香港浸會大學
HONG KONG BAPTIST UNIVERSITY

DEPARTMENT OF
COMPUTER SCIENCE
HONG KONG BAPTIST UNIVERSITY
香港浸會大學計算機科學系

Improving Personalized Explanation Generation through Visualization

Lei Li

Hong Kong Baptist University

csleili@comp.hkbu.edu.hk

Dec. 8, 2022

Invited Talk at University of Luxembourg

Outline

- **Explainable Recommendation**
- Natural Language Explanation (ACL’21)
- Visual Explanation (ACL’22)
- Future Work

Recommendations Everywhere

E-commerce
(taobao.com)

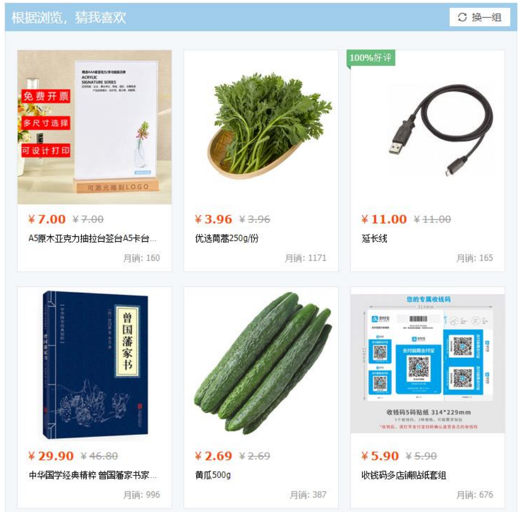

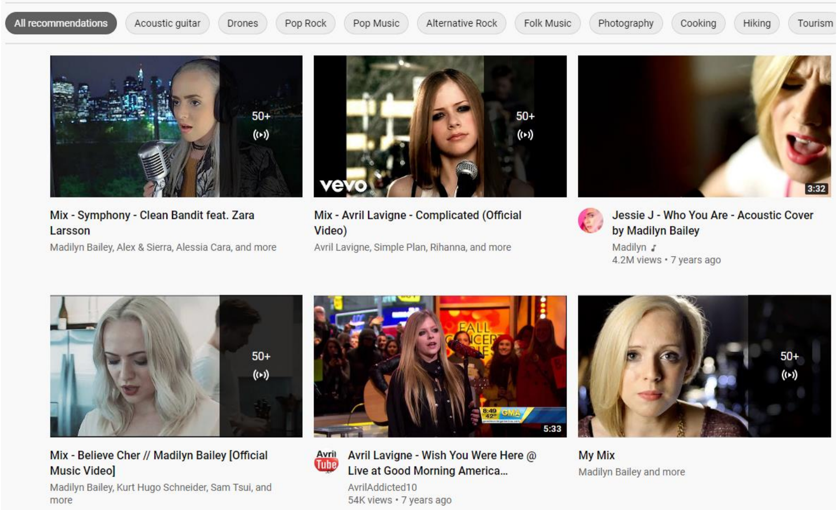

Video-streaming
(youtube.com)

Social Network
(instagram.com)

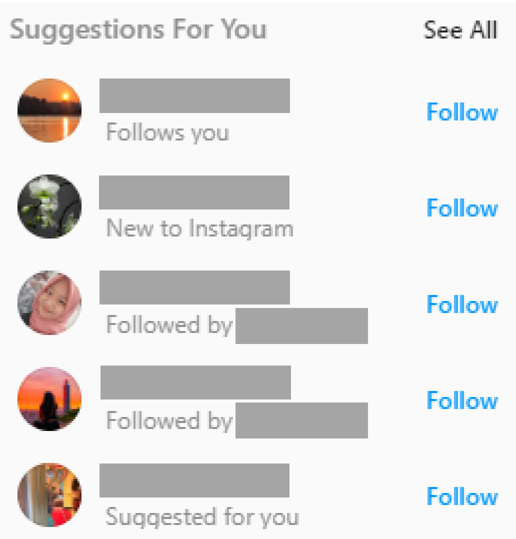

Movie
(movie.douban.com)

Industrial Applications of Explanations

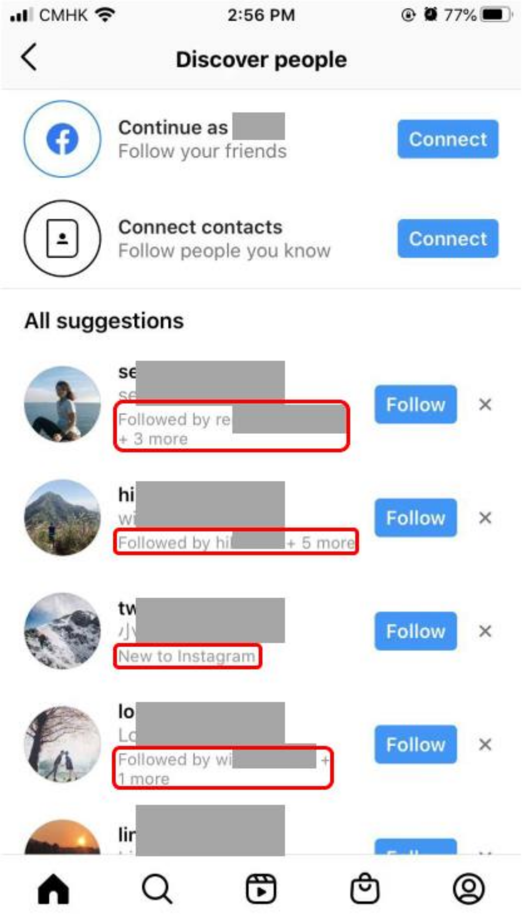

Instagram
(instagram.com)

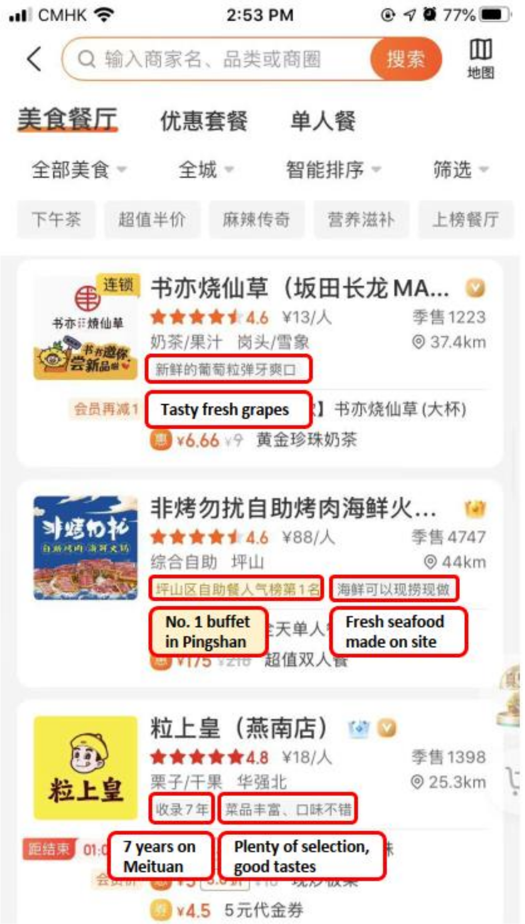

Meituan
(meituan.com)

