shot, 1 shot, Few shot
  - Negative as well as positive examples
- Generation temperature varied
- 10-20 samples generated

Instruction: Construct a code-mixed Hindi-English pun based on the steps below.

Steps:
Step 1. Create a triplet of:
a. English word **EN**
b. Hindi word **HI**, homophone to **EN**
c. **HI** translated to English, labeled **HIToEN**
Step 2. If **HIToEN** == **EN**, redo step 1. Otherwise, proceed to step 3
Step 3. Construct short sentences (less than 10 words) with **EN** as the object of the sentence.
Step 4. Replace **EN** with **HI**.
Step 5. Replace the noun phrase at the start of the sentence with a contextualized phrase that is closely related to the **HIToEN** word.

Example:
Step 1: EN: **dude**, HI: दूध [d̪ud̪ʱɐ] (milk), HIToEN: **milk**
Step 2: Since **milk** != **dude**, proceed to step 3.
Step 3: Construct sentence with **EN** as the object: "Jack asked - What's up dude?"
Step 4: Replace **EN** with **HI**: "Jack asked - What's up दूध [d̪ud̪ʱɐ] (dude)?"
Step 5: Replace the noun phrase with a contextualized phrase: "The cow asked - What's up दूध [d̪ud̪ʱɐ] (dude)?"

He, He, Nanyun Peng, and Percy Liang. "Pun Generation with Surprise."
Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers). 2019.

Bottleneck I : Identifying Homophones

LLMs struggling to find homophones

Step 1: EN: **pen**, HI: पेन (pen) [pɛn], HIToEN: **pen**
Step 2: Since **pen** == **pen**, go back to step 1.

Step 1 (redo): EN: **apple**, HI: एप्पल (apple) ['æpəl], HIToEN: **apple**
Step 2: Since **apple** == **apple**, go back to step 1.

Step 1 (redo): EN: **ball**, HI: बॉल (ball) [bɔl], HIToEN: **ball**
Step 2: Since **ball** == **ball**, go back to step 1.
...

Stuck in a recursive loop - Excerpt from a recursive output with 70 retries !

Bottleneck II : Token counts in Devanagari

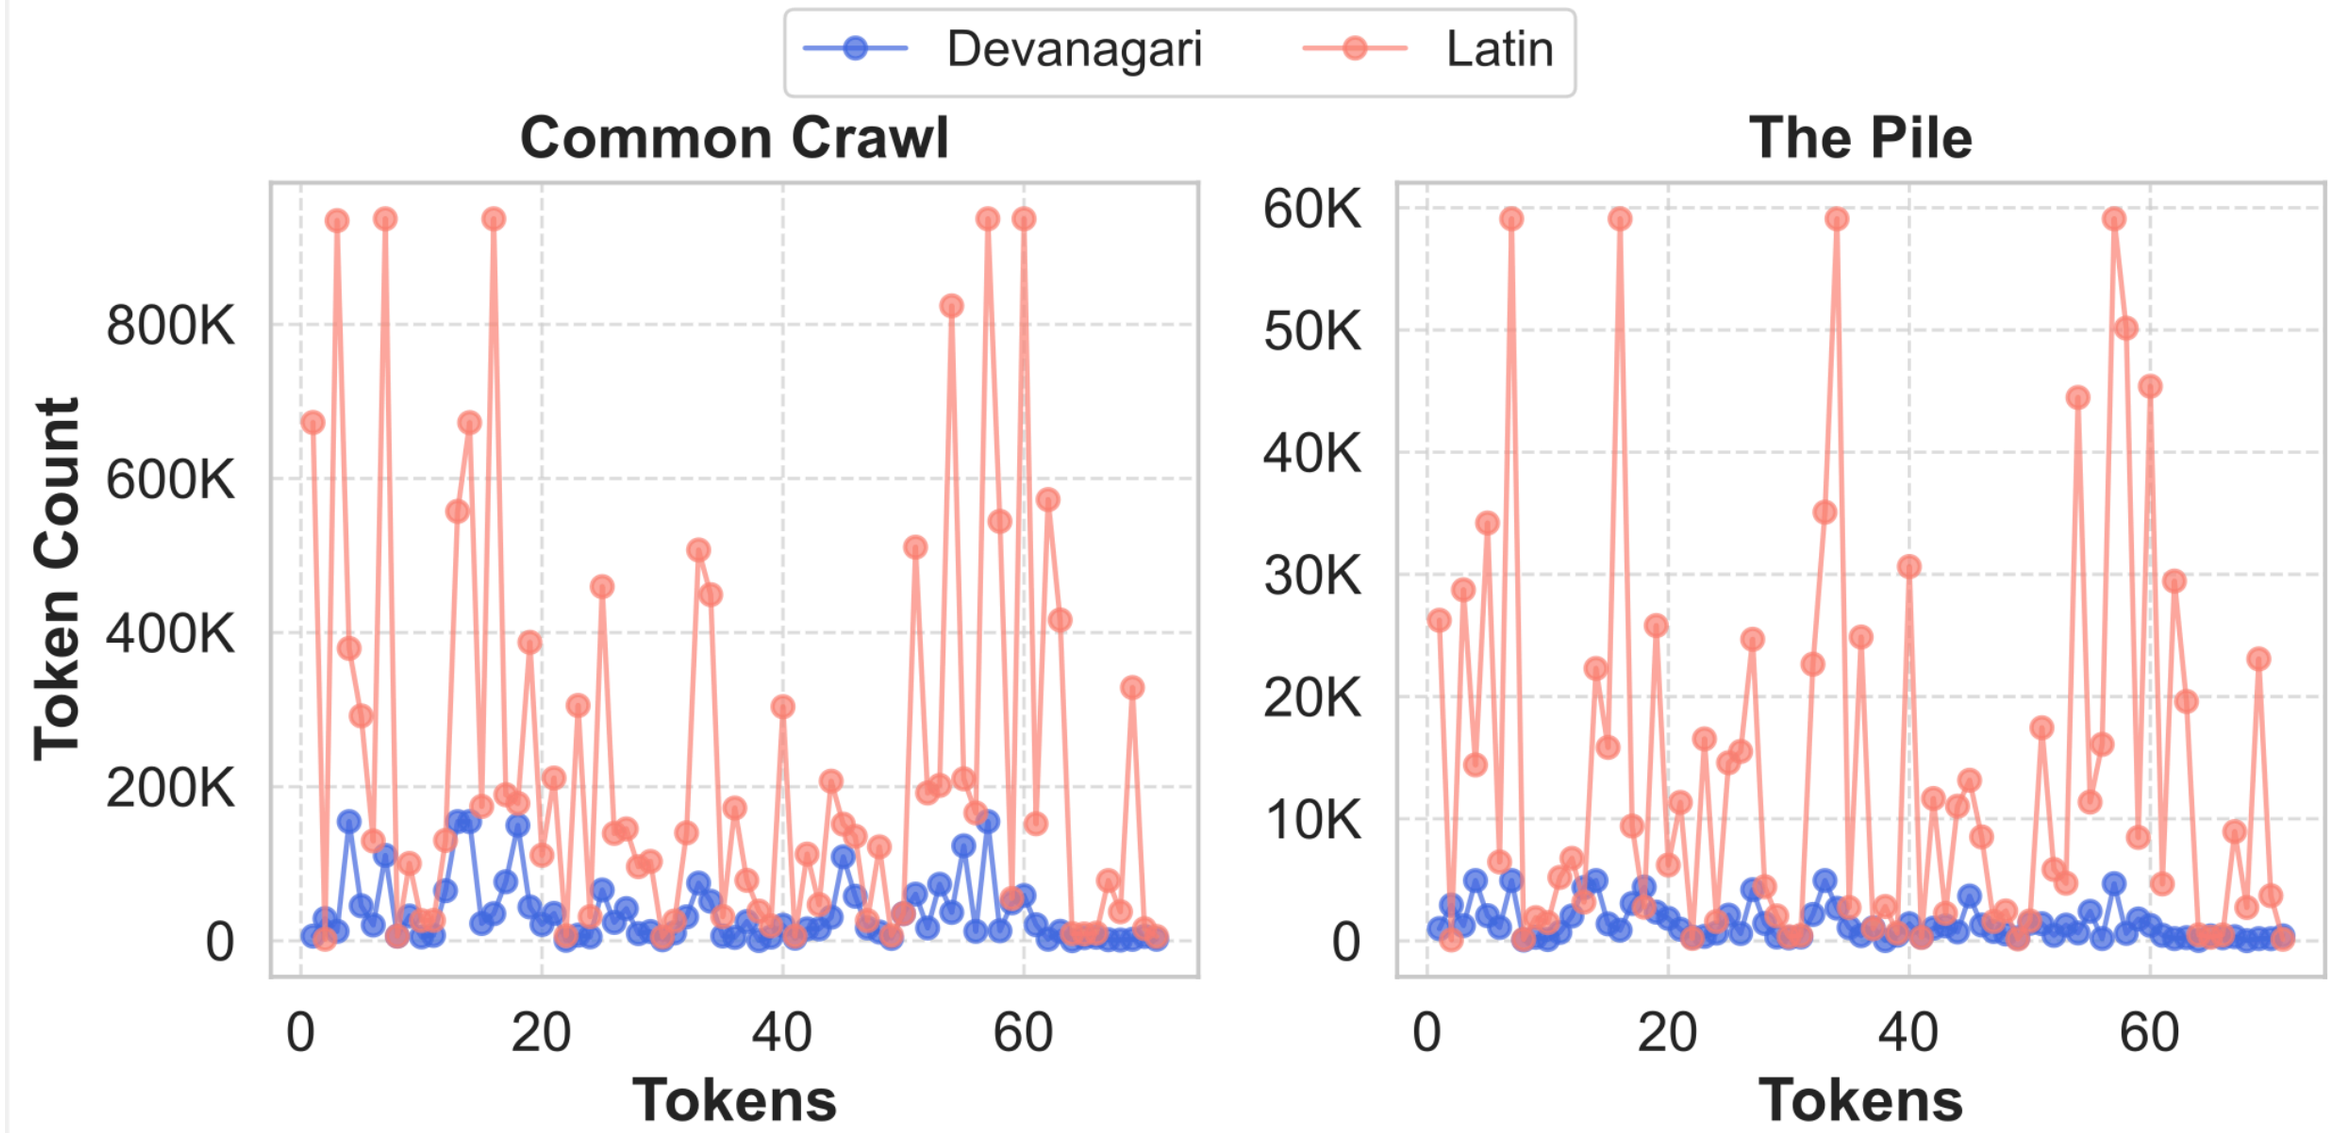

Counts in Devanagari << Counts in Latin

(using WIMBD [1])

[1] Elazar, Yanai, et al. "What's In My Big Data?." *The Twelfth International Conference on Learning Representations*.

Homophone + Transliteration + Prompt

UNIVERSITÄT
DES
SAARLANDES

Homophone + Transliteration + Prompt

- Find Homophones — using Phonetic similarity with Epitran
- Novel, low cost Transliteration Algorithm
- Prompt LLMs with this hybrid approach *

UNIVERSITÄT DES SAARLANDES

Homophone + Transliteration + Prompt

- Find Homophones — using Phonetic similarity with Epitran
- Novel, low cost Transliteration Algorithm
- Prompt LLMs with this hybrid approach *

* Self Refine etc. still used, figure is just for illustration.

UNIVERSITÄT DES SAARLANDES

Homophone + Transliteration + Prompt

- Find Homophones — using Phonetic similarity with Epitran
- Novel, low cost Transliteration Algorithm
- Prompt LLMs with this hybrid approach *

Example Prompt for Hybrid Approach

Task: Generate a code-mixed Hindi-English pun based on the homophones provided as input. Some example input-output pairs are provided as reference.

Input: 'Submit', 'Sab Mit' (everything gets erased)
Output: "Exam ki answer sheet return karte hi SUBMIT jata hai"

Input: <EnWord>, <HiTransliteratedHomophone><(EnglishTranslation)>

* Self Refine etc. still used, figure is just for illustration.

UNIVERSITÄT DES SAARLANDES

Evaluation : Survey

Evaluation : Survey

- 56 participants

Evaluation : Survey

- 56 participants
- 4-5 random samples per category

Evaluation : Survey

- 56 participants
- 4-5 random samples per category
- Human jokes + Distractors added for comparison

Evaluation : Survey

- 56 participants
- 4-5 random samples per category
- Human jokes + Distractors added for comparison
- Insincere submissions discarded

Evaluation : Survey

- 56 participants
- 4-5 random samples per category
- Human jokes + Distractors added for comparison
- Insincere submissions discarded
    - Mean rating ~ 1 for all

Evaluation : Survey

- 56 participants
- 4-5 random samples per category
- Human jokes + Distractors added for comparison
- Insincere submissions discarded
  - Mean rating ~ 1 for all
  - High rating ~ 5 for all

Evaluation : Survey

- 56 participants
- 4-5 random samples per category
- Human jokes + Distractors added for comparison
- Insincere submissions discarded
  - Mean rating ~ 1 for all
  - High rating ~ 5 for all
- 39 Submissions remained

Evaluation : Survey

- 56 participants
- 4-5 random samples per category
- Human jokes + Distractors added for comparison
- Insincere submissions discarded
  - Mean rating ~ 1 for all
  - High rating ~ 5 for all
- 39 Submissions remained

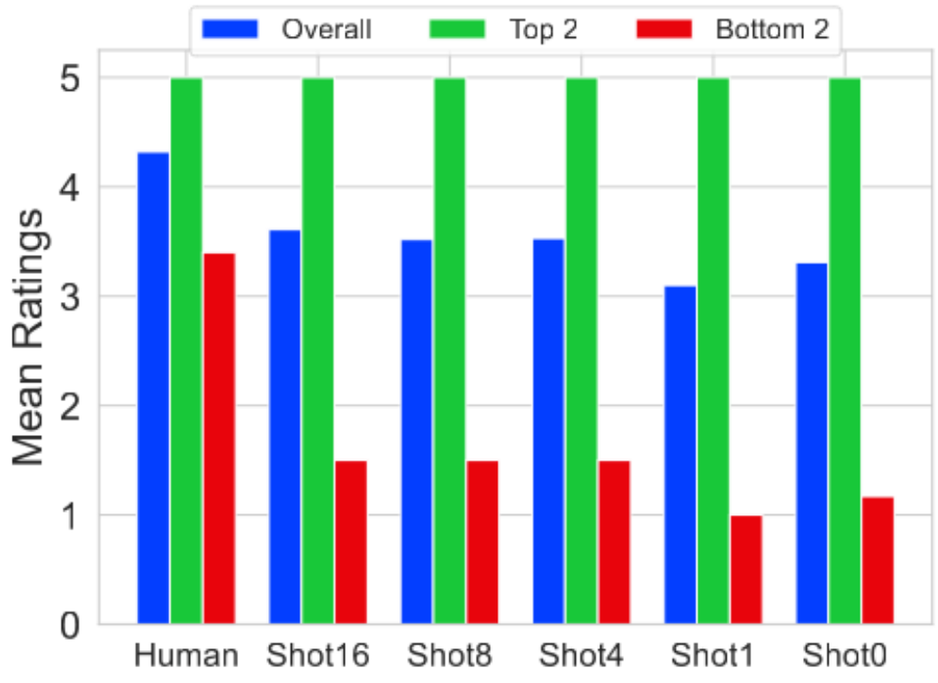

Evaluation : Survey

- 56 participants
- 4-5 random samples per category
- Human jokes + Distractors added for comparison
- Insincere submissions discarded
  - Mean rating ~ 1 for all
  - High rating ~ 5 for all
- 39 Submissions remained

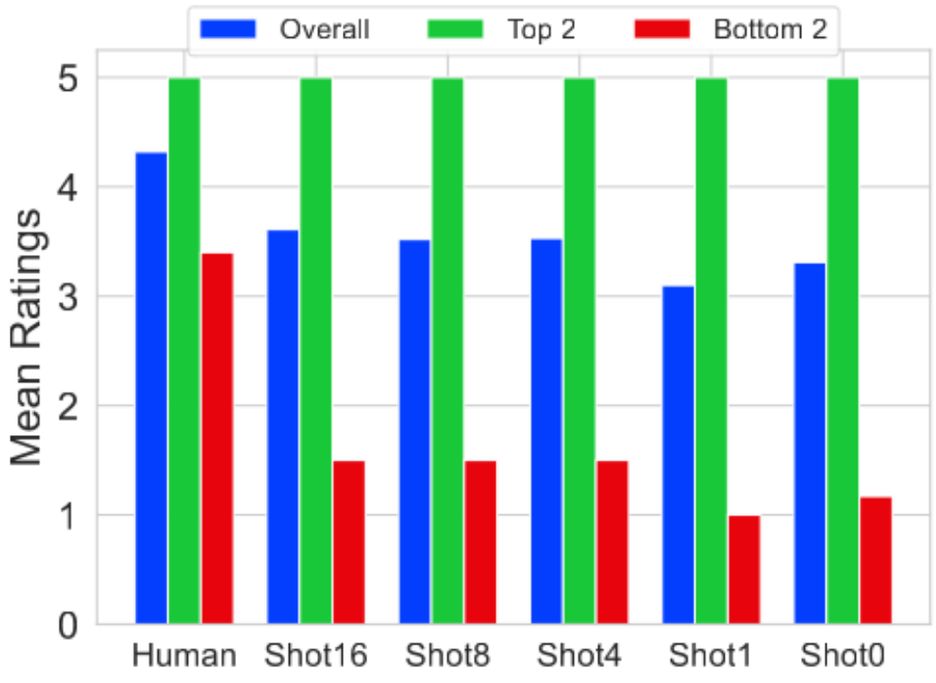

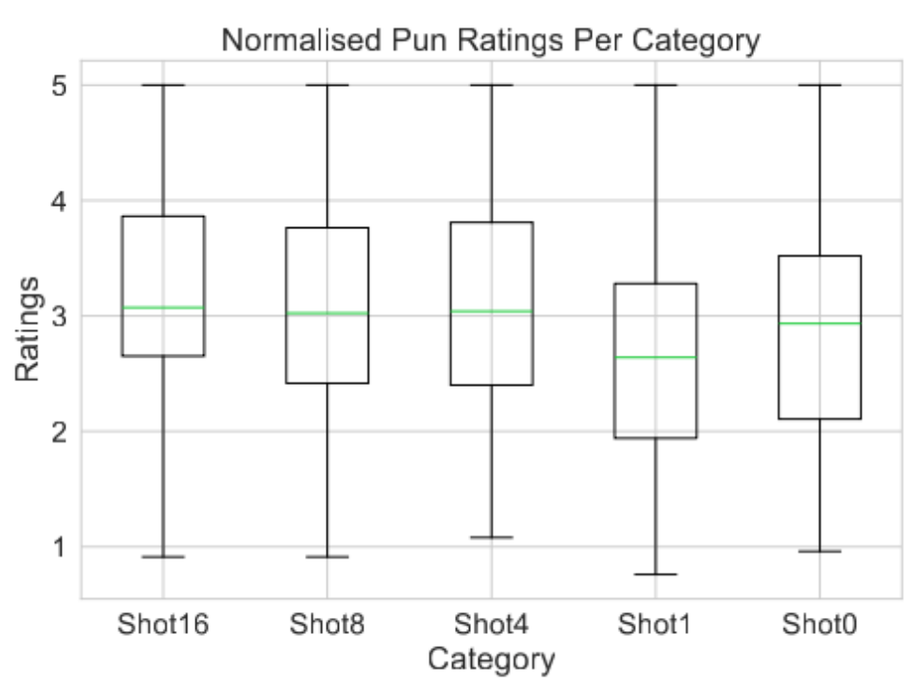

Evaluation : Survey

- 56 participants
- 4-5 random samples per category
- Human jokes + Distractors added for comparison
- Insincere submissions discarded
  - Mean rating ~ 1 for all
  - High rating ~ 5 for all
- 39 Submissions remained

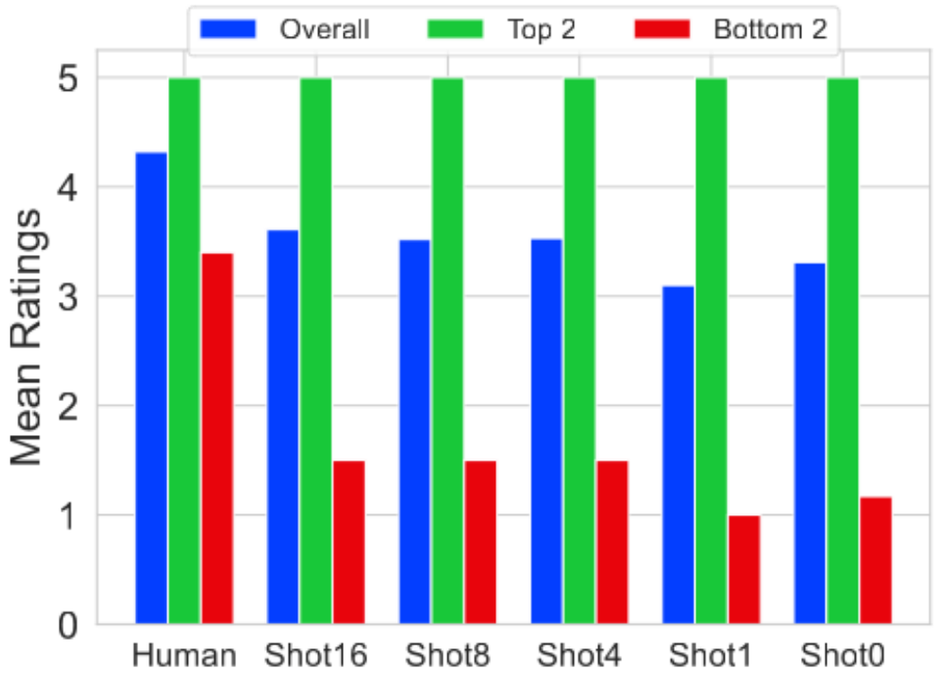

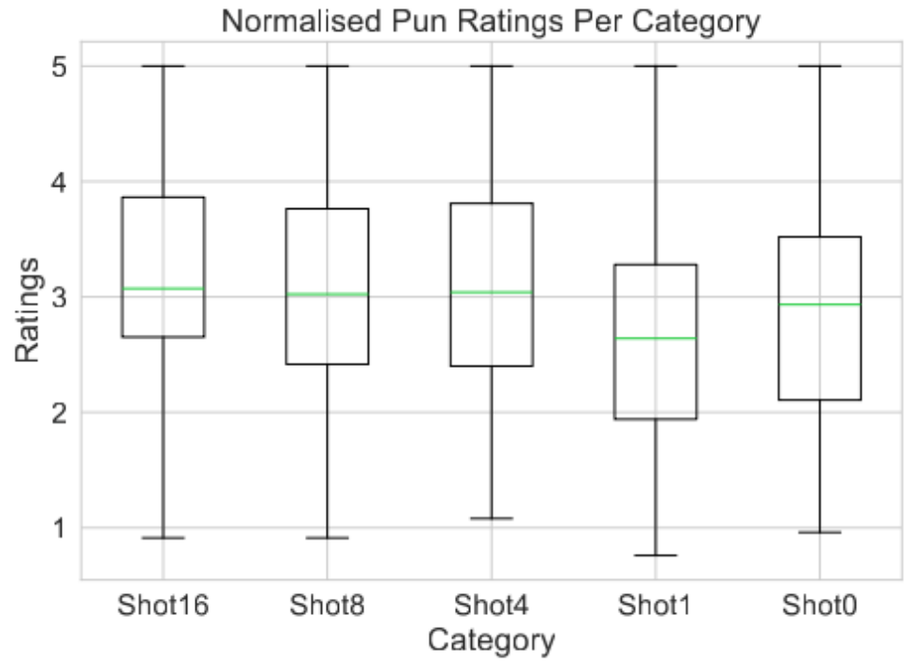

Conclusion : LLMs can generate engaging Puns but lack consistency

Limitations

Limitations

- Current Homophone generation — restricted to single words
  - I will see vs Aalsi (lazy) difficult to generate

UNIVERSITÄT DES SAARLANDES

Limitations

- Current Homophone generation — restricted to single words
  - I will see vs Aalsi (lazy) difficult to generate
- Results only for GPT 3.5

UNIVERSITÄT DES SAARLANDES

Limitations

- Current Homophone generation — restricted to single words
  - I will see vs Aalsi (lazy) difficult to generate
- Results only for GPT 3.5
- Limited Survey (voluntary participation)

UNIVERSITÄT DES SAARLANDES

Limitations

- Current Homophone generation — restricted to single words
  - I will see vs Aalsi (lazy) difficult to generate
- Results only for GPT 3.5
- Limited Survey (voluntary participation)
- Hinglish datasets have insufficient scale + inadequate maintenance

UNIVERSITÄT DES SAARLANDES

Limitations

- Current Homophone generation — restricted to single words
  - I will see vs Aalsi (lazy) difficult to generate
- Results only for GPT 3.5
- Limited Survey (voluntary participation)
- Hinglish datasets have insufficient scale + inadequate maintenance

Word	IPA	Present	Actual
इनके	[ɪnkɛ]	unke	inke
में	[me:n]	be	mein
है	[hɛ]	ahai	hai

UNIVERSITÄT DES SAARLANDES

Limitations

- Current Homophone generation — restricted to single words
  - I will see vs Aalsi (lazy) difficult to generate
- Results only for GPT 3.5
- Limited Survey (voluntary participation)
- Hinglish datasets have insufficient scale + inadequate maintenance

Word	IPA	Present	Actual
इनके	[ɪnkɛ]	unke	inke
में	[me:n]	be	mein
है	[hɛ]	ahai	hai

Annotation Mistake in Dakshina dataset by Google

UNIVERSITÄT DES SAARLANDES

Conclusion

Conclusion

- English based humour generation techniques can be adapted to generate puns in Code-Mixed settings

UNIVERSITÄT
DES
SAARLANDES

Conclusion

- English based humour generation techniques can be adapted to generate puns in Code-Mixed settings
- Using Chain of Thought + Self Refine can uncover hidden challenges in Humour generation

Conclusion

- English based humour generation techniques can be adapted to generate puns in Code-Mixed settings
- Using Chain of Thought + Self Refine can uncover hidden challenges in Humour generation
- For Hinglish : Homophone Identification + Transliteration help

UNIVERSITÄT DES SAARLANDES

Conclusion

- English based humour generation techniques can be adapted to generate puns in Code-Mixed settings
- Using Chain of Thought + Self Refine can uncover hidden challenges in Humour generation
- For Hinglish : Homophone Identification + Transliteration help
- With Hybrid Prompting :: High Quality puns can be generated, but not with consistency

UNIVERSITÄT DES SAARLANDES

Conclusion

- English based humour generation techniques can be adapted to generate puns in Code-Mixed settings
- Using Chain of Thought + Self Refine can uncover hidden challenges in Humour generation
- For Hinglish : Homophone Identification + Transliteration help
- With Hybrid Prompting :: High Quality puns can be generated, but not with consistency
- Pun Generation in Code Mixed settings — Has many exciting future research directions !

UNIVERSITÄT DES SAARLANDES