


Towards Useful AI Interpretability for Humans via Interactive AI Explanations

Hua Shen

 huashen@umich.edu  @huashen218

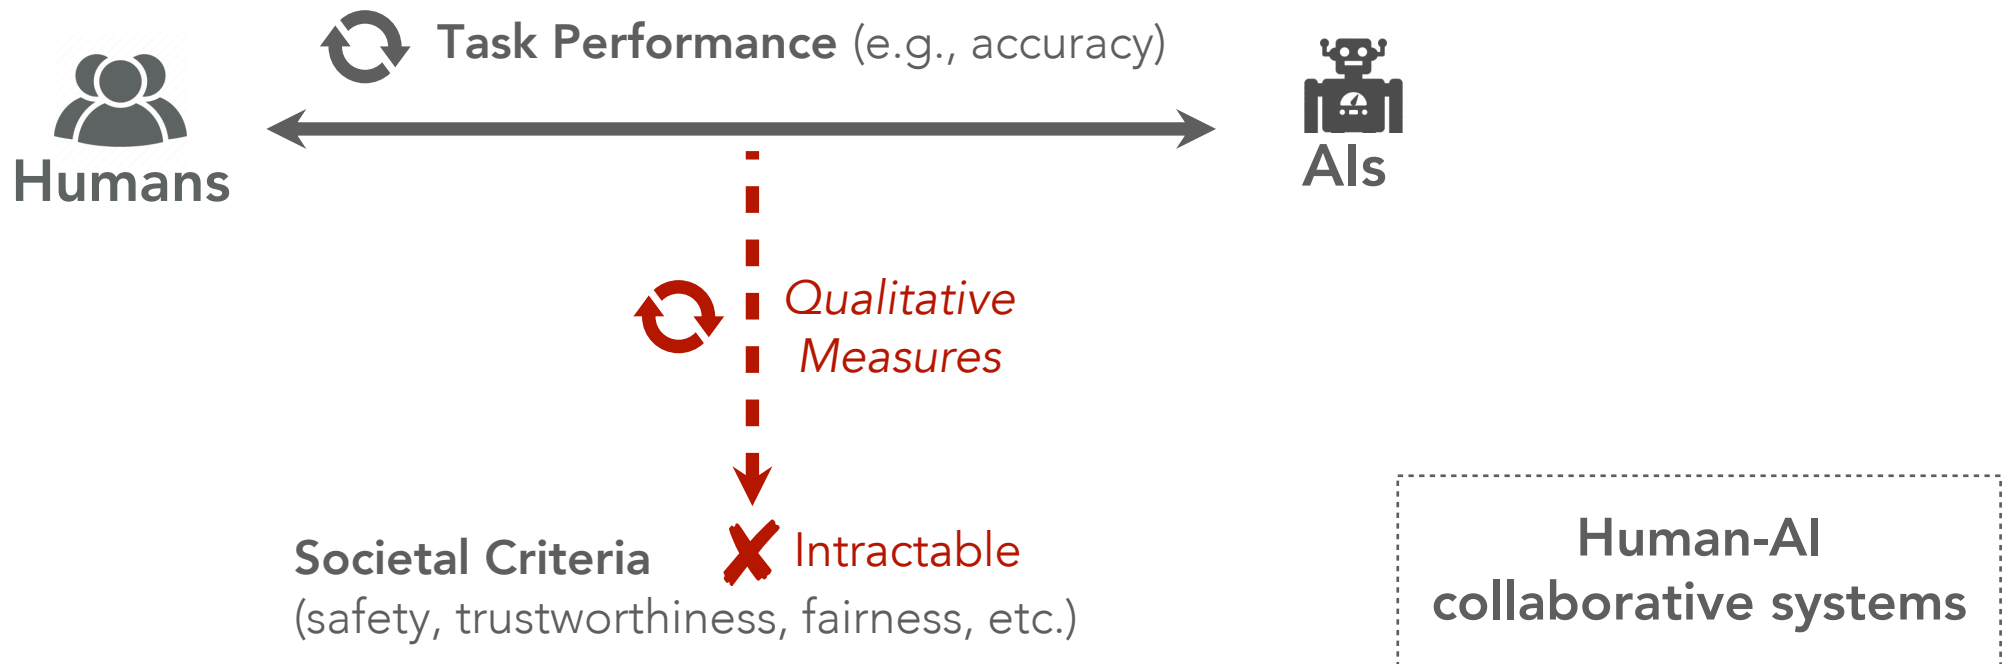
University of Michigan

Background &
Motivation



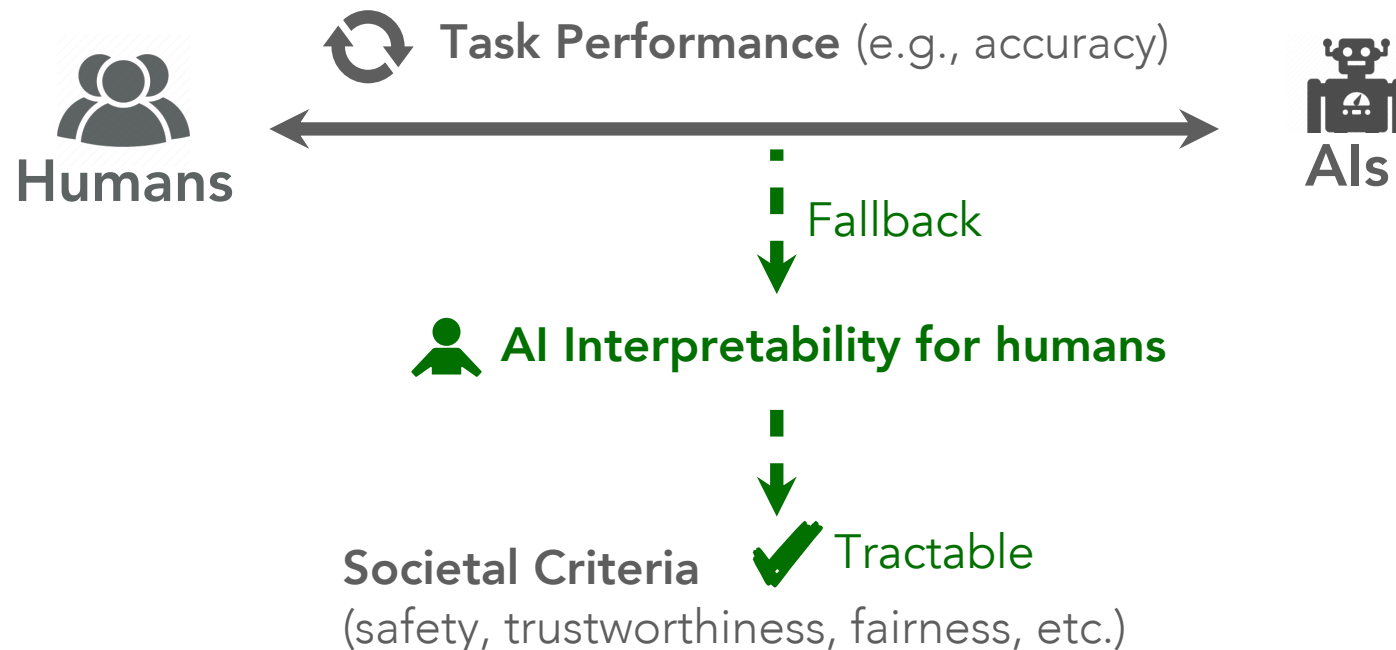
Why do we need AI interpretability?

Human-AI collaborative systems are not only **optimized** for **task performance** (e.g., accuracy), but also are required to **satisfy** vital **societal criteria** (e.g., trustworthiness, safety, fairness, etc.).