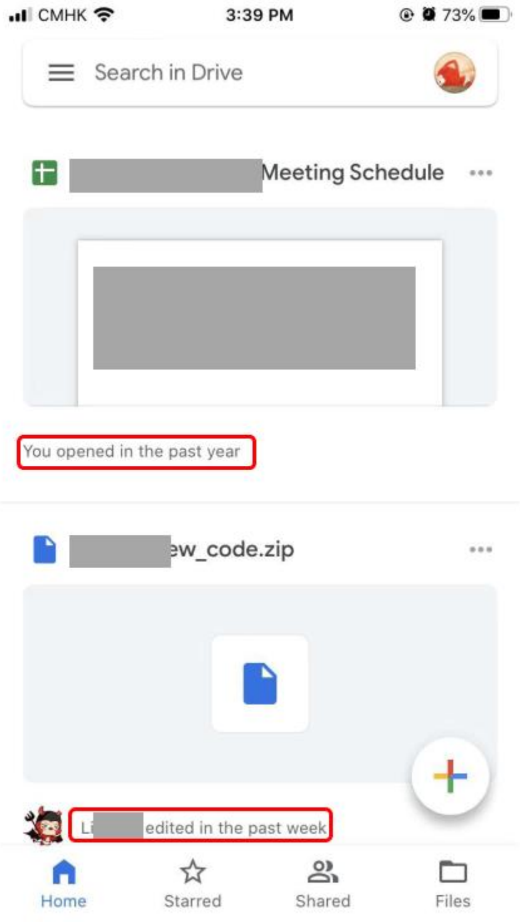

Google Drive
(drive.google.com)

Explanatory Goals (Tintarev and Mashoff, 2015)

- **Trust**: increase users’ confidence in the system
- **Effectiveness**: help users make good decisions
- **Persuasiveness**: convince users to try or buy
- **Efficiency**: help users make decisions faster
- **Satisfaction**: increase the ease of use or enjoyment

User-centric

- **Transparency**: explain how the system works
- **Scrutability**: allow users to tell the system it is wrong

System-centric

Typical Explanation Styles

- Item Features
- Templates
- Highlights
- Review Segments
- Generated Text
- ……

Item Features

- Selected features about the user or the item
- Typical models
  - Tripartite Graph (He et al., CIKM’15)
  - Decision Tree (Wang et al., WWW’18)

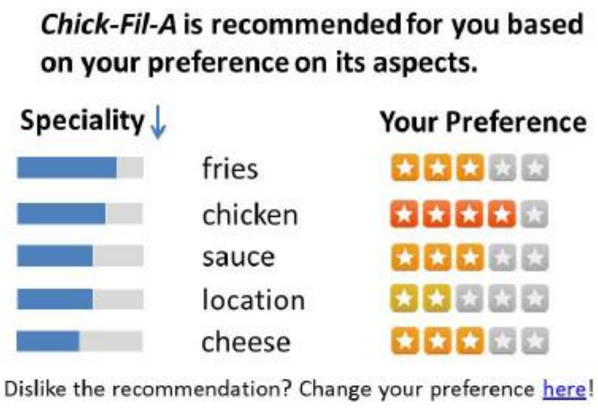

Courtesy image from TriRank (He et al., CIKM’15)

v_1	[User Country=UK] & [User Style=Art and Architecture Lover] ⇒ [Item Attribute=Concerts and Shows] & [Item Tag=Imelda Staunton]
v_{22}	[User Age=35-49] & [User Country=UK] ⇒ [Item Tag=**Camden Town**] & [Item Rating=4.0]
v_{130}	[User Age≠ 25-34] & [User Gender=Female] & [User Style=Peace and Quiet Seeker] ⇒ [Item Attribute=Sights & Landmarks] & [Item Tag=**Walk Around**]
v_{148}	[User Age≠ 50-64] & [User Country≠USA] ⇒ [Item Tag=**Top Deck & Canary Wharf**]
v_{336}	[User Age=35-49] & [User Country=UK] & [User Style=Art and Architecture Lover] ⇒ [Item Tag=Royal Opera House] & [Item Tag=Interval Drinks]

Courtesy image from TEM (Wang et al., WWW’18)

Templates (1)

- Item features filled in pre-defined templates
- Typical methods
  - Matrix Factorization (Zhang et al., SIGIR'14)
  - Tensor Factorization (Wang et al., SIGIR'18)
  - Counterfactual Reasoning (Tan et al., CIKM'21)

You might be interested in [feature], on which this product performs well.
You might be interested in [feature], on which this product performs poorly.

Courtesy image from EFM (Zhang et al., SIGIR'14)

• **Amazon** Recommendation: *Superleggera/Dual/Layer/Protection/case*
Explanation: *Its* ***grip*** *is [firmer] [soft] [rubbery]. Its* ***quality*** *is [sound] [sturdy] [smooth]. Its* ***cost*** *is [original] [lower] [monthly].*
• **Yelp** Recommendation: *Smash/Kitchen&Bar*
Explanation: *Its* ***decor*** *is [neat] [good] [nice]. Its* ***sandwich*** *is [grilled] [cajun] [vegan]. Its* ***sauce*** *is [good] [green] [sweet].*

Courtesy image from MTER (Wang et al., SIGIR'18)

Templates (2)

- Data sources
  - Concept Graph (Gao et al., AAAI'19)
  - Knowledge Graph (Wang et al., AAAI'19)
- Typical models
  - Attention (Gao et al., AAAI'19)
  - Long Short-Term Memory (Wang et al., AAAI'19)

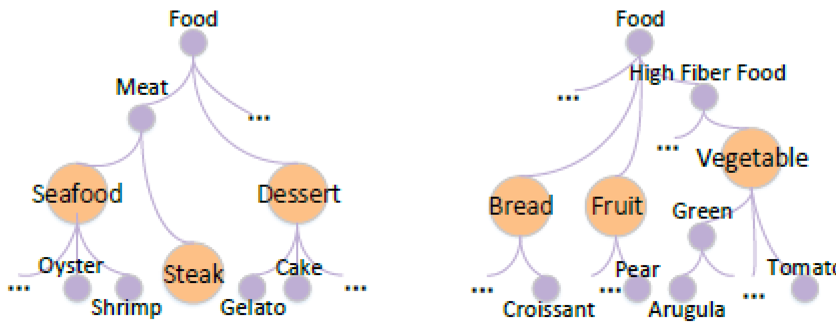

You might be interested in [features in E], on which this item performs well.

Courtesy image from DEAML (Gao et al., AAAI'19)

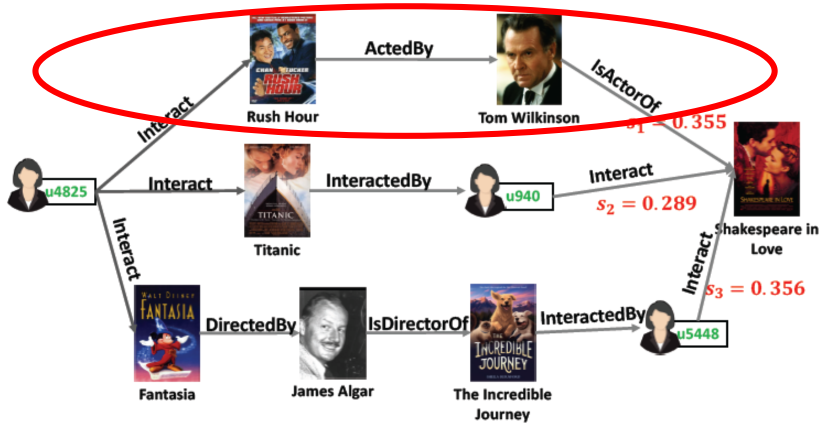

Shakespeare in Love is recommended since you have watched **Rush Hour** acted by the same actor **Tom Wilkinson**.

Courtesy image from KPRN (Wang et al., AAAI'19)

Highlights

- Words in reviews or regions in images highlighted by **attention**
- Typical models
  - Convolutional Neural Network (Seo et al., RecSys'17; Chen et al., SIGIR'19)
  - Gated Recurrent Unit (Lu et al., WWW'18)

I bought this ebook (16G) for my kindergarten and elementary children to read books on trips and my old child to check emails. It does the jobs well until now. I personally like it very much for its excellent hardware performance. With low price, fast response, and light weight, book size, and Barn and Nobles support, It is the best device for children when you want to have something as an alternative for your computer. I am a fan of Amazon.com and meant to buy a kindle for my children. But the displayed sample on my local Bestbuy store showed me that the nook tablet responded much faster than kindle fire. It could not be generally true, but based on the displayed tablets, I had to choose the nook. Until now, I have had mainly happy experiences with it. It has apps for children, one of them has many Smithsonian videos that my small children love the most.

Courtesy image from TARMF (Lu et al., WWW'18)

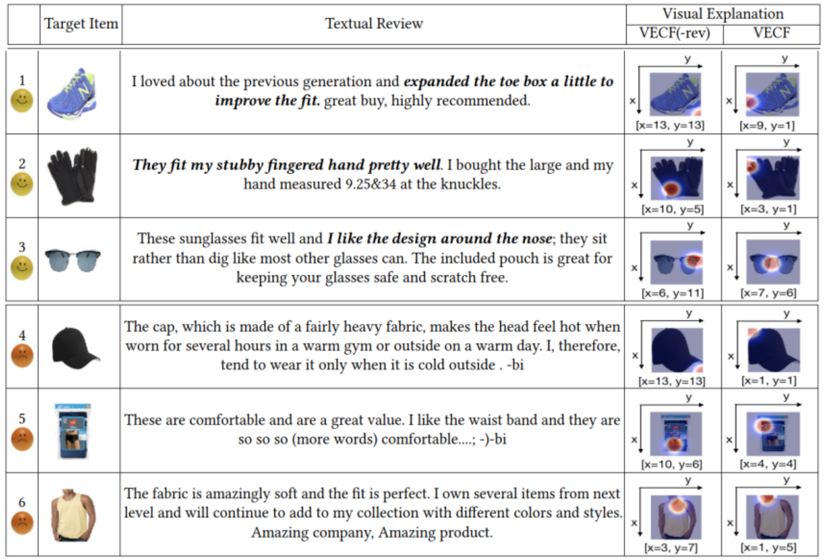

	Target Item	Textual Review	Visual Explanation	
			VECF(-rev)	VECF
1		I loved about the previous generation and ***expanded the toe box a little to improve the fit.*** great buy, highly recommended.	[x=13, y=13]	[x=9, y=1]
2		***They fit my stubby fingered hand pretty well.*** I bought the large and my hand measured 9.25&34 at the knuckles.	[x=10, y=5]	[x=3, y=1]
3		These sunglasses fit well and ***I like the design around the nose***; they sit rather than dig like most other glasses can. The included pouch is great for keeping your glasses safe and scratch free.	[x=6, y=11]	[x=7, y=6]
4		The cap, which is made of a fairly heavy fabric, makes the head feel hot when worn for several hours in a warm gym or outside on a warm day. I, therefore, tend to wear it only when it is cold outside . -bi	[x=13, y=13]	[x=1, y=1]
5		These are comfortable and are a great value. I like the waist band and they are so so so (more words) comfortable....; -)-bi	[x=10, y=6]	[x=4, y=4]
6		The fabric is amazingly soft and the fit is perfect. I own several items from next level and will continue to add to my collection with different colors and styles. Amazing company, Amazing product.	[x=3, y=7]	[x=1, y=5]

Courtesy image from VECF (Chen et al., SIGIR'19)

Review Segments

- Selected reviews or their segments (mostly by **attention**)
- Typical methods
  - Convolutional Neural Network (Chen et al., WWW'18; Catherine and Cohen, RecSys'17)
  - Reinforcement Learning (Wang et al., ICDM'18)
  - Gated Recurrent Unit (Chen et al., AAAI'19)

These brushes are great quality for children's art work. They seem to last well and the bristles stay in place very well even with tough use.
I bought it for my daughter as a gift.
From beginning to end this book is a joy to read. Full of mystery, mayhem, and a bit of magic for good measure. Perfect flow with excellent writing and editing.
I like reading in my spare time, and I think this book is very suitable for me.

Courtesy image from NARRE (Chen et al., WWW'18)

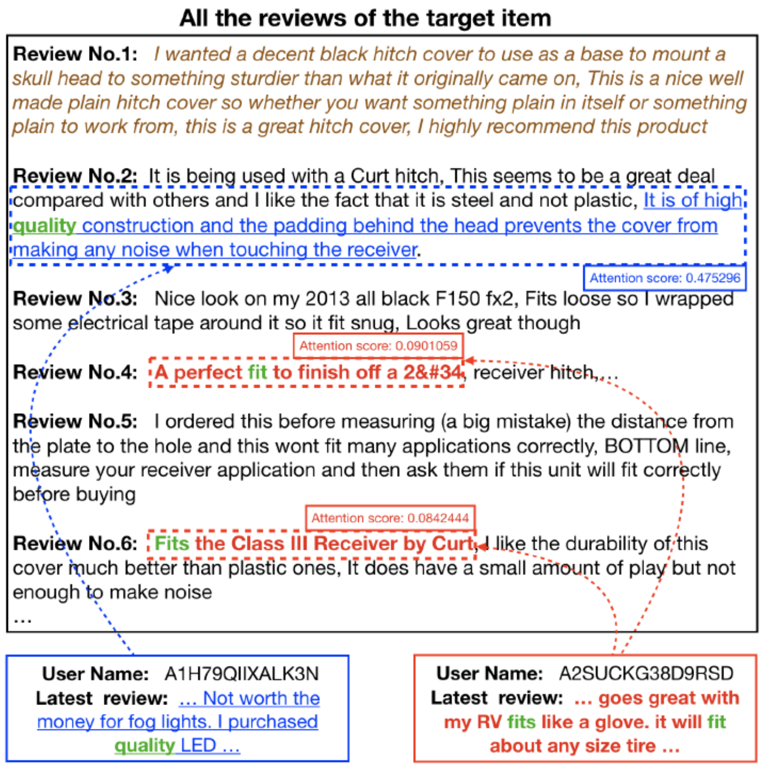

Courtesy image from DER (Chen et al., AAAI'19)

Generated Text

- Generated review sentences by natural language generation
- Typical models
  - Generative Adversarial Nets (Lu et al., RecSys'18)
  - Gated Recurrent Unit (Li et al., CIKM'20)
  - Transformer (Li et al., ACL'21)
  - Pre-trained Language Model (Li et al., arXiv'22)

Explanation
the **rooms** are spacious and the bathroom has a large tub the pool area is nice and the gym is very well equipped <eos> the rooms were clean and comfortable <eos>
beautiful **lobby** and nice bar the bathroom was large and the shower was great <eos> the lobby was very nice and the rooms were very comfortable <eos>

PETER (Li et al., ACL'21)

Rating	Feature	Explanation
4		***The rooms are spacious and the bath-room has a large tub.***
3.90	bathroom	The **bathroom** was large and had a sep-arate shower.
	tub	The bathroom had a separate shower and **tub**.
	rooms	The **rooms** are large and comfortable.
4		***The rooms are brilliant and ideal for business travellers.***
4.13	rooms	The **rooms** are very spacious and the **rooms** are very comfortable.
2		***The broken furniture and dirty sur-faces are a dead giveaway.***
2.96	furniture	The **furniture** is worn.
4		***Ideal for plane spotters and very close to the airport.***
2.76	airport	It is not close to the **airport**.

NETE (Li et al., CIKM'20)

Outline

- Explainable Recommendation
- **Natural Language Explanation (ACL'21)**
- Visual Explanation (ACL'22)
- Future Work

Autoregressive Natural Language Generation

- Predict future tokens based on past tokens
  - Generate an output sequence, based on the given input sequence

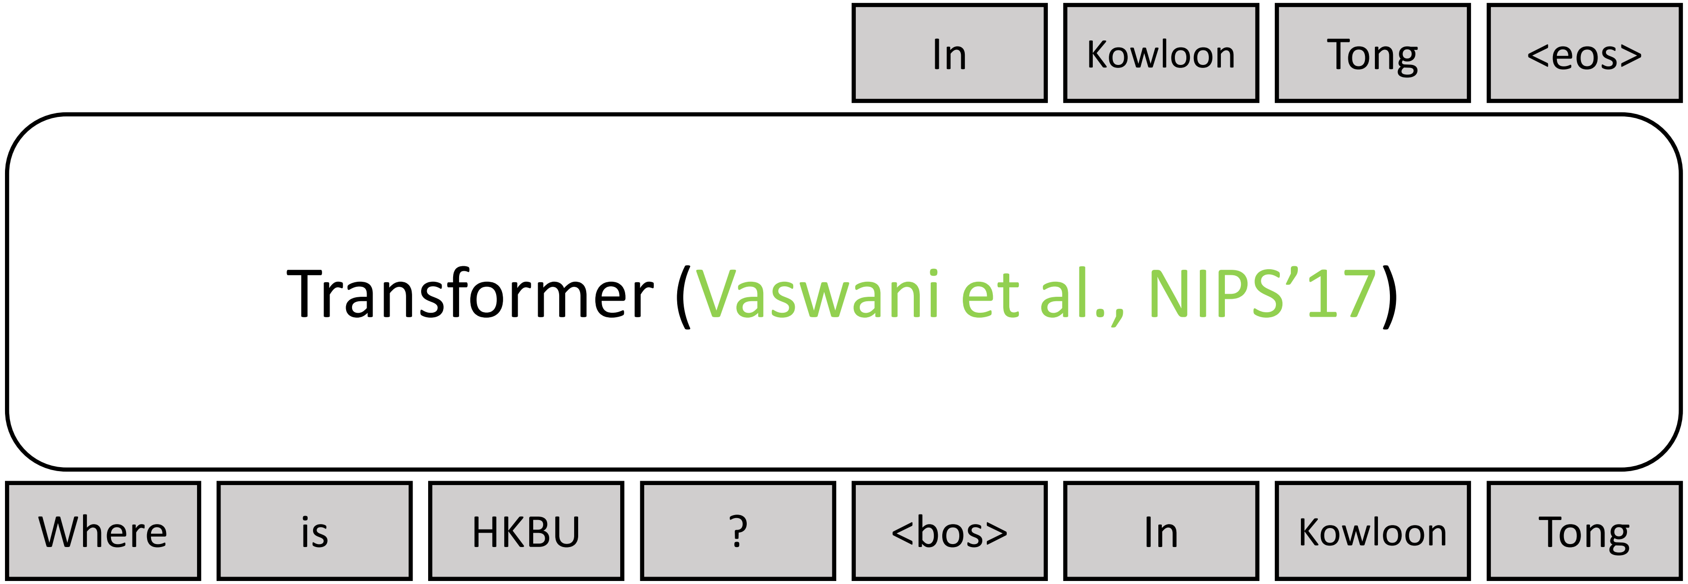

Explanation Generation with Transformer

- Regard the user-item pair as an input sequence
  - Treat the IDs as tokens, just like words

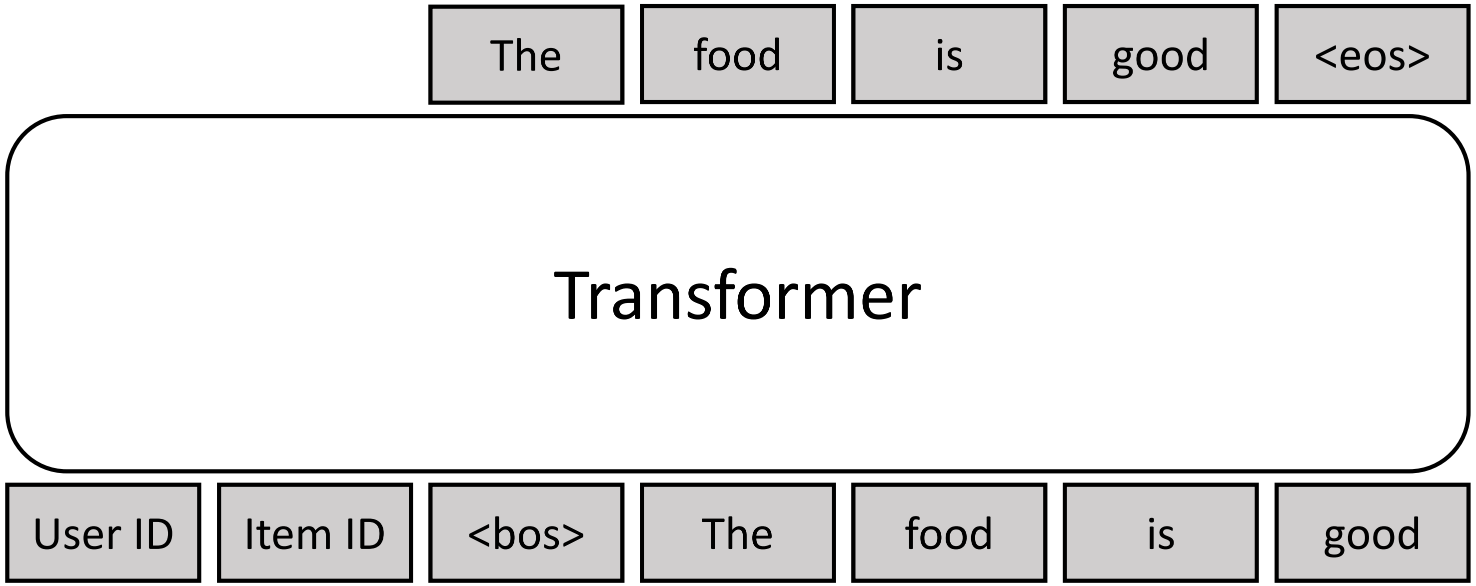

Problem Identification

- Identical generated explanations for every user-item pair
    - Adam Main Canteen
    - Beth Renfrew Cafeteria
    - Carol Bistro Bon
    - David Harmony Cafeteria
    - ……

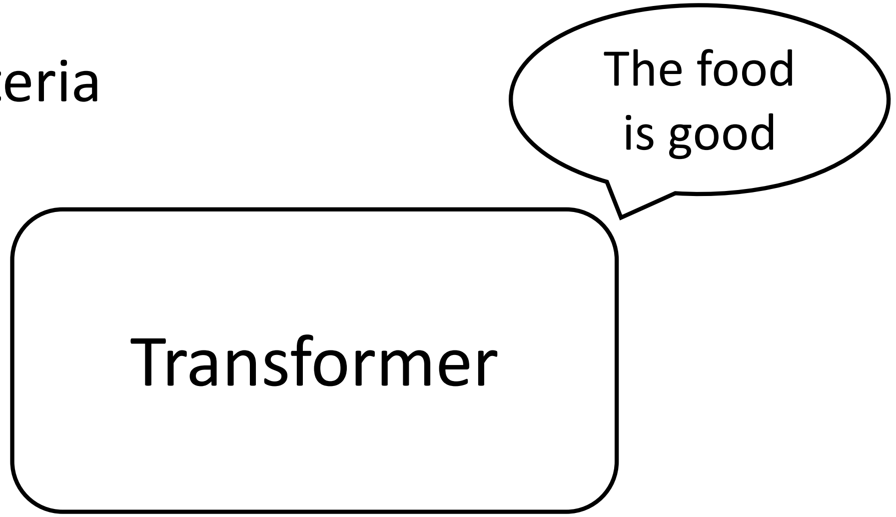

- **Less useful, if unable to explain the key specialty of each recommendation**
- **May cause negative effects on users (**Tintarev and Mashoff, 2015**)**

Attention Visualization

- The generation relies heavily on <bos>
  - The reason why all explanations are identical
- No attention weights on user ID and item ID
  - Model insensitive to IDs

Why insensitive?

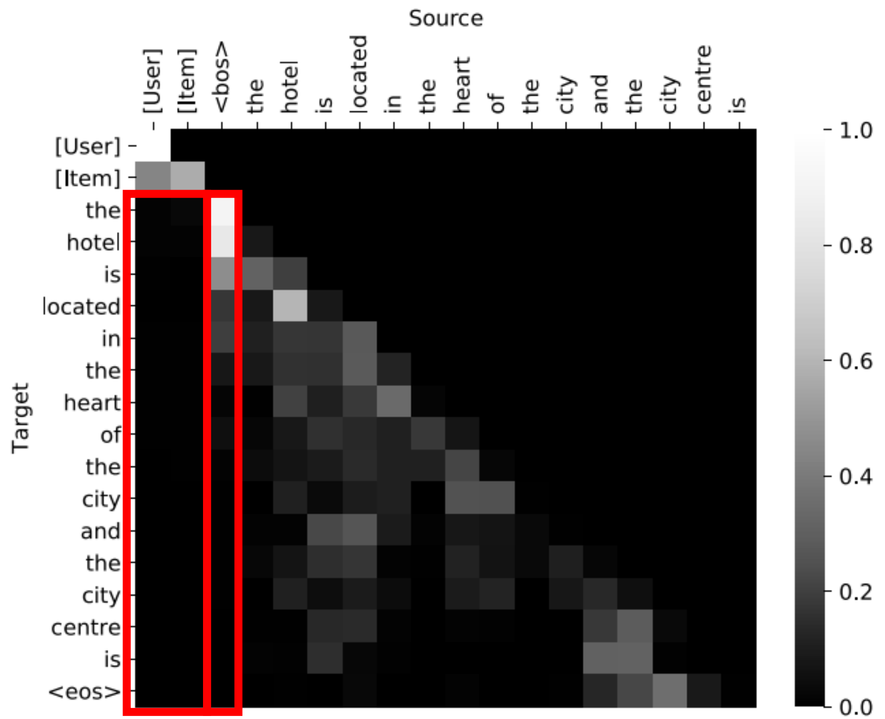

Problem Analysis

- Frequency mismatch between IDs and words
  - One user/item ID vs. hundreds of words in a review
  - An ID appears in only a few reviews
- IDs being regarded as uncommon words (OOV tokens)

Shelby S. Elite '2021
Saint Paul, MN
166 354 771

12/4/2019

6 photos

Ho Lee Fook was one of the best food spots I went to in HK. At first I was skeptical because sometimes the fusion or westernized type Asian restaurants are all for the look but don't taste great. But, Ho Lee Fook was beautiful inside and the food was amazing. We ordered the pan fried thick rolled noodles and the massive bone steak (forgot the actually name) but you won't miss it on the menu. The noodles were crispy and seasoned just right. The steak was so tender and delicious. It came with a jalapeño sauce on the plate which complimented it so well.

While being here I forgot I was in HK because everyone spoke English and the menu was also in English! The entrance is so cute with the lucky cats all on the walls.

If you are visiting HK or live there I definitely recommend giving this place a try! It is a little on the pricey side but for the atmosphere it is expected.

Restaurant review
(yelp.com)

Solution: Context Prediction

- Bridge IDs and words with this task

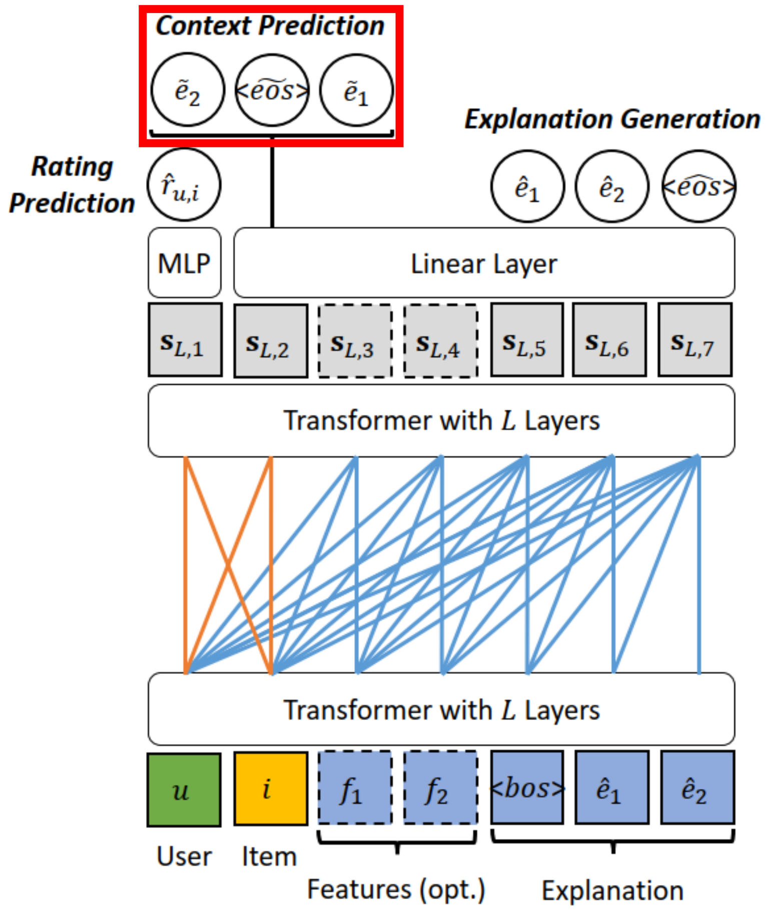

Attention Visualization Again

- PETER can utilize IDs for explanation generation

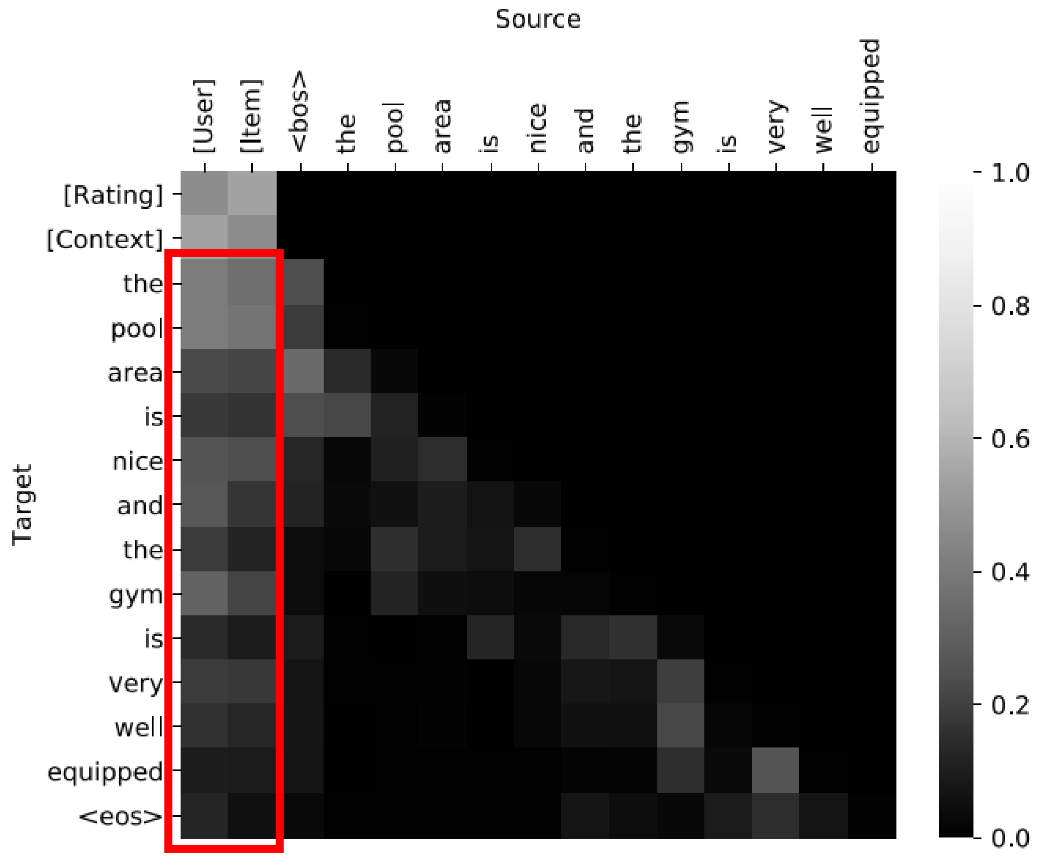

Context Prediction vs Explanation Generation

- Context prediction: predict explanation words in one step
- Explanation generation: generate them one by one

$$\mathcal{L}_c = \frac{1}{|\mathcal{T}|} \sum_{(u,i)\in\mathcal{T}} \frac{1}{|E_{u,i}|} \sum_{t=1}^{|E_{u,i}|} -\log c_{2}^{e_t}$$

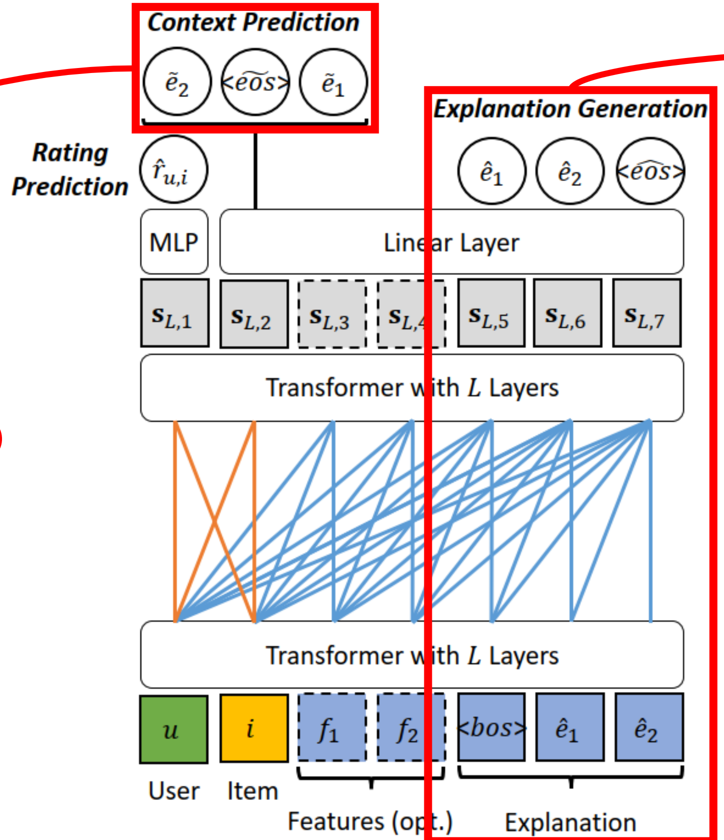

$$\mathcal{L}_e = \frac{1}{|\mathcal{T}|} \sum_{(u,i)\in\mathcal{T}} \frac{1}{|E_{u,i}|} \sum_{t=1}^{|E_{u,i}|} -\log c_{2+|F_{u,i}|+t}^{e_t}$$

Recommendation & Targeted Explanation

- Predict a rating score for the user-item pair
- Incorporate item features for targeted explanation generation
  - E.g., conversational recommendation (Chen et al., IJCAI'20)

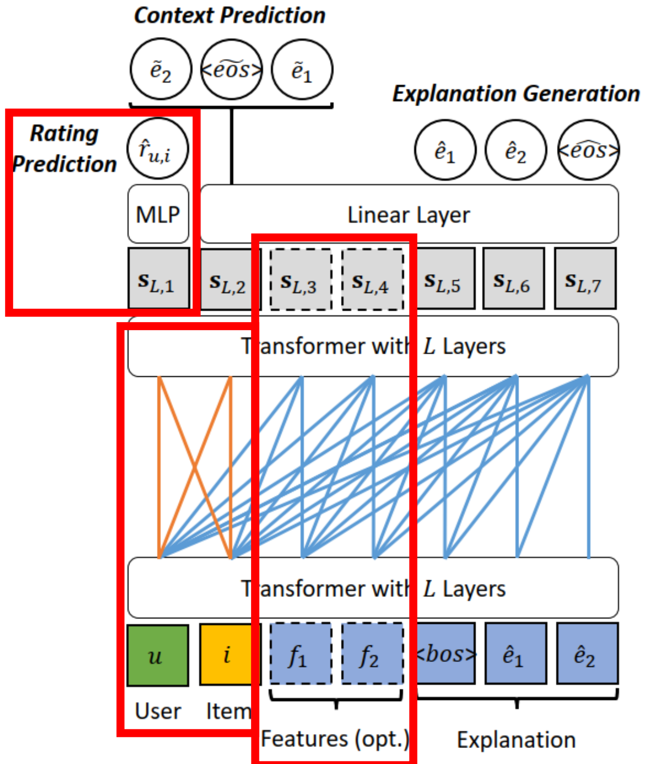

Outline

- Explainable Recommendation
- Natural Language Explanation (ACL’21)
- **Visual Explanation (ACL’22)**
- Future Work

Visual Tokens

- An image can also be represented as a sequence of discrete tokens.
  - The codebook (vocabulary) is constructed by vector-quantization.

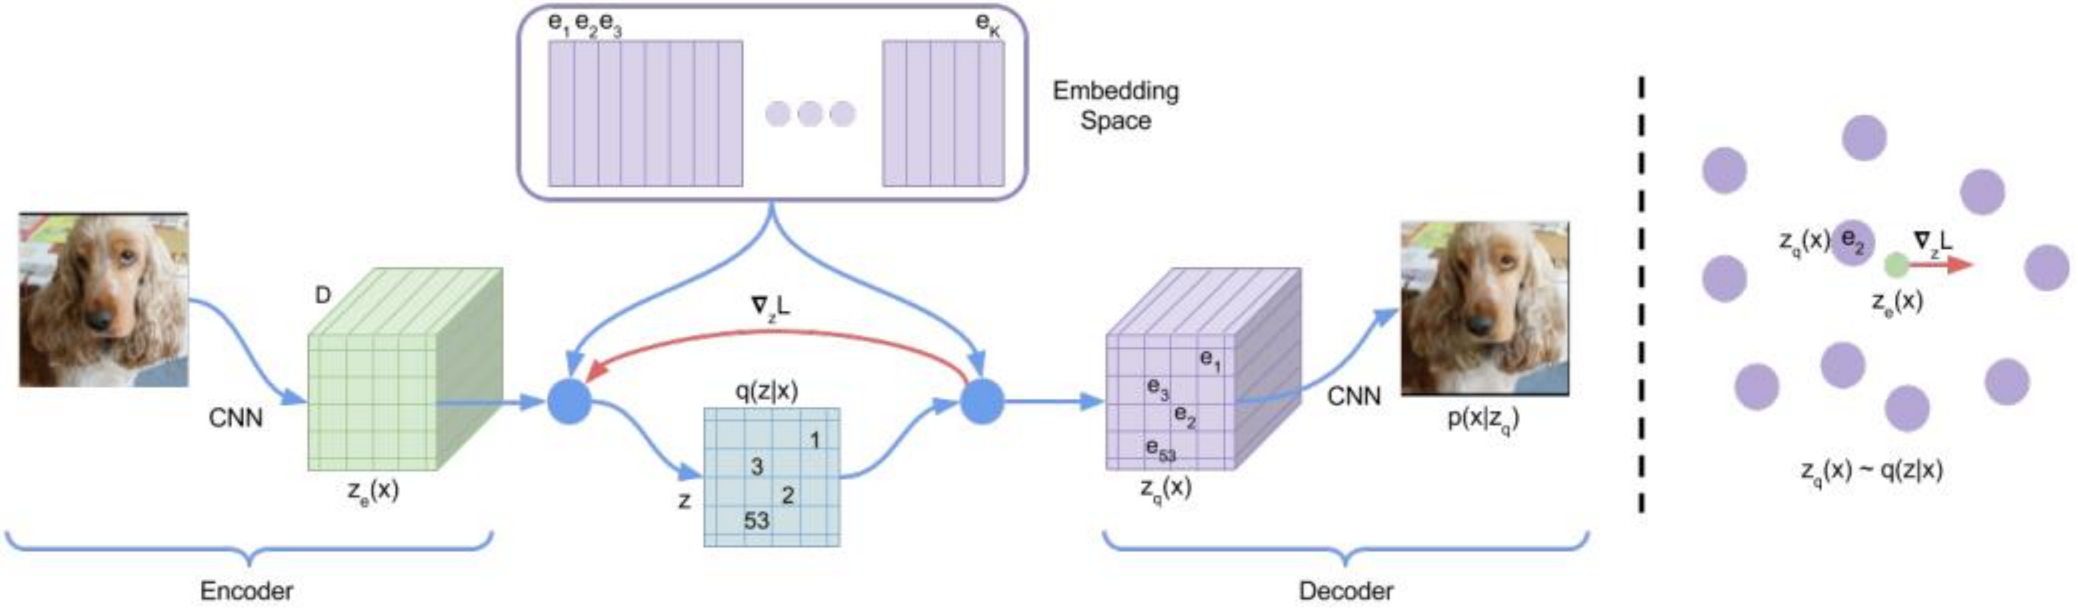

Courtesy image from VQ-VAE (Oord et al., NIPS'17)

Technical Details

- Suppose we have trained the encoder $\mathcal{E}$, decoder $\mathcal{G}$, and codebook $\mathcal{Z} = \{z_k\}_{k=1}^{K} \in \mathbb{R}^d$
  - VQ-GAN (Esser et al., CVPR'21) is adopted in implementation

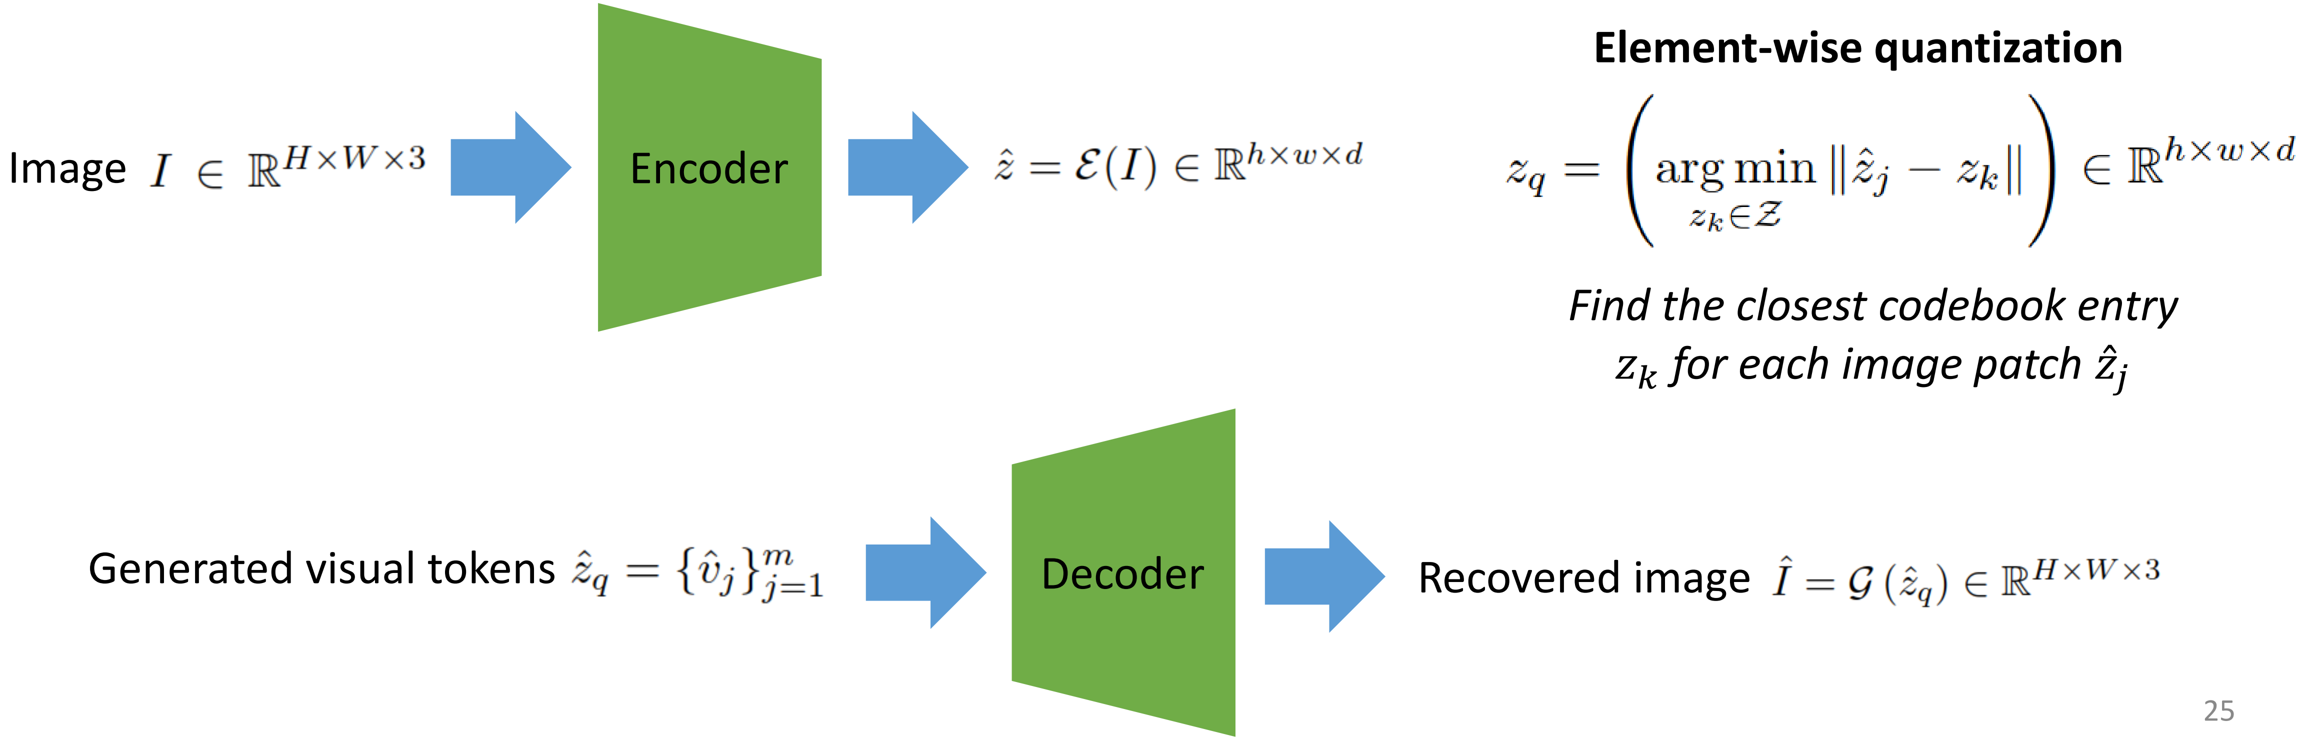

Visual Explanation

- "A picture is worth a thousand words."
  - Look at the pictures when choosing a restaurant
  - See the room layout when booking a hotel
  - Go to the framework figure when reading a paper
  - ……

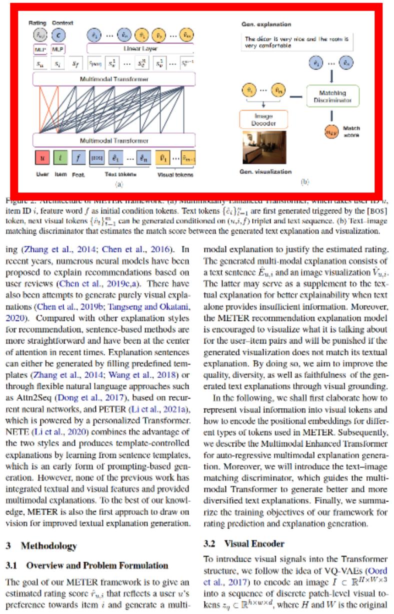

METER
(Geng et al., ACL'22)

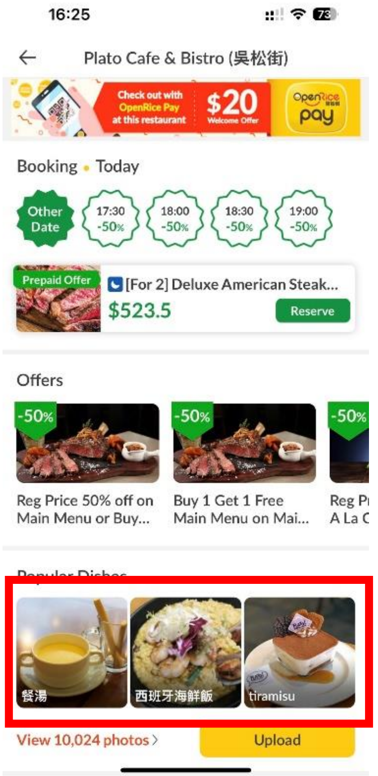

OpenRice
(openrice.com)

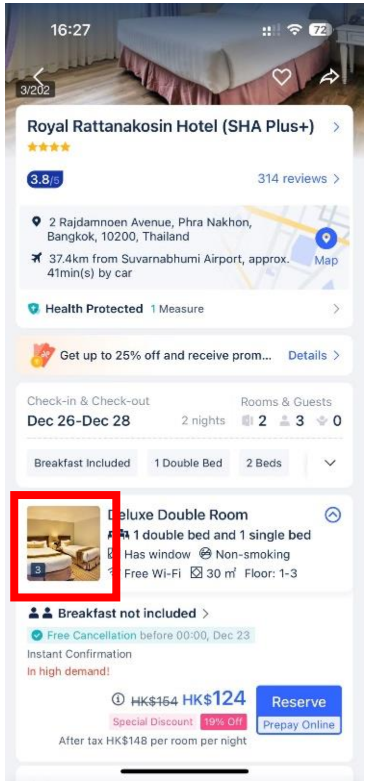

Trip.com
(trip.com)

Image Generation with Transformer

- Perform autoregressive generation as natural language generation
  - First text, and then image

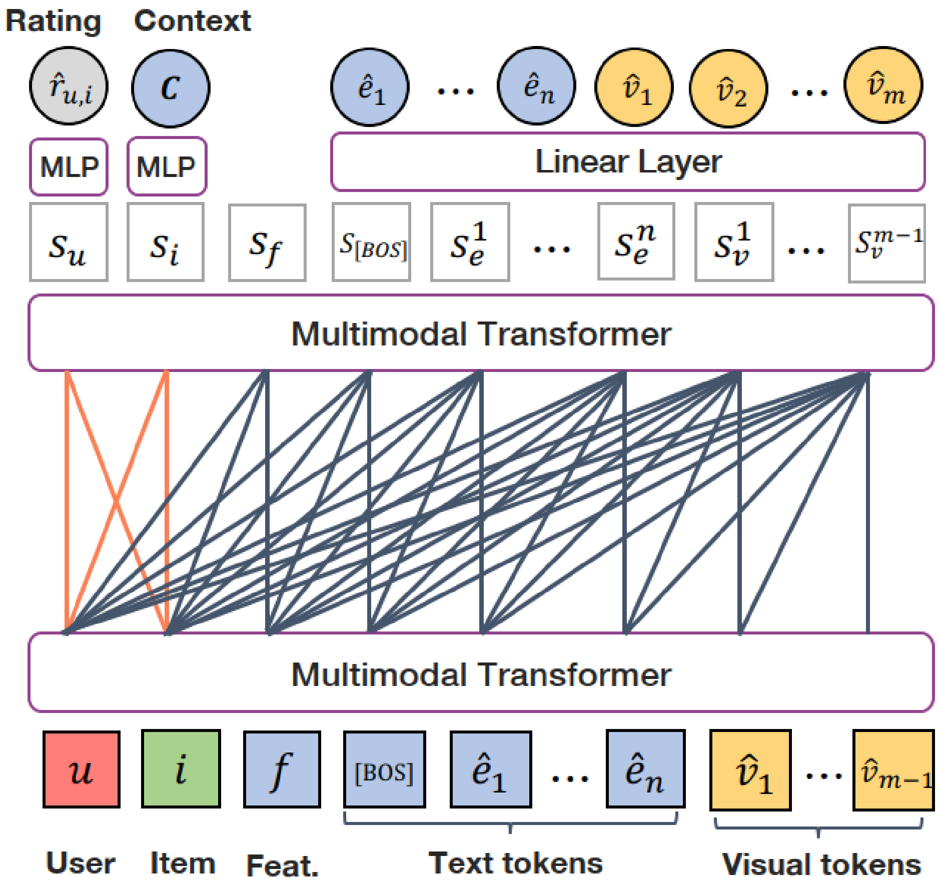

Text-Image Matching

- Design a discriminator to measure the degree of consistency between text and image
  - Two Transformer encoders for visual token sequence and text sequence