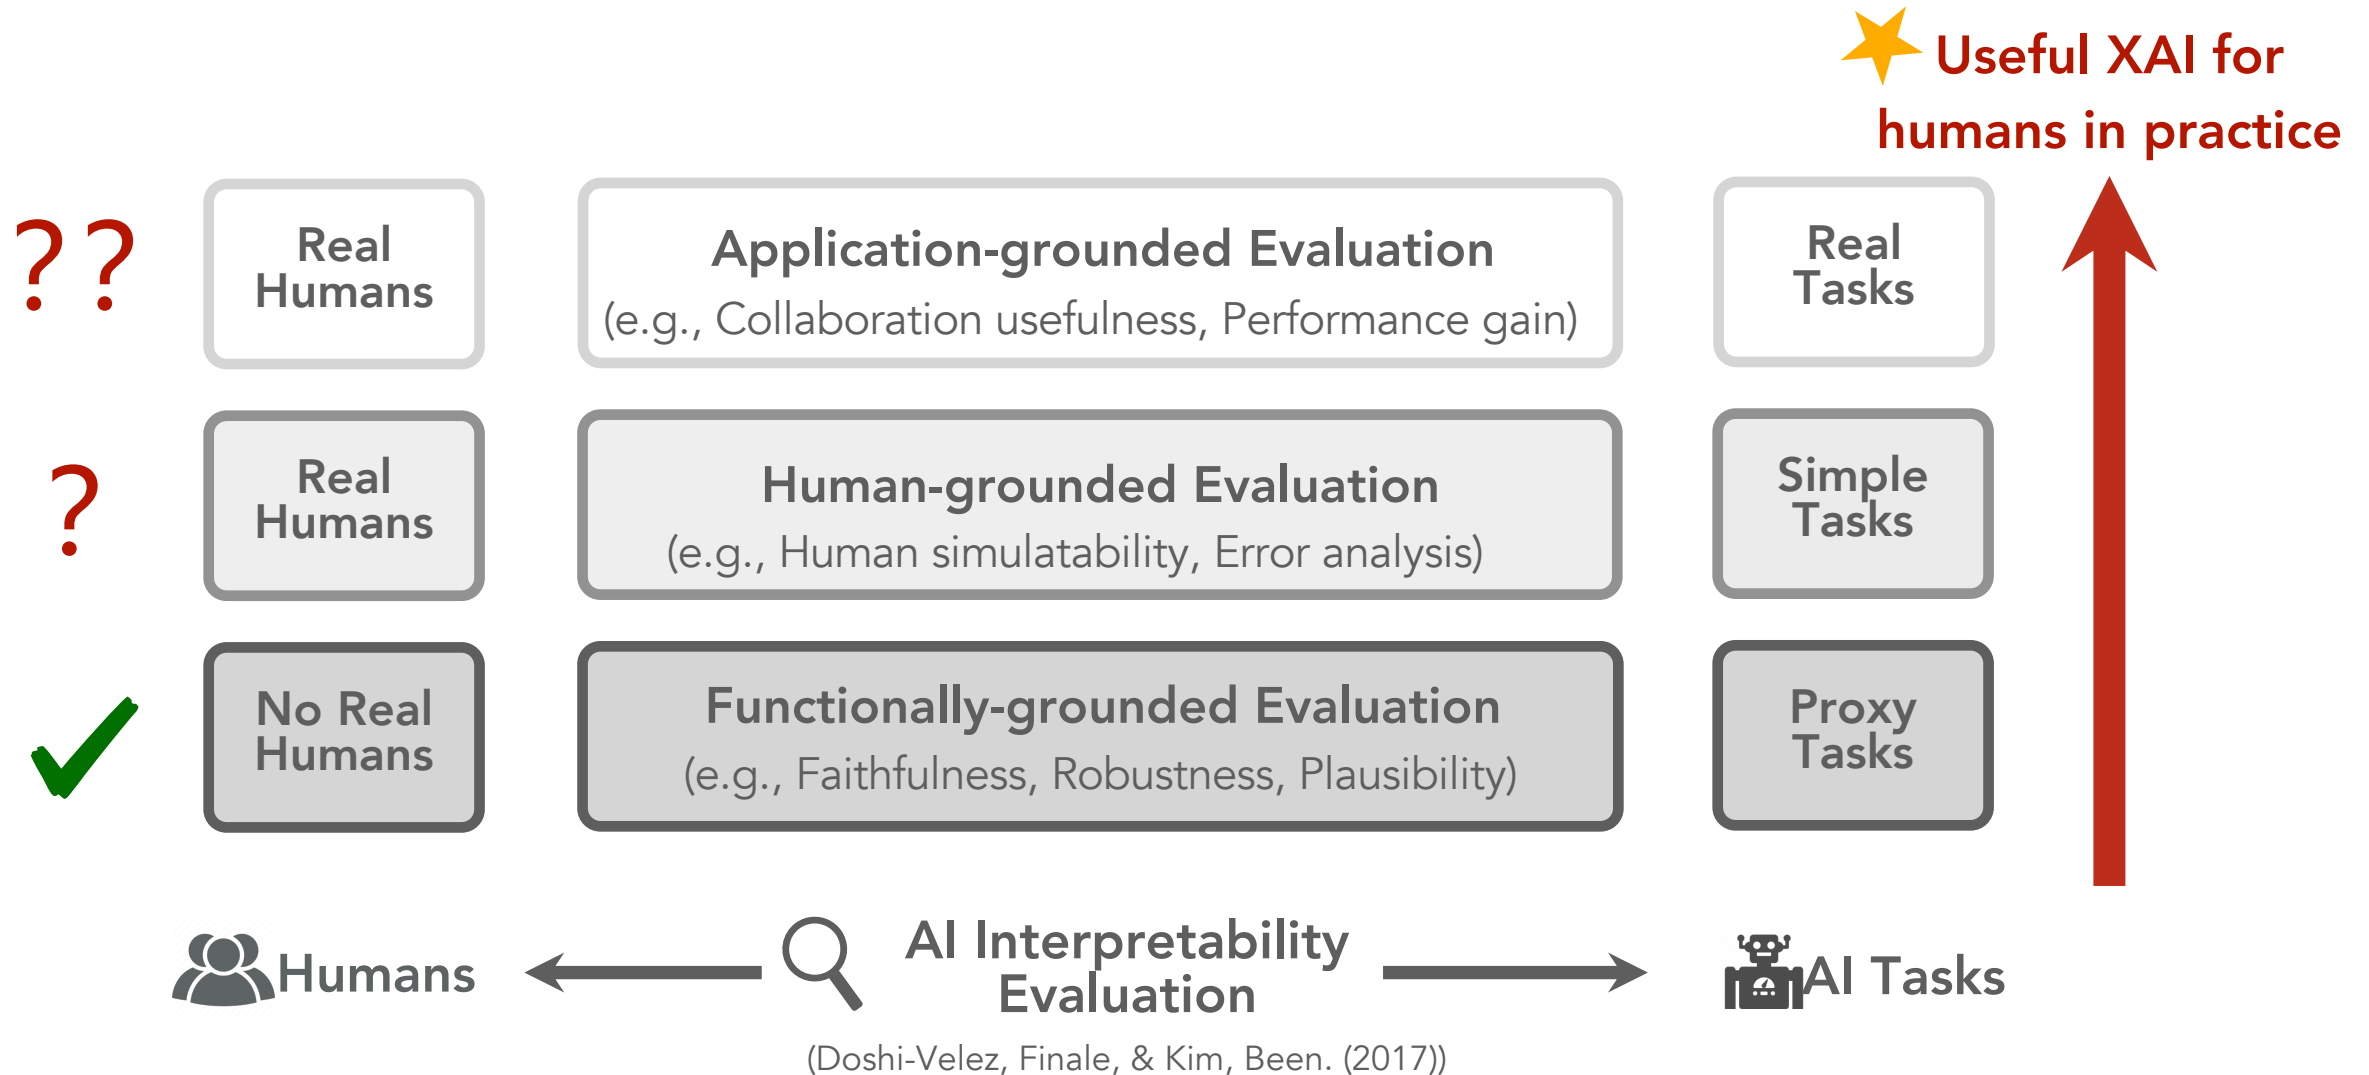


The usefulness of XAI for humans is crucial

“AI interpretability is a **fallback** to be **used by humans** to **gauge the AI model reasoning** and **assess the societal measurements**”



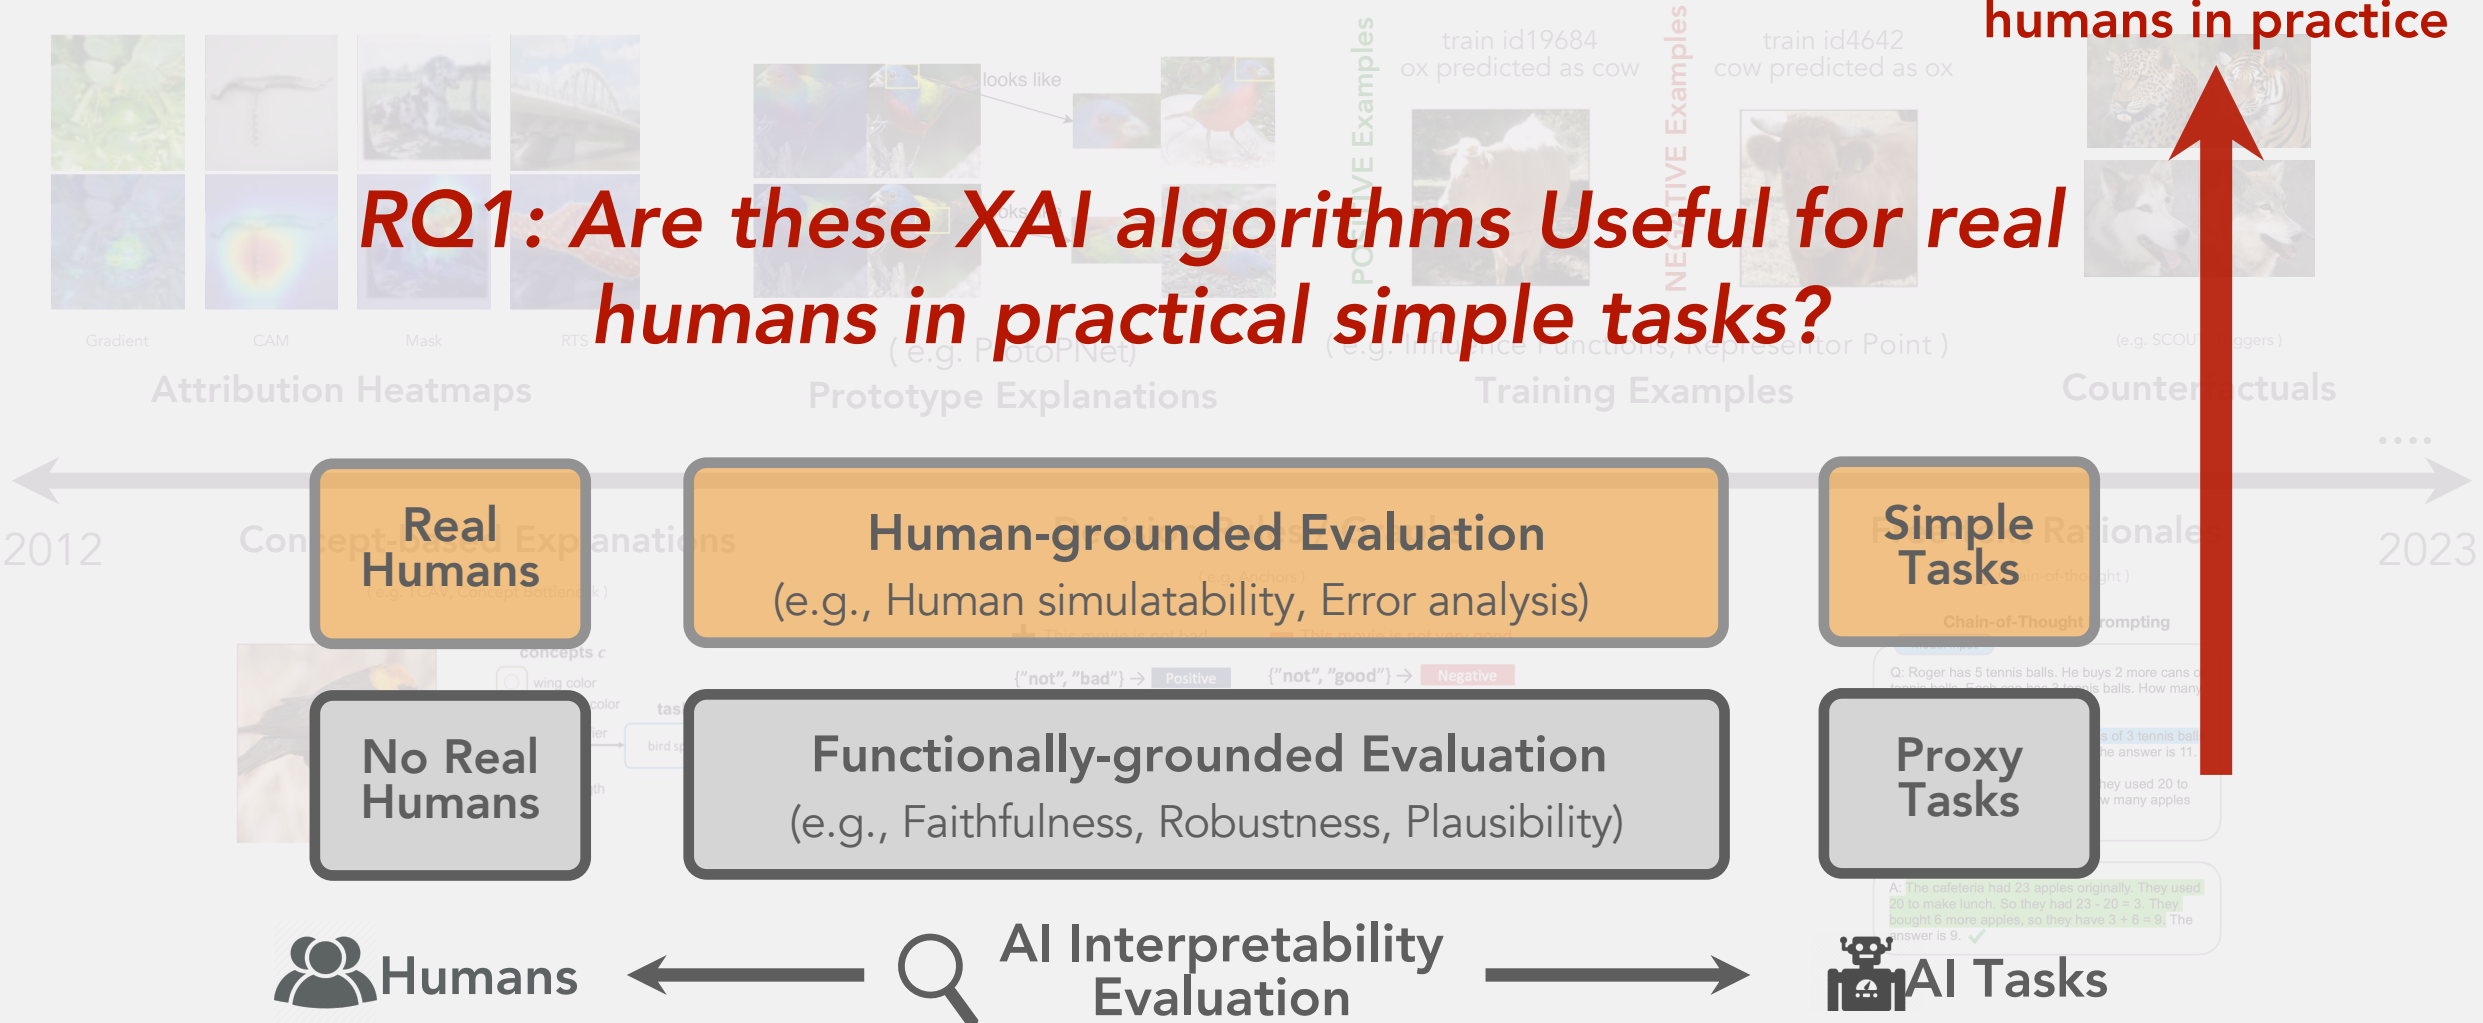
Evaluation of XAI usefulness




Under-Explored: human evaluation of XAI usefulness

★ Useful XAI for humans in practice

RQ1: Are these XAI algorithms Useful for real humans in practical simple tasks?



Background & Motivation



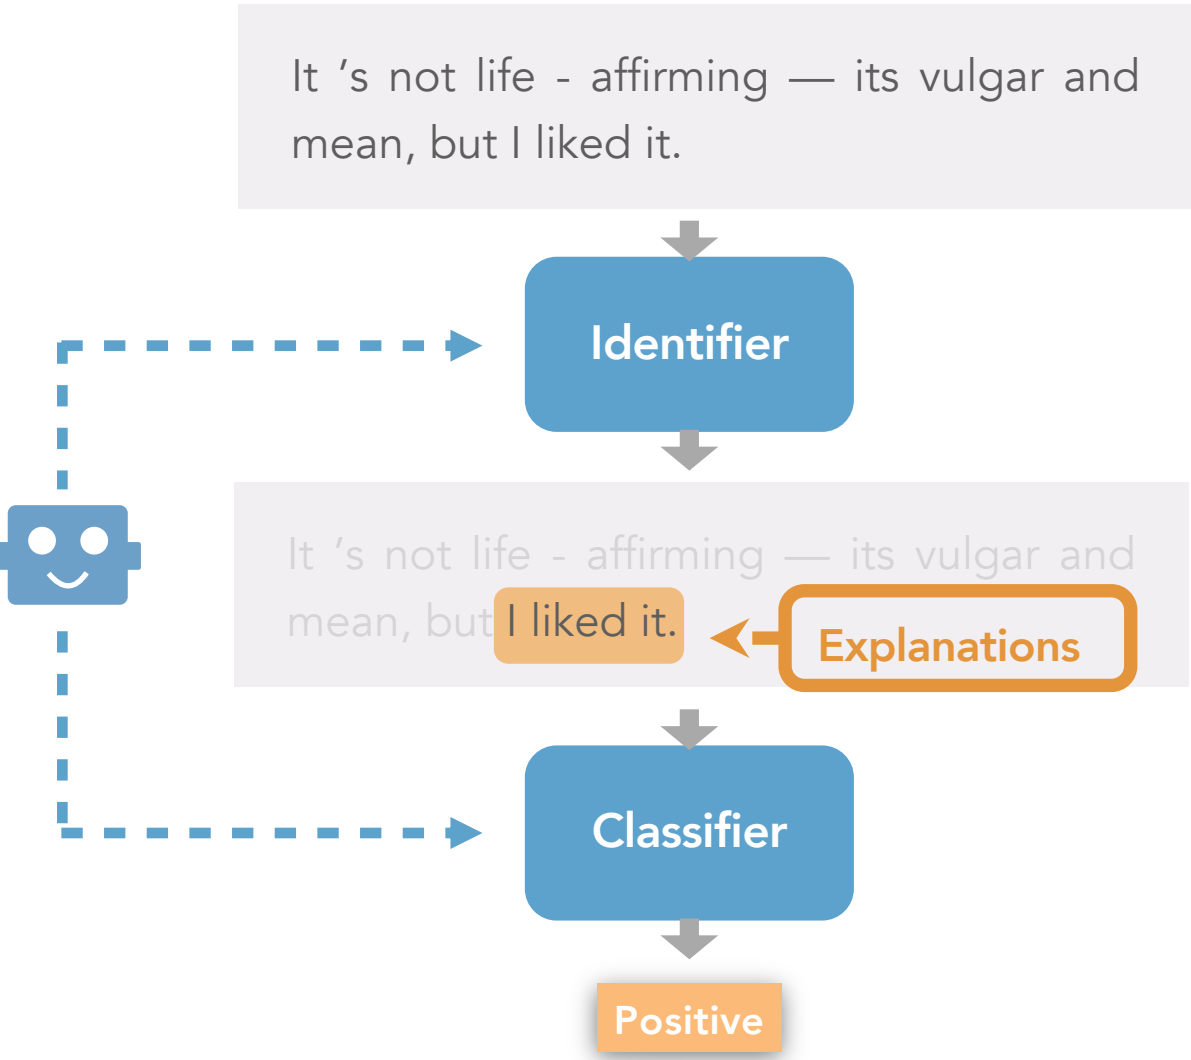
RQ1: Are XAI Useful for Humans?



NLP Interpretability  ACL 2022
22ND - 27TH MAY | 60TH MEETING | DUBLIN

Vision Interpretability

Self-Explaining Language Models



Explanations:
A sufficient **subset** of input **words**, that are **short** and **coherent**, yet **sufficient** to make the **correct** model's **prediction**.



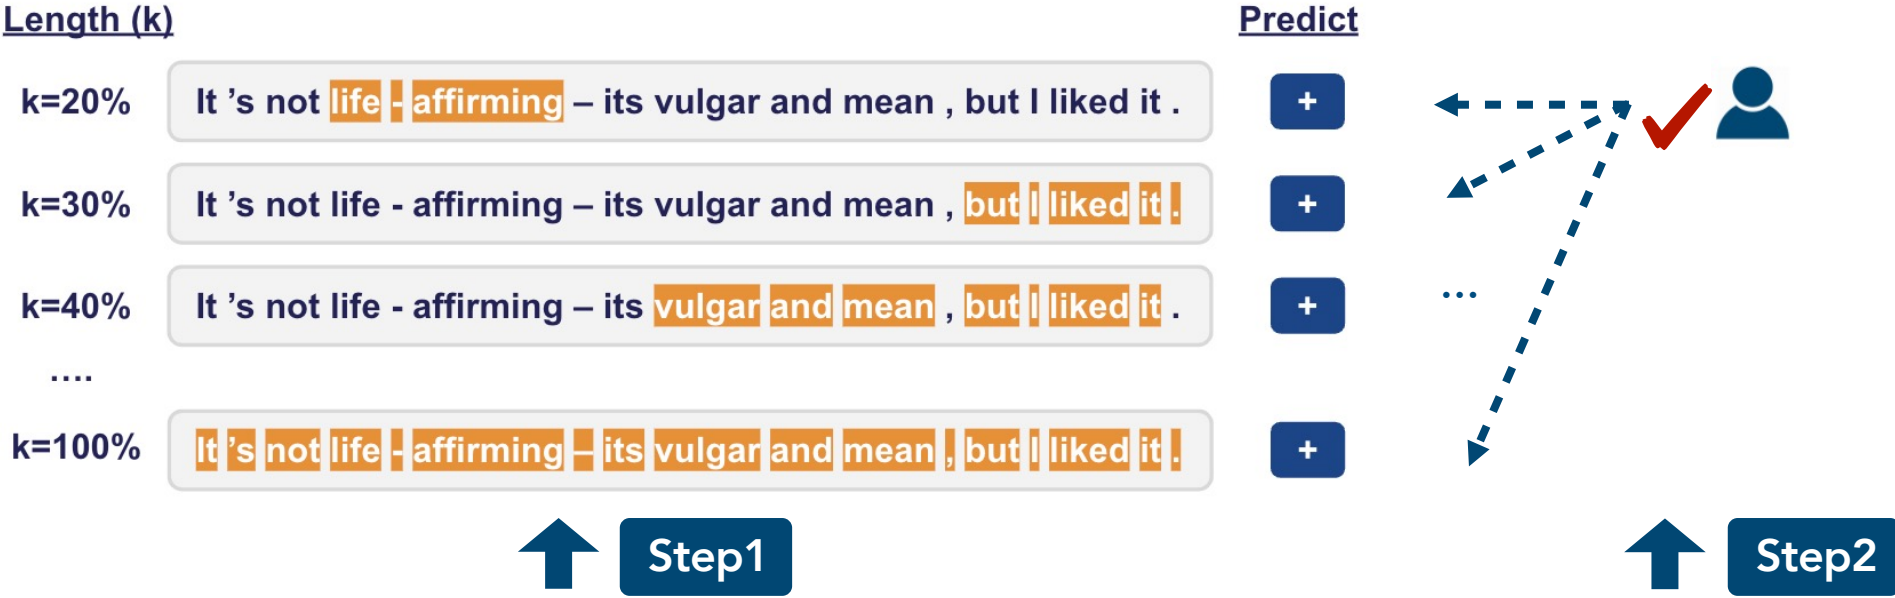
AI Researchers' Assumption

Shorter Explanations are Better for End Users.

? Yet to be validated by human studies!

Lei, Tao, Regina Barzilay, and Tommi Jaakkola. "Rationalizing neural predictions." EMNLP, 2016.
Vafa, Keyon, et al. "Rationales for sequential predictions." EMNLP, 2021.
Bastings, Jasmijn, et al. "Interpretable neural predictions with differentiable binary variables." ACL, 2019.

Are *Shortest* AI Explanations the *Most Useful* for Human Understanding?



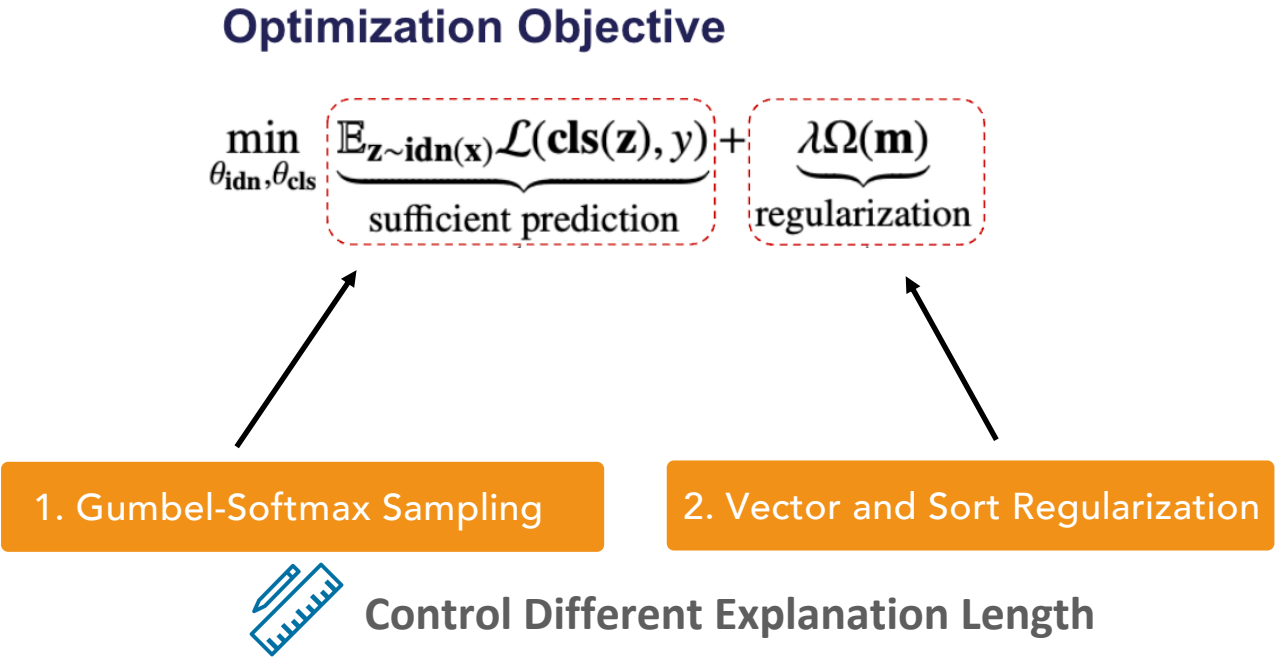
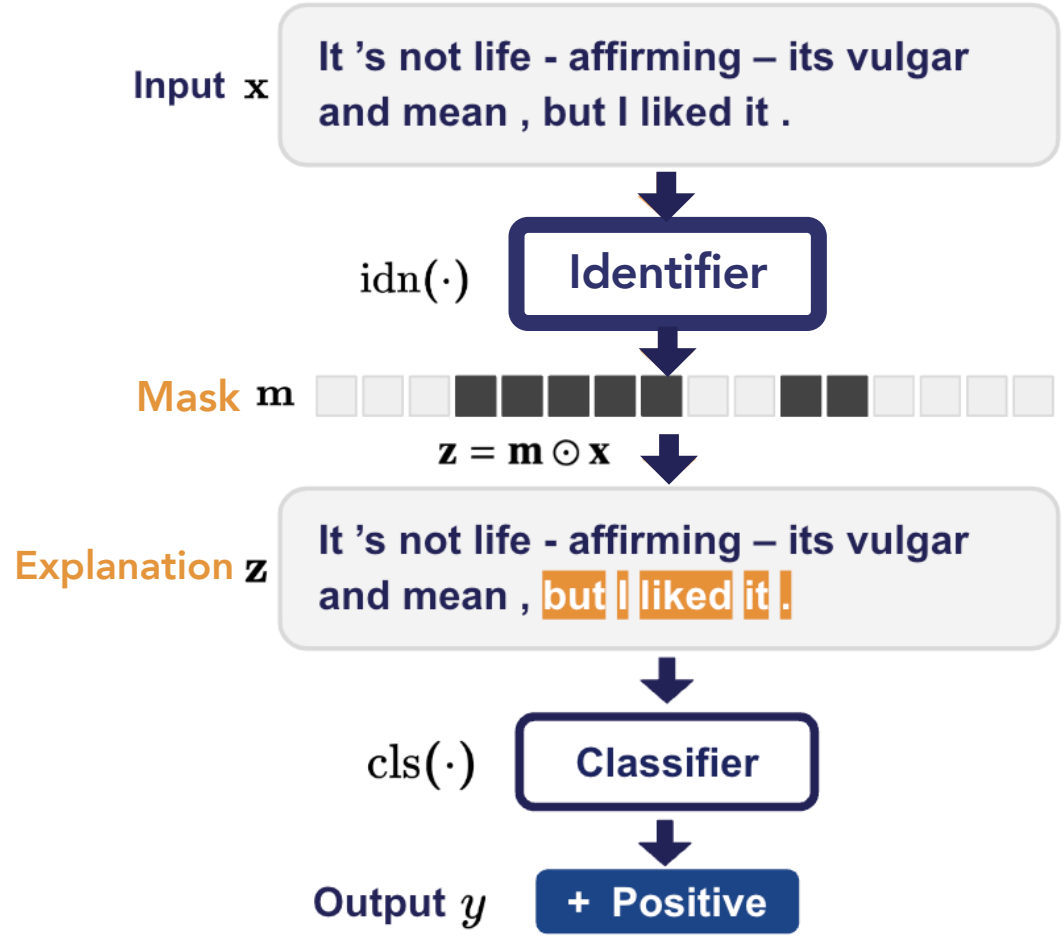
Propose a **novel self-explaining LM** to generate explanations with different lengths

Humans guess the labels with explanations of different lengths

Contribution

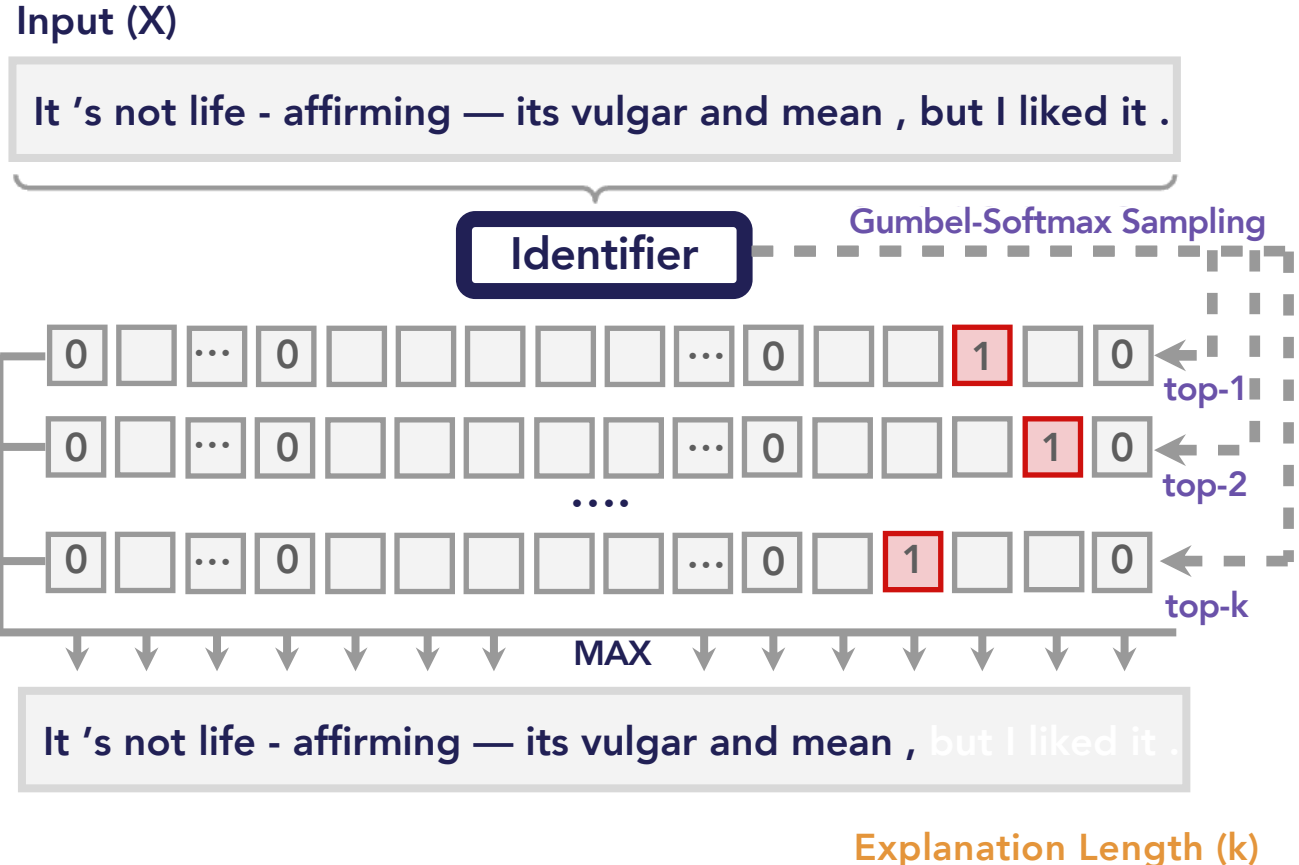
A novel self-explaining model **human interactively guess and select the LM output**

LimitedInk: A novel self-explaining LM

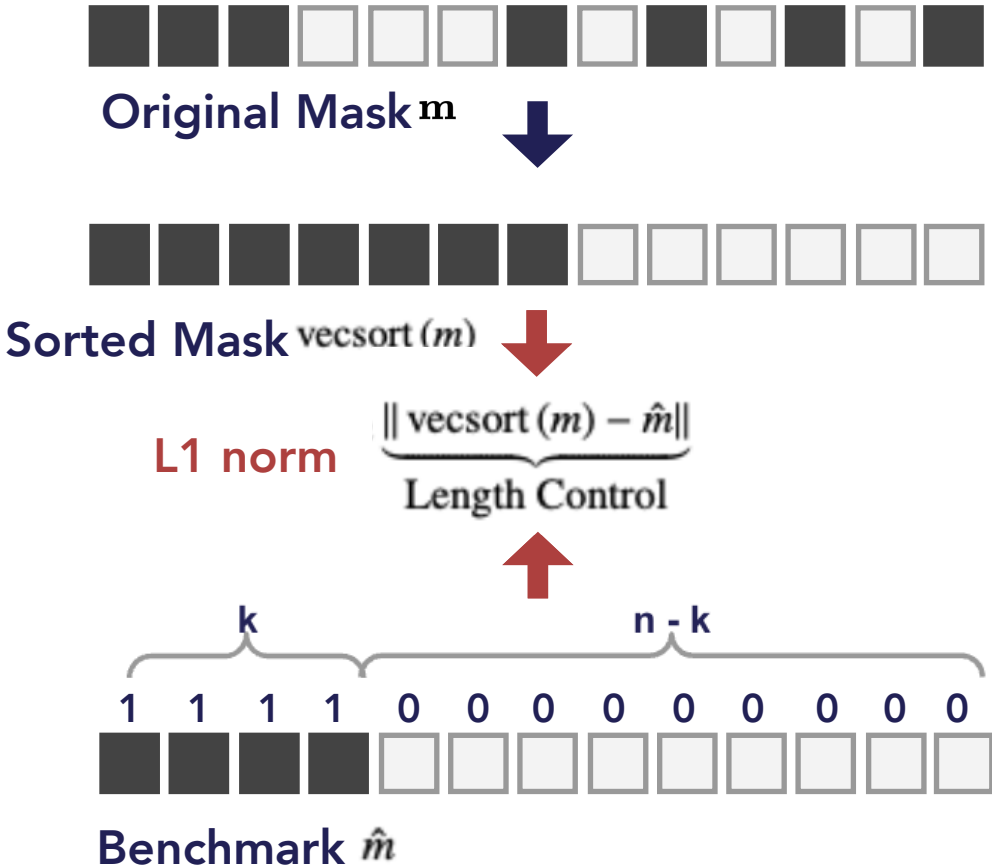


How to control explanation length in LimitedInk

1. Gumbel-Softmax Sampling



2. Vector and Sort Regularization



Can LimitedInk **perform well on classification?**

- End-task classification: **Task**, weighted average F1
- Human Plausibility with annotated dataset: **P**recision, **R**ecall, Token-level **F1**

Method	Movies				BoolQ				Evidence Inference				MultiRC				FEVER			
	Task	P	R	F1	Task	P	R	F1	Task	P	R	F1	Task	P	R	F1	Task	P	R	F1
Full-Text	.91	-	-	-	.47	-	-	-	.48	-	-	-	.67	-	-	-	.89	-	-	-
Sparse-N	.79	.18	.36	.24	.43	.12	.10	.11	.39	.02	.14	.03	.60	.14	.35	.20	.83	.35	.49	.41
Sparse-C	.82	.17	.36	.23	.44	.15	.11	.13	.41	.03	.15	.05	.62	.15	.41	.22	.83	.35	.52	.42
Sparse-IB	.84	.21	.42	.28	.46	.17	.15	.15	.43	.04	.21	.07	.62	.20	.33	.25	.85	.37	.50	.43
LIMITEDINK Length Level	.90	.26	.50	.34	.56	.13	.17	.15	.50	.04	.27	.07	.67	.22	.40	.28	.90	.28	.67	.39
		50%				30%				50%				50%				40%		

LimitedInk **performed compatible with three SOTA baselines** on the two common rationale metrics in five ERASER text classification benchmark datasets.

Step2 - Human Study Setups

LimitedInk Explanations

k=10%
... not

k=20%
..... life - affirming

k=30%
..... but || liked it

k=40%
..vulgar and mean but || liked it

k=50%
..... its vulgar and mean , but || liked it , ...

Random text spans (similar length)

k=10%
... affirming

k=20%
..... affirming - its ...

k=30%
.... - its vulgar and mean

k=40%
..... not life - affirming ... || liked it

k=50%
... life - affirming - its vulgar and mean , but

Only highlight explanations & hide other texts!

Five-level explanations:
10%, 20%, 30%, 40%, 50%

We conducted **user studies** to investigate the **human understanding** on **LimitedInk** and **Baseline** (random sampled tokens).

User Interface for Human Interaction

Select Sentiment and Confidence of the Displayed Parts of Movie Review

Please select the **sentiment label** of the displayed parts of the movie review and provide your **confidence** on the selection.

Parts of the Movie Review 1

..... recall hearing species 2 described as "erotic." I would love to know who used with that adjective for this a woman's abdomen as an alien baby claws its way free , splat blood and gore in all directions . anyone turned on by that

Question1: Is the movie review **Positive** or **Negative**? Please guess based on the parts of texts you see.

(Empty reviews are usually caused by data processing errors)

Question2: How **Confident** are you in your above selection?

- The displayed texts show clear attitude, and reflects the core sentiment (like/dislike) of the full review.

- The displayed texts show attitude towards the movie, but not very clear to reflect the core sentiment.

- The displayed texts seem positive/negative, but I cannot guess if it's representative of the full review.

- The displayed texts are ambiguous. I am not confident on the attitude towards the movie.

- The displayed texts are too trivial and does not reflect on the larger themes.



Sentiment Analysis:

we randomly sampled **100** reviews (correct prediction) from the **Movie review** test set



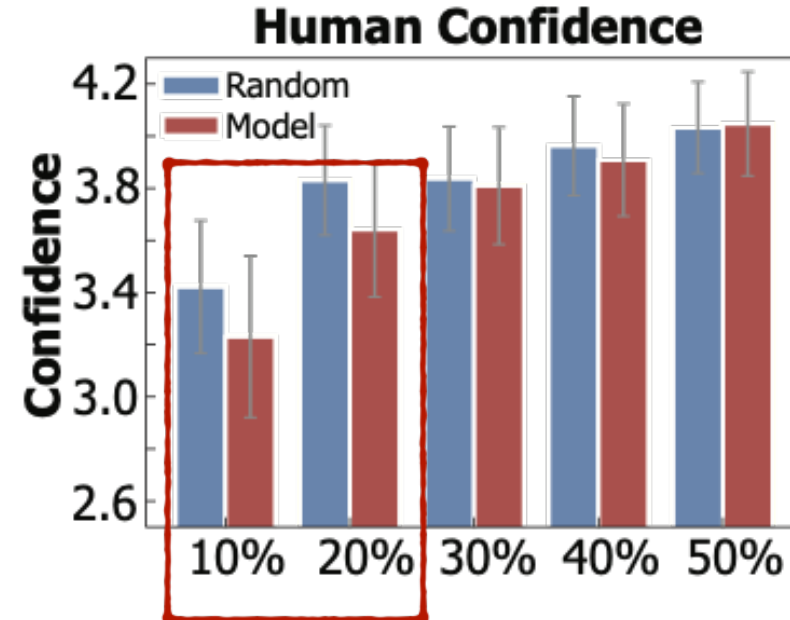
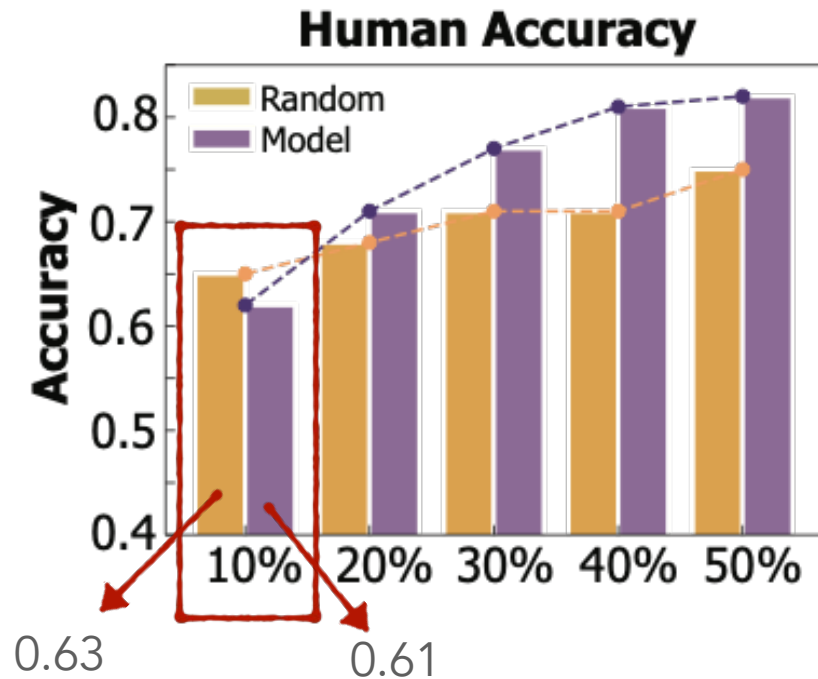
1. simulate model predictions



2. provide the confidence




Key Findings



Human **accuracy** and **confidence**, at the shortest level (i.e., 10% length), are **lower than** the random baseline.

The **shortest AI explanations** are **NOT always Useful** for humans to understand the AI's decision-making.

Background & Motivation

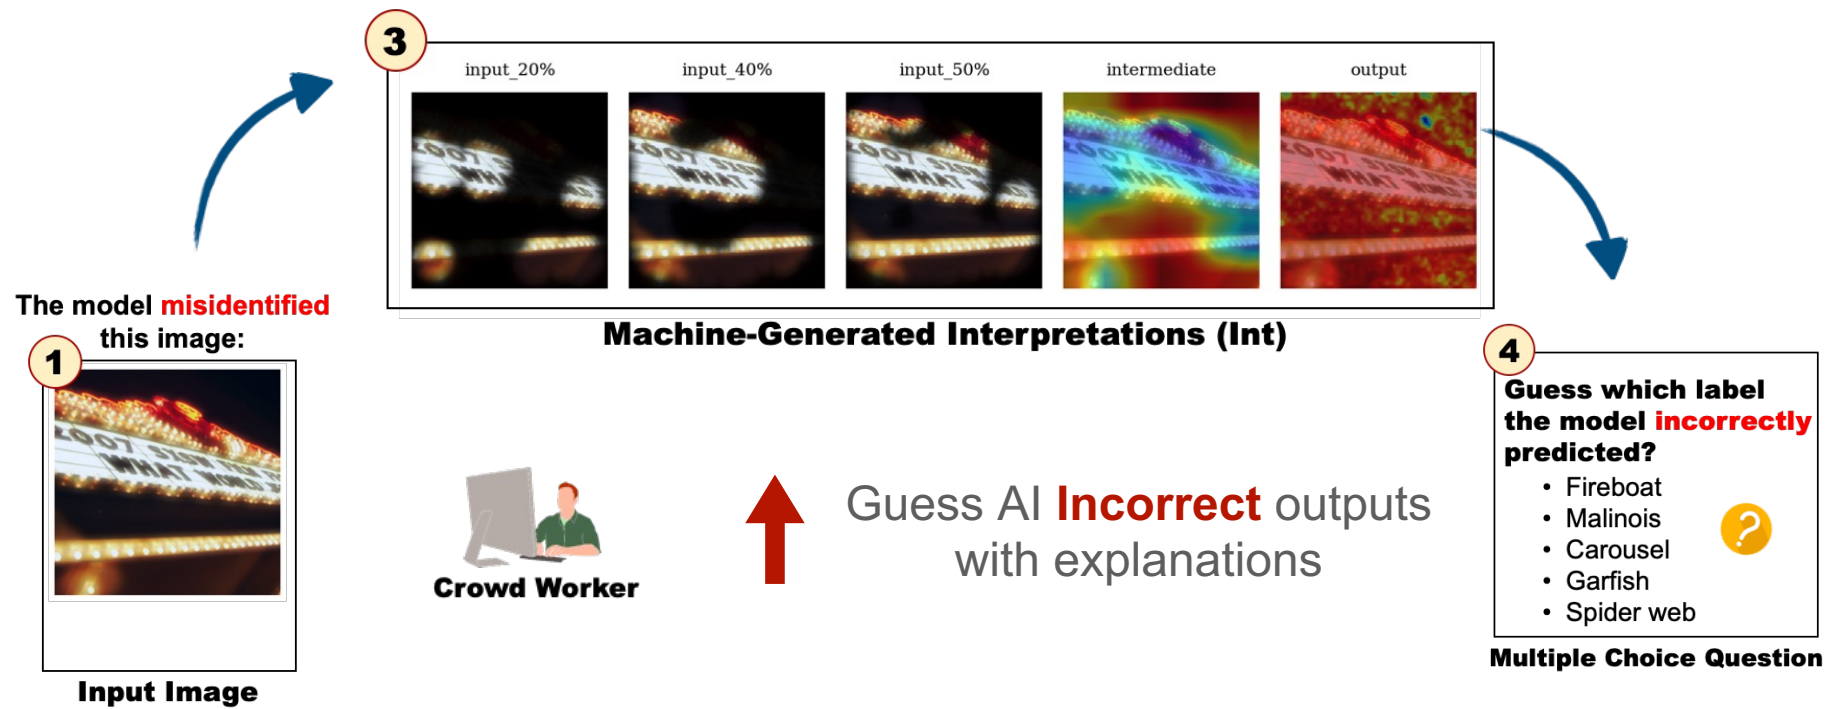


RQ1: Are XAI Useful for Humans?

NLP Interpretability

Vision Interpretability

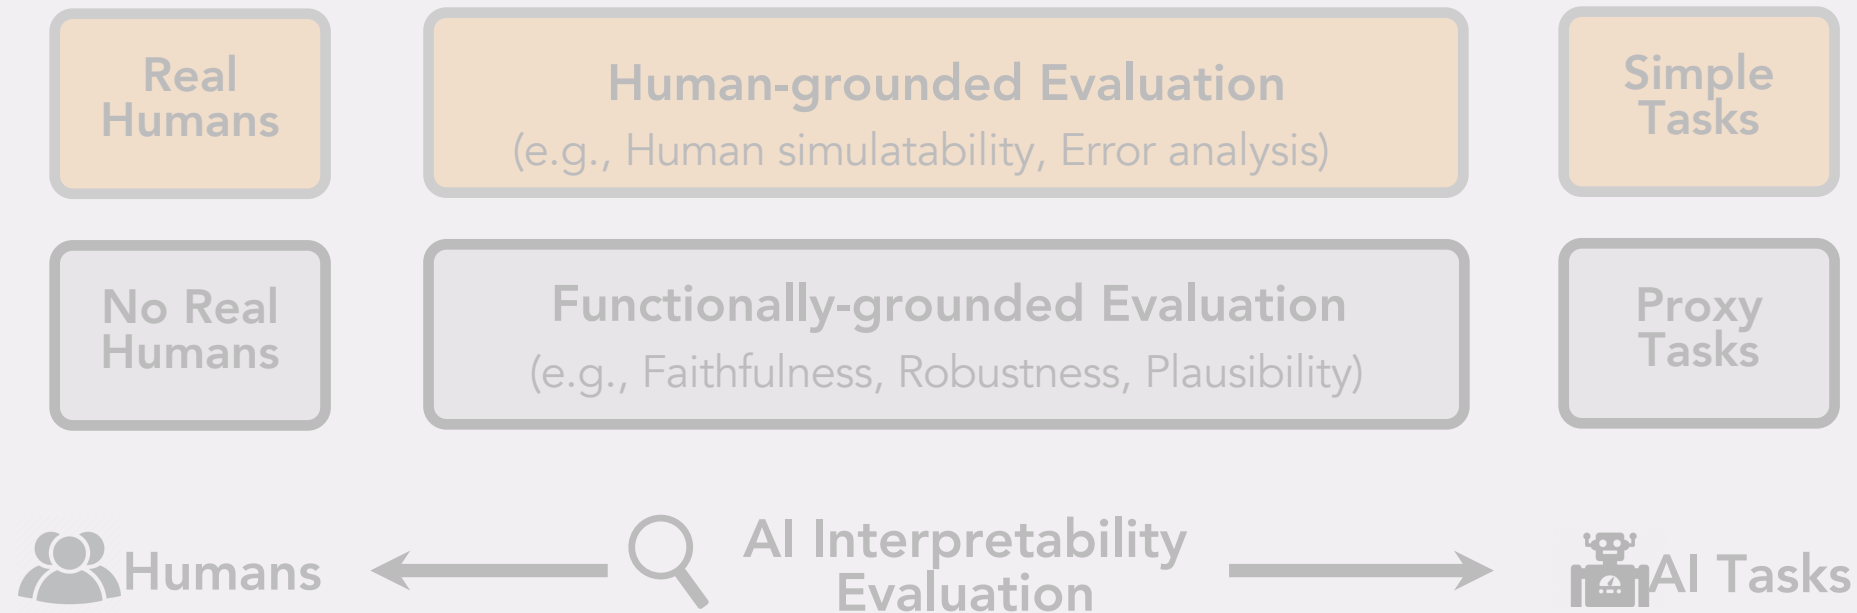




Visual AI explanations **did not increase**, but rather **decreased**, the **human's accuracy** in guessing the AI's **incorrect** decision-making.

XAI is **NOT** always **Useful** for Humans


AI explanations are **NOT always useful** for **humans** to understand the decision-making of **AI models** (including both language and vision models).



ACL 2022
22ND – 27TH MAY | 60TH MEETING | DUBLIN

HCOMP
2020

Background & Motivation



RQ1: Are XAI Useful for Humans?

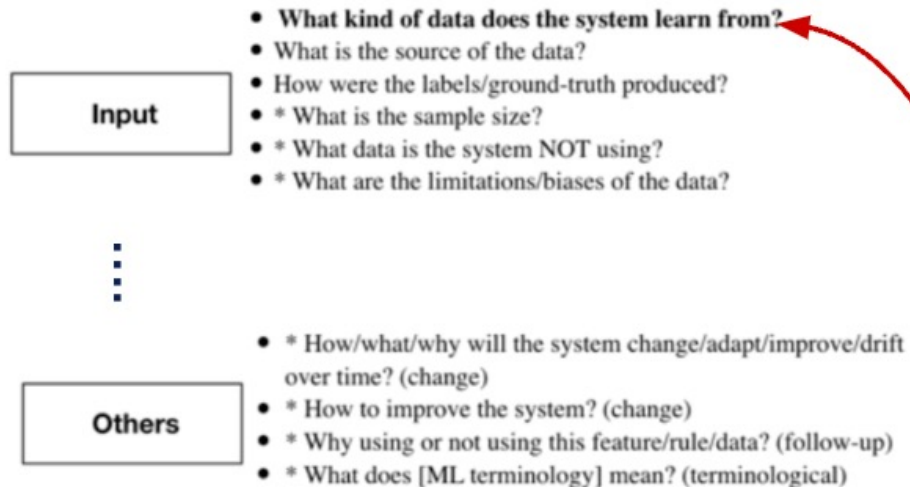
RQ2: Why?
(CHI 2021 Workshop)



Disparity between XAI with Humans?

43 User Questions in Practice

(Liao, Q. V., Gruen, D., & Miller, S. 2020)

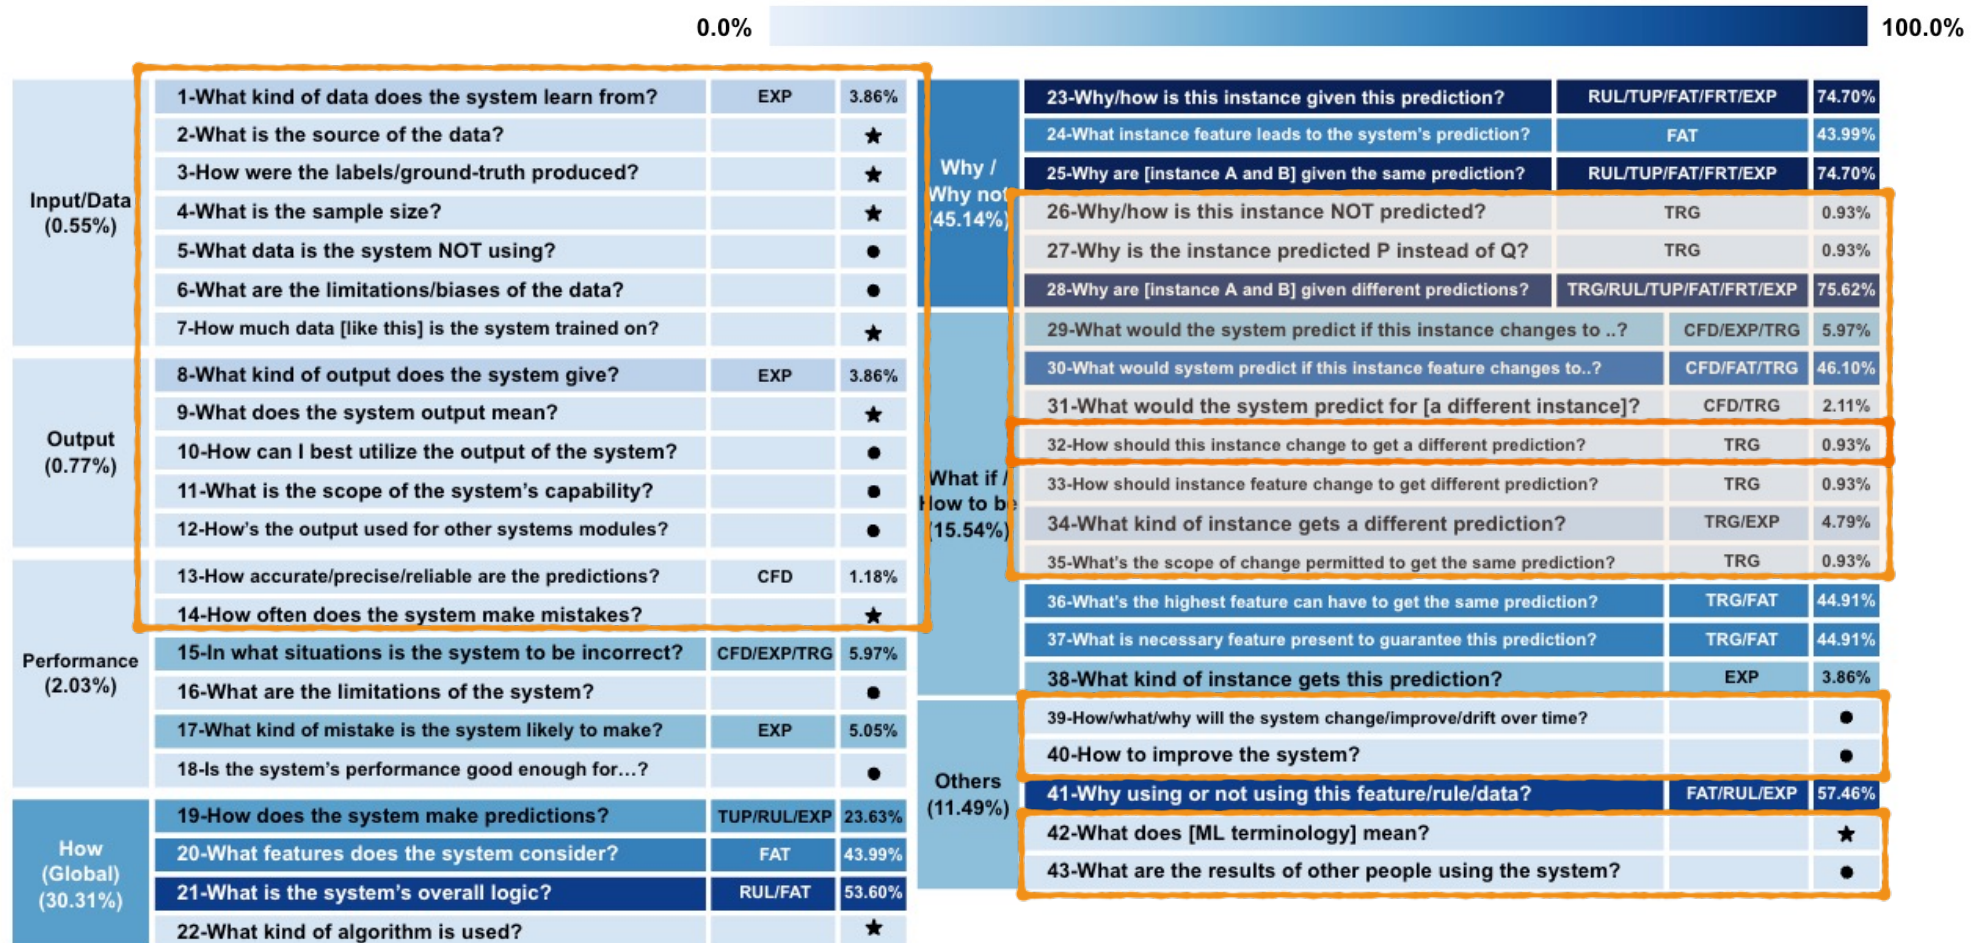


218 XAI Papers in NLP

ID	Title	Year	Venue	Paper URL
1	" Why should I trust you?" Explaining the predictions of any classifier	2016	KDD	https://arxiv.org/pdf/1602.04938
2	Visualizing and Understanding Neural Models in NLP	2016	NAACL	https://www.aclweb.org/anthology
3	Rationalizing Neural Predictions	2016	EMNLP	https://people.csail.mit.edu/taole
4	BERT Rediscovered the Classical NLP Pipeline	2019	ACL	https://www.aclweb.org/anthology
5	Attention is not Explanation	2019	NAACL	https://arxiv.org/pdf/1902.10186
⋮				
214	How much should you ask? On the question structure in QA systems	2018	BlackboxNLP	https://arxiv.org/pdf/1809.03734
215	Interpretable Multi-dataset Evaluation for Named Entity Recognition	2020	EMNLP	https://arxiv.org/pdf/2011.06854
216	A Survey of the State of Explainable AI for Natural Language Processing	2020	AACL-IJCNLP	https://arxiv.org/pdf/2010.00711
217	Explaining Simple Natural Language Inference	2019	ACL	https://www.aclweb.org/anthology
218	Understanding Neural Abstractive Summarization Models via Uncertainty	2020	EMNLP	https://arxiv.org/pdf/2010.07882

We match the **disparity** between the existing **200+ XAI papers** with **43 practical user questions!**

Existing XAIs largely Ignored...

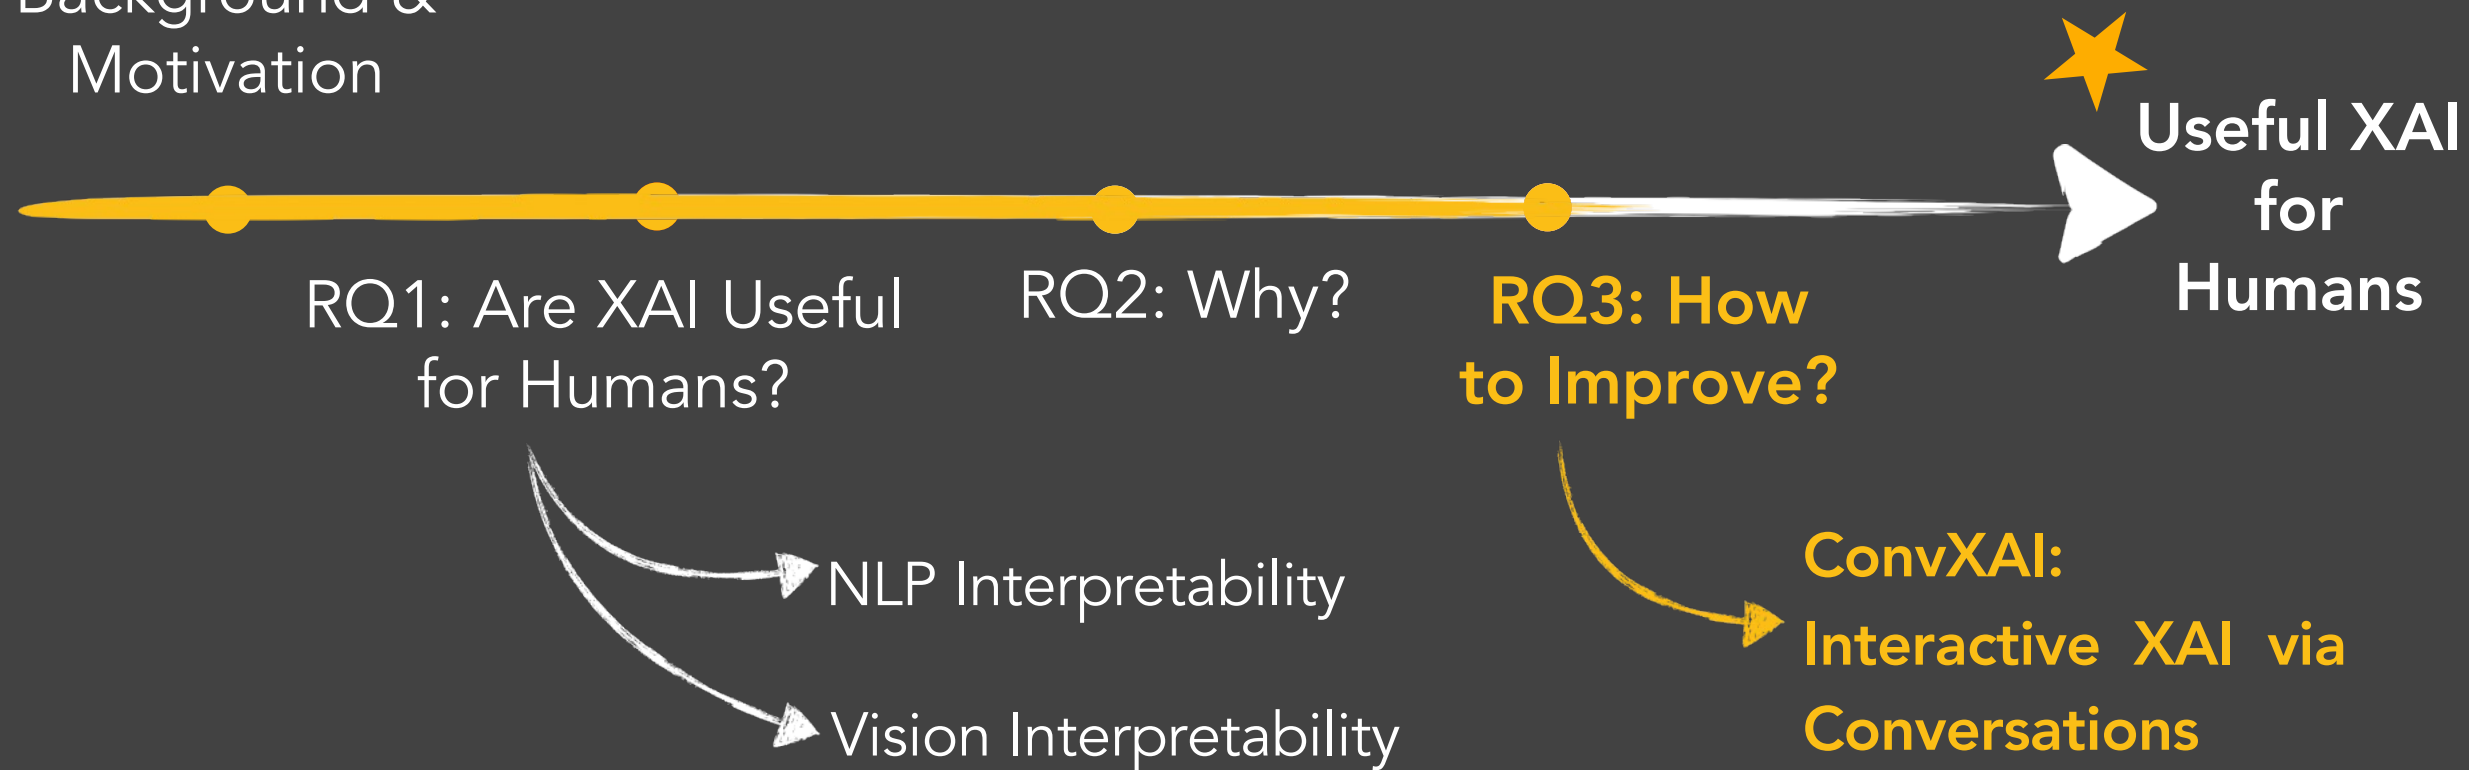


What AI systems **CANNOT** achieve (e.g., counterfactuals).



Diverse information across the whole AI lifecycle (data, model, deployment, etc.)

Background & Motivation



Best Demo

Challenges of Existing XAI

Humans




AI



Diverse User Needs
(Shen & Huang, CHI HCXAI, 2021)



ONE Explanation



Question	Response	Confidence	Explanation
1. What kind of data does the system have?
2. What is the reason for the result?
3. How were the features used in prediction?
4. What is the impact of the result?
5. What advice can the system give to improve the result?
6. How can I check the result of the system?
7. What does the system do?
8. How can I use the system?
9. How can I get help from the system?
10. How can I contact the system developer?
11. How can I report a bug to the system developer?
12. How can I provide feedback to the system developer?
13. How can I learn more about the system?
14. How can I get more information about the system?
15. How can I get more information about the system?
16. How can I get more information about the system?
17. How can I get more information about the system?
18. How can I get more information about the system?
19. How can I get more information about the system?
20. How can I get more information about the system?



Needs are NOT satisfied

?

?

?

- Showing **ONE** specific **explanation** might **NOT** meet **diverse XAI user needs**.

Challenges of Existing XAI



Diverse User Needs

MANY Explanations



Cognitive Overload

- Showing ONE specific explanation might NOT meet diverse XAI user needs.
- Showing **MANY explanations** at one time may lead to **cognitive overload** for humans

Solution: Conversational XAI



- Showing ONE specific explanation might NOT meet diverse XAI user needs.
- Showing MANY explanations at one time may lead to cognitive overload for humans

Diverse User Needs

XAI Candidate Pool

Question	Answer	Score
1. What kind of data does the system have access to?
2. How is the system trained?
3. How does the system handle errors?
4. How does the system handle user input?
5. How does the system handle user feedback?
6. How does the system handle user privacy?
7. How does the system handle user security?
8. How does the system handle user accessibility?
9. How does the system handle user performance?
10. How does the system handle user satisfaction?

Human-centered Conversational XAI empowers humans to interactively inquire the specific explanation with minimal cognitive load.

ConvXAI Demo:



Which conference are you most likely to submit this paper abstract to:

CHI (Human-Computer Interaction Domain)

Select an abstract example to try:

Select an abstract example

Or Edit your paper abstract:

Normal B I S U Q " ' < > < > < > < >

While various AI explanation (XAI) methods have been designed to gain insights into AI systems , it is still hard for users to acquire the information they need .

Prior work suggested using chatbots to dynamically cater to human needs , but little has been explored about how conversational AI should be designed .

S3.aspect=purpose
In this paper , we examine the Conversational XAI potential in the context of scientific writing .

Informed by human linguistics and formative studies , we identify four design principles of Conversational XAI : address various user questions (' multifaceted ') , provide details on-demand (' controllability ') , proactively suggest and accept follow-up questions (' mix-initiative ' and ' context-aware drill-down ') .

We instantiate them into an interactive prototype , CONVXAI , which allows writers to interact with various explanations through a chatbot interface .

Through 13 user studies , we show that 9 out of 13 participants preferred CONVXAI over the static interface baseline SelectXAI .

CONVXAI is promising to help users think through and address their diverse questions .

We are also aware of the limitation of CONVXAI , such as a steeper learning curve than baseline .

We conclude by discussing implications and challenges of conversational XAI systems .

Click to Submit Your Writing

Click below buttons to switch the model's prediction on each sentence.

Writing Structure Model Writing Style Model

A good paper abstract should describe comprehensive research aspects, this model (i.e., a SciBERT-based) classifies each sentence into one of the five aspect labels.

Background Purpose Method Finding/Contribution Other

Conversational Explainable AI (XAI) Assistant

To improve, you can check the most important words resulting in the prediction and further check how to revise input into another label . See XAI questions below:

Label Distribution Prediction Confidence

Similar Published Sentences

Which words are most important for this prediction?

How can I revise the sentence to get a different label?

How are the structure labels distributed?

We use the Research Aspects Model to generate aspect sequences of all 9935 paper abstracts. Then we cluster these sequences into five patterns as below. We compare your writing with these patterns for review.

Types	Patterns
Pattern1	'background' (42.9%) -> 'purpose' (14.3%) -> 'finding' (42.9%)
Pattern2	'background' (22.2%) -> 'purpose' (11.2%) -> 'method' (33.3%) -> 'finding' (33.3%)
Pattern3	'background' (33.3%) -> 'purpose' (16.7%) -> 'method' (16.7%) -> 'finding' (33.3%)
Pattern4	'background' (33.3%) -> 'method' (16.7%) -> 'finding' (50%)
Pattern5	'background' (20%) -> 'finding' (6.7%) -> 'background' (13.3%) -> 'purpose' (6.7%) -> 'background' (13.3%) -> 'finding' (6.7%) -> 'method' (6.7%) -> 'finding' (26.7%)

You can ask below XAI questions for the selected sentence:

In this paper , we examine the Conversational XAI potential in the context of scientific writing .

Send

Four Design Principles for useful conversational XAI

P1

Multifaceted XAI

Contain multiple XAI types that explain AI from various aspects

P2

Mixed-Initiative

Proactively send users XAI tutorials or hints to teach them "how to use XAI"

P3

Context-aware Drill-down

Maintain the conversation history to generate responses with user needs

P4

Controllability

Enable humans to customize XAI with personalized needs

Technical Challenges & Contributions

Challenges:

1. No unified approach for various XAI
2. No dialog system to parse XAI user questions and customization

Technical Contribution

- A Unified conversational XAI API for various XAI types that enable user to customize AI explanations.

```
In [3]: """Human-ConvXAI Interaction with the unified API"""  
### 10 Types of AI Explanation Questions  
user_xai_questions = {  
    "global-ask-data": "What data did the system learn from?",  
    "global-ask-model": "What kind of models are used?",  
    "global-ask-quality-score": "What's the range of the style quality scores?",  
    "global-ask-label-distribution": "How are structure labels distributed in the dataset?",  
    "global-ask-sent-length": "What's the statistics of the sentence lengths?",  
    "local-ask-model-confidence": "How confident is this prediction?",  
    "local-ask-xai-example": "What are some published sentences that look similar to mine semantically?",  
    "local-ask-feature-attribution": "Which words in this sentence are most important for this prediction?",  
    "local-ask-counterfactual": "How can I revise the input to get a different prediction?",  
    "others": "Who are you?"  
}
```

Global AI Explanations -- generating AI explanation for meta information

[XAI Type 1] - global AI explanation for describing dataset

```
In [4]: user_xai_question = user_xai_questions['global-ask-data']  
response = convxai_agent.explain(  
    user_xai_question,  
    ai_input,  
    ai_predict_output,  
    conference,  
    visualize=visualize  
)
```

what data did the system learn from?

Sure! We are comparing your writing with our collected **CHI Paper Abstract** dataset to generate the above review. The dataset includes **21643 sentences** in **3235 papers**.

[XAI Type 2] - global AI explanation for describing models

```
In [5]: user_xai_question = user_xai_questions['global-ask-model']  
response = convxai_agent.explain(  
    user_xai_question,  
    ai_input,  
    ai_predict_output,  
    conference,  
    visualize=visualize  
)
```

what kind of models are used?

Of course! The **Writing Structure Model** is a **SciBERT** based classifier finetuned with the **CODA-19** dataset. Also, the **Writing Style Model** is a **GPT-2** based generative model finetuned with **9935 abstracts** from **CHI**, **ACL** and **ICLR** papers (click the terms to view more).

[XAI Type 3] - global AI explanation for describing quality scores

Evaluate ConvXAI with real human studies

Who
is
studied



13 graduate researchers

Task1



8 researchers

Task2

When

09/2022 (90min)

12/2022 (90min) (**rejoin**)

How
it's
studied

1. **Two** think-aloud **scientific writing tasks**:
 - **Within-Subjects** Study: ConvXAI vs. Baseline
 - **Improve a paper's abstract**;
 - Paper **domains**: NLP, or HCI, or AI
2. Post **Survey** - Questionnaires
3. Semi-**structured Interviews**

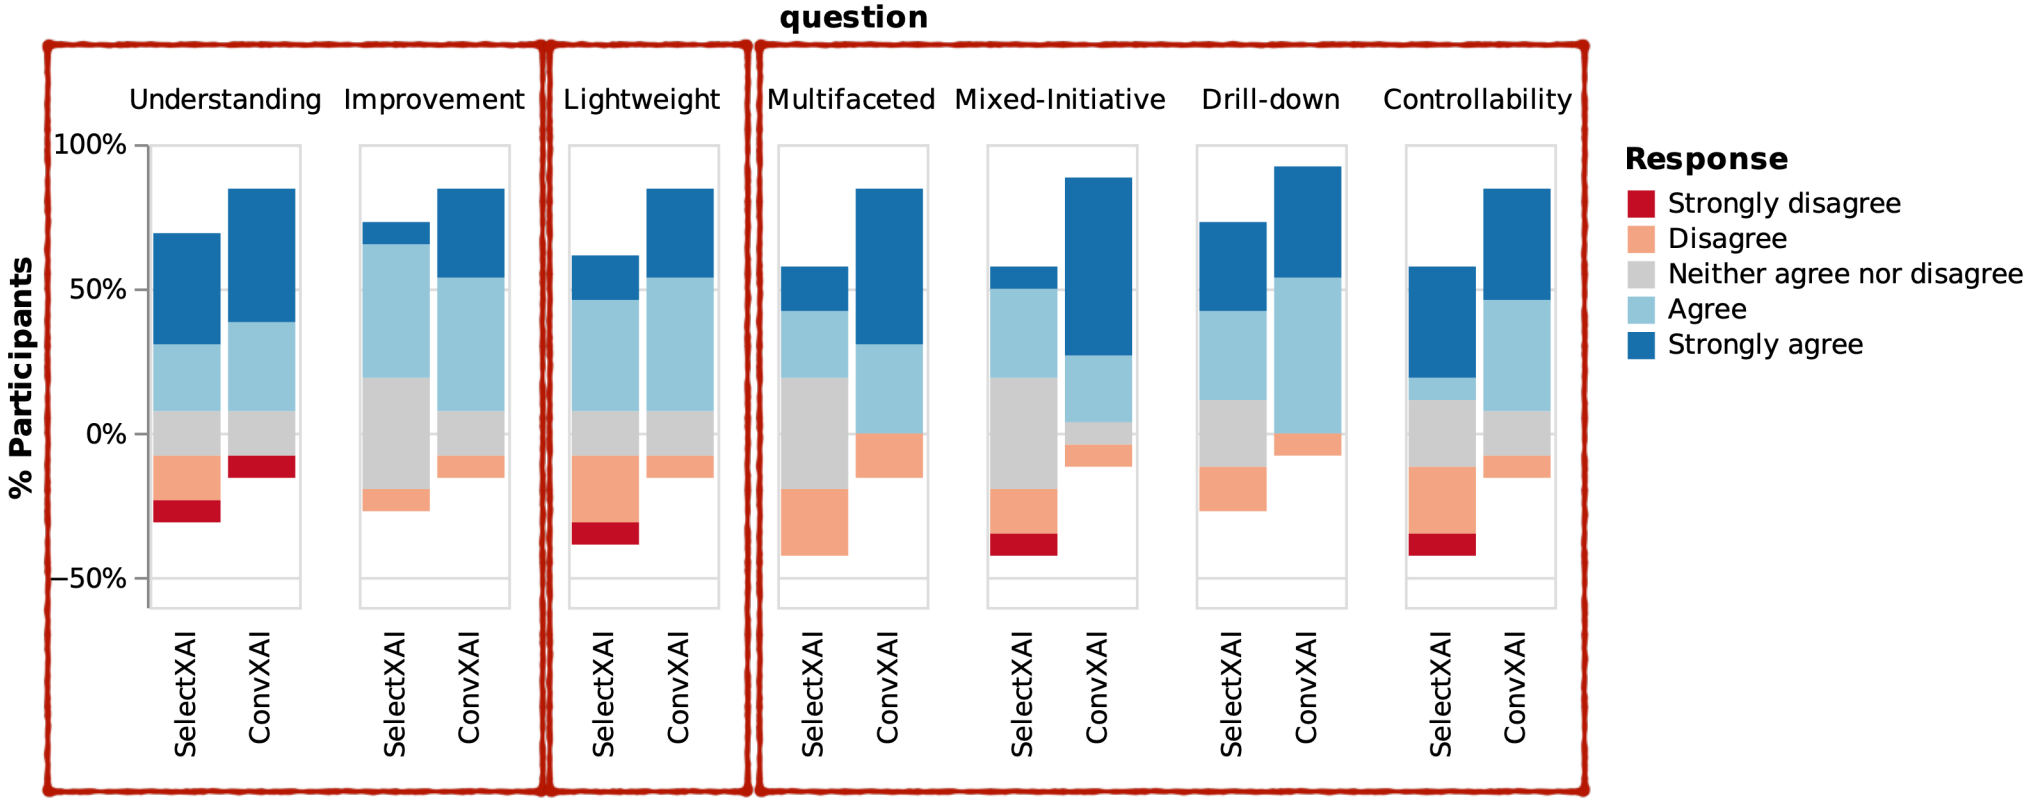
Baseline System (SelectXAI)

Within-Subjects Study Design

The image shows a composite screenshot of the 'Scientific Writing Support' and 'AI Explanation (XAI) Panel' interfaces. The left panel, titled 'Scientific Writing Support', includes a dropdown for 'Which conference are you most likely to submit this paper abstract to?' (set to 'CHI (Human-Computer Interaction Domain)'), another dropdown for 'Select an abstract example to try:', and a text editor for 'Or Edit your paper abstract:'. The text editor contains a paragraph of text with a highlighted sentence: 'In this paper, we examine the Conversational XAI potential in the context of scientific writing.' A red dashed box highlights this sentence, and a tooltip labeled 'A' points to it. Below the text editor is a 'Click to Submit Your Writing' button and a 'Click below buttons to switch the model's prediction on each sentence.' section with 'Writing Structure Model' and 'Writing Style Model' buttons. A tooltip labeled 'B' points to the 'Writing Structure Model' button. Below this are five buttons: 'Background', 'Purpose', 'Method', 'Finding/Contribution', and 'Other'. The right panel, titled 'AI Explanation (XAI) Panel', features a 'Writing Feedback Summary' section with a score of 3 out of 5. Below this is a 'Structure Suggestions' section with a tooltip labeled 'C' pointing to a highlighted sentence: 'Based on the sentence labels' percentage and order...'. Below the suggestions is a list of XAI options: 'Data Statistics (What data did the system learn from?)', 'Model Description (What kind of models are used?)', 'Quality Score (What's the range of the style quality scores?)', 'Aspect Distribution (How are the structure labels distributed?)', 'Sentence Length (What's the statistics of the sentence lengths?)', 'Prediction Confidence (How confident is the model for this prediction?)', 'Similar Examples (What are the most similar examples in the trainset?)', 'Important Words (Which words in this sentence are most important for this prediction?)', and 'Counterfactual Inputs (How can I revise the input to get a different prediction label?)'. A red arrow points from the 'Data Statistics' option in the XAI panel to the highlighted sentence in the 'Structure Suggestions' section.

Survey results of human study in Task1

Finding#1: **ConvXAI is a useful approach** to help end users understand and collaborate with AI models.



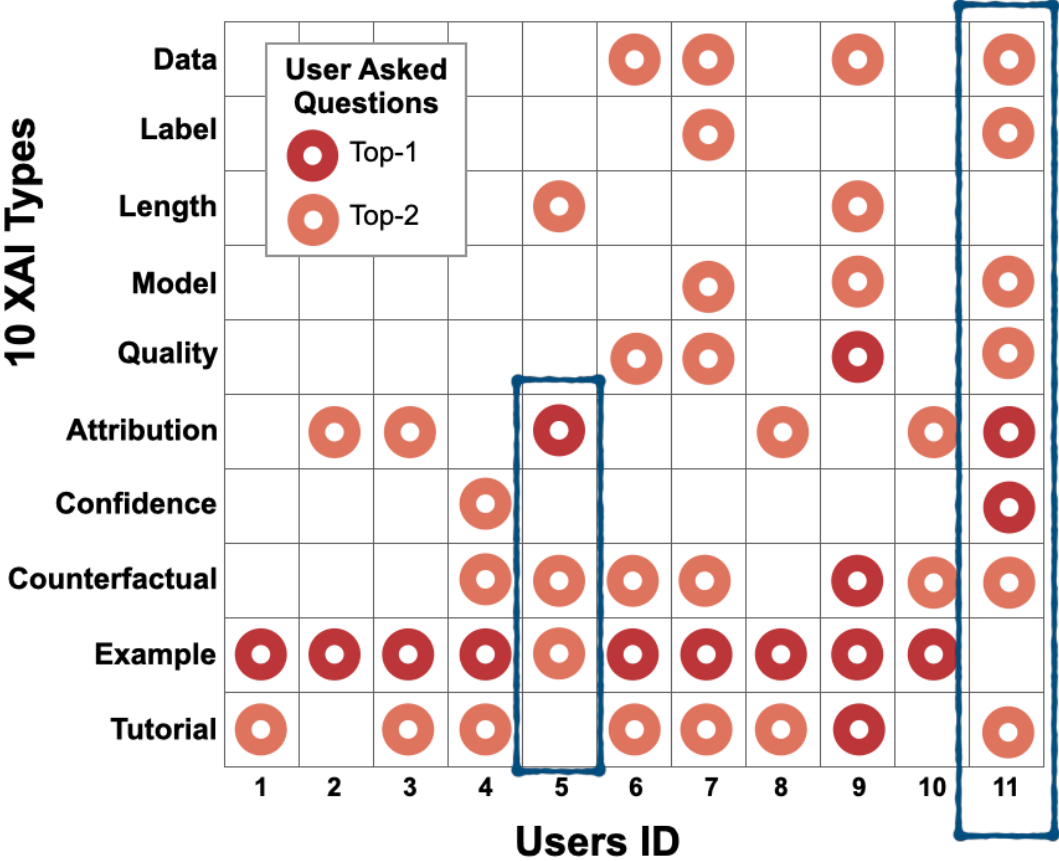
Useful in Understanding & Improving Writing

Less Cognitive Load

More aligned with human-centered design rationales


Humans' XAI usage patterns in Task1

Finding#2: Different users prefer to use different XAI formats in the real-world tasks.



Task1 v.s. Task2: user needs changed along time

Finding#3: **Users XAI needs changed along time** and converged to instance-wise XAIs.

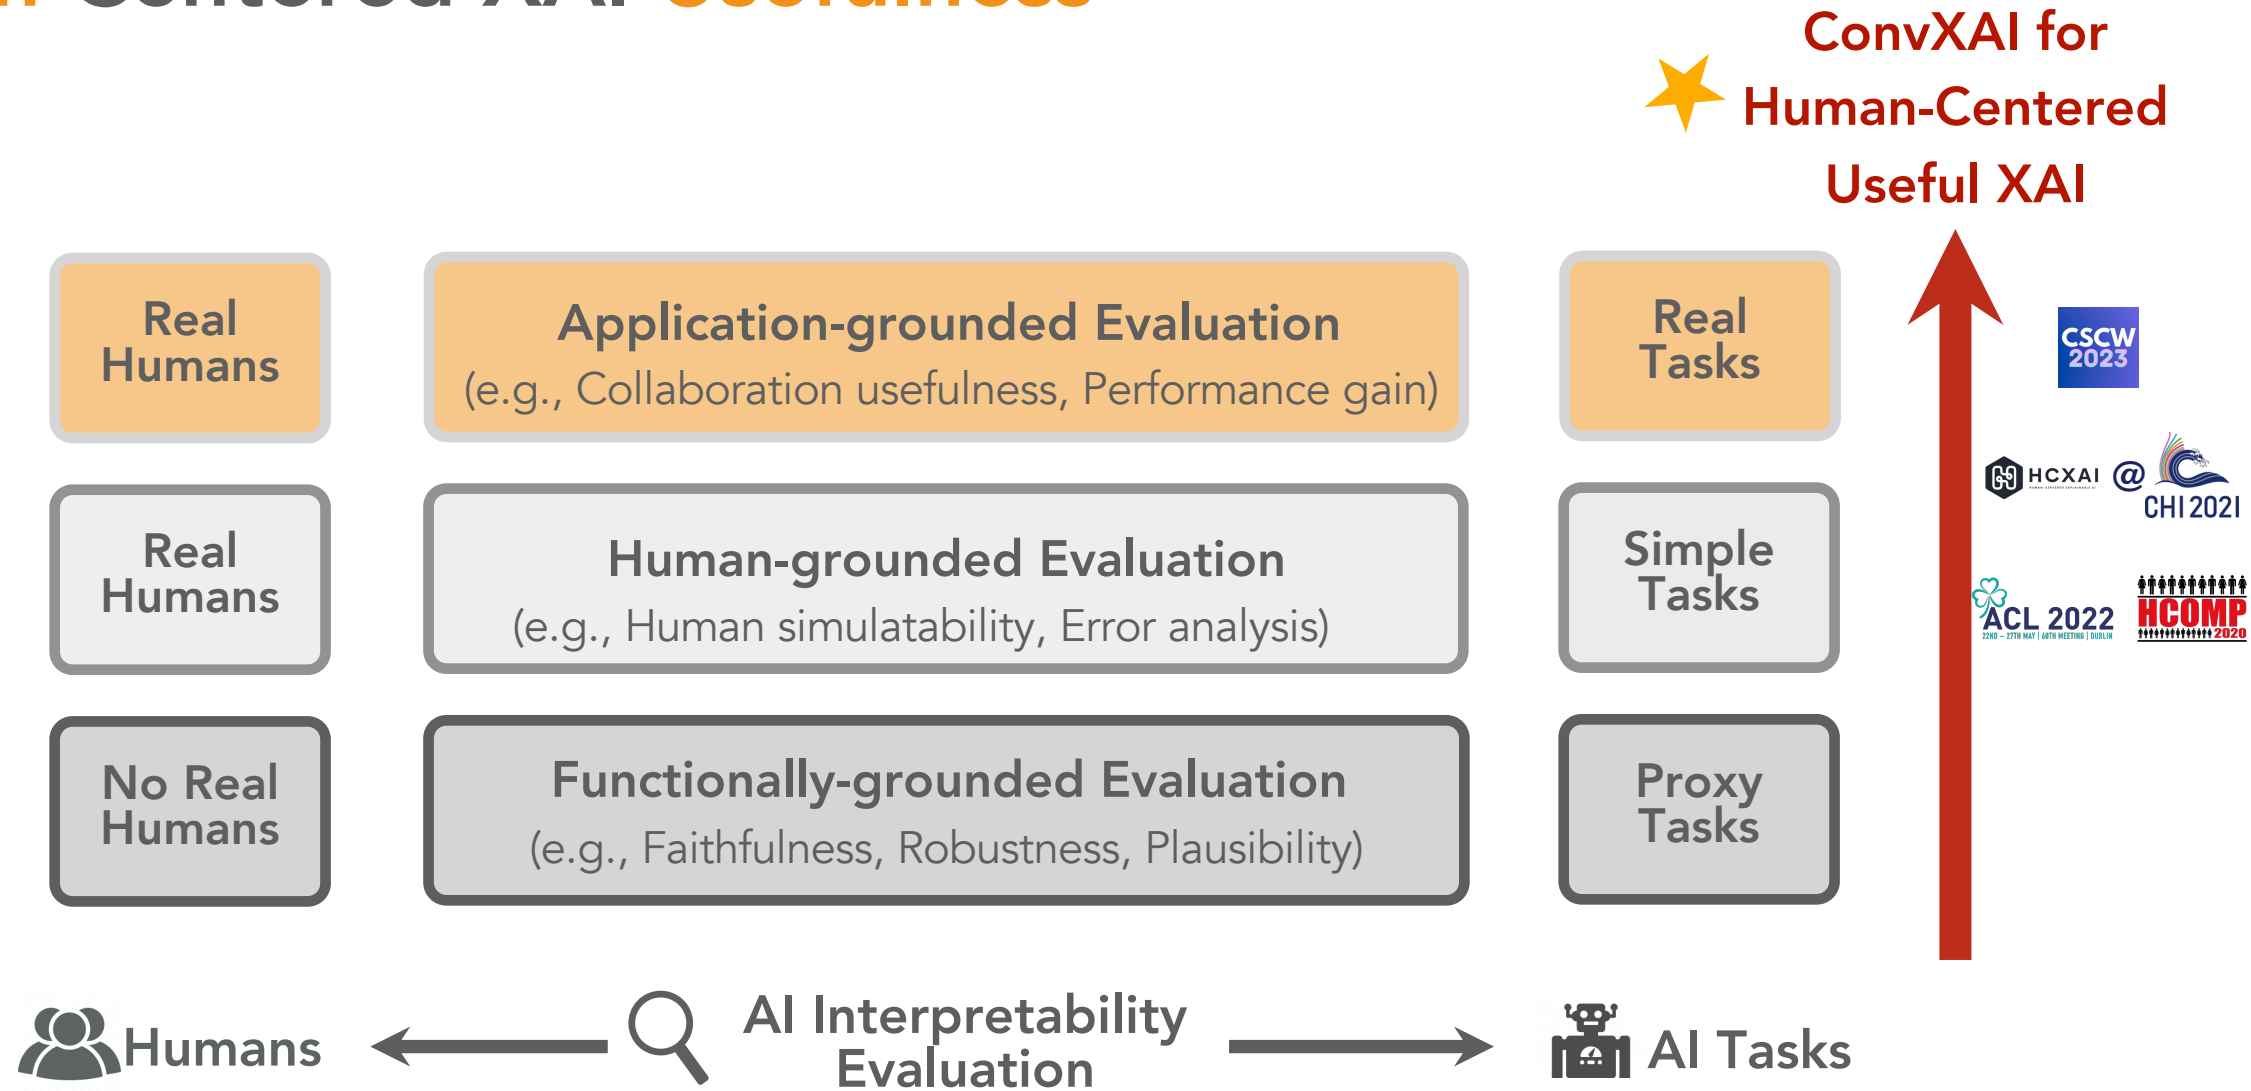


Finding#4: User-oriented **XAI Customization is important** in many XAI types.

Takeaway

ConvXAI is a potentially **useful human-centered XAI** approach that empowers humans to interactively inquire **heterogeneous AI Explanations via a simple conversation interface.**

Human-Centered XAI Usefulness



What's *Next* ...



ConvXAI: A *Start* of *Useful XAI* for *Humans*

- 1 Tools** *How to construct scalable interactive/conversational XAI tools for a wider range of human-AI collaboration tasks?*
- 2 Useful for Humans** *How to measure usefulness for humans and tailor interactive XAI to improve human performance?*
- 3 Useful for AIs** *How to collect human feedback from interactive XAI to improve AI model performance?*

Other Human-centered AI papers (2020 - 2023)

Keywords

Human-annotated AI dataset

Conversational XAI for Human

Human-AI Interactive System

Human Evaluation on LLM

Human-AI Co-writing Eval

Human Eval on NLP XAI

Human-AI Agent Interact Tool

Fairness on Speaker Verification

Human-in-the-loop Speech

Survey of 200+ XAI Papers

Human Eval on CV XAI

XAI Robustness

AI Adversarial & Security

1. [Hua Shen](#), Vicky Zayats, Johann Rocholl, Dan Walker, and Dirk Padfield. MultiTurnCleanup: A Benchmark for Multi-Turn Spoken Conversational Transcript Cleanups. [EMNLP 2023](#) 🏆 **Google Research Scholarships**
2. [Hua Shen](#), Chieh-Yang Huang, Tongshuang Wu, Ting-Hao (Kenneth) Huang. ConvXAI: Delivering Heterogeneous AI Explanations via Conversations to Support Human-AI Scientific Writing. [CSCW 2023 Demo](#). 🏆 **Best Demo Award**
3. Tongshuang Wu, [Hua Shen](#), Daniel S Weld, Jeffrey Heer, Marco Tulio Ribeiro. ScatterShot: Interactive In-context Example Curation for Text Transformation. [IUI 2023](#). 🏆 **Best Paper Honorable Mention**
4. [Hua Shen](#)*, Adaku Uchendu*, Jooyoung Lee*, Thai Le, Ting-Hao 'Kenneth' Huang, and Dongwon Lee. Does Human Collaboration Enhance the Accuracy of Identifying Deepfake Texts? [AAAI HCOMP 2023](#)
5. [Hua Shen](#), Tongshuang Wu. Parachute: Evaluating Interactive Human-LM Co-writing Systems. [CHI 2023 In2Writing](#) Workshop
6. [Hua Shen](#), Tongshuang Wu, Wenbo Guo, Ting-Hao (Kenneth) Huang. Are Shortest Rationales the Best Explanations For Human Understanding? [ACL 2022](#)
7. Binfeng Xu, Xukun Liu, [Hua Shen](#), Zeyu Han, Yuhan Li, Murong Yue, Zhiyuan Peng, Yuchen Liu, Ziyu Yao, Dongkuan Xu. Gentopia.AI: A Collaborative Platform for ToolAugmented LLMs. [EMNLP 2023 Demo](#)
8. [Hua Shen](#)*, Yuguang Yang*, Guoli Sun, Ryan Langman, Eunjung Han, Jasha Droppo, Andreas Stolcke. Improving Fairness in Speaker Verification via Group-adapted Fusion Network. [ICASSP 2022](#).
9. Shih-Hong Huang, Chieh-Yang Huang, Yuxin Deng, [Hua Shen](#), Szu-Chi Kuan, and TingHao 'Kenneth' Huang. Too Slow to Be Useful? On Incorporating Humans in the Loop of Smart Speakers. [AAAI HCOMP 2022 WiP/Demo](#)
10. [Hua Shen](#), Ting-hao (Kenneth) Huang. Explaining the Road Not Taken. [CHI 2021 HCXAI](#) Workshop
11. [Hua Shen](#), Ting-hao (Kenneth) Huang. How Useful Are the Machine-Generated Interpretations? A Human Evaluation on Guessing the Wrongly Predicted Labels. [AAAI HCOMP 2020](#)
12. Xinyang Zhang, Ningfei Wang, [Hua Shen](#), Shouling Ji, Ting Wang. Interpretable Deep Learning under Fire. [USENIX 2020](#)
13. Ren Pang, [Hua Shen](#), Xinyang Zhang, Shouling Ji, Yevgeniy Vorobeychik, Xiapu Luo, Alex X. Liu, Ting Wang. The Tale of Evil Twins: Adversarial Inputs versus Poisoned Models. [ACM CCS 2020](#)
14. 15. CHI 2024 Under Review.....

Acknowledgment!!




Advisors & Ph.D. Committee



RisingStar Mentor & Organizers



Mentors & Collaborators





Hua Shen

 huashen@umich.edu  @huashen218

University of Michigan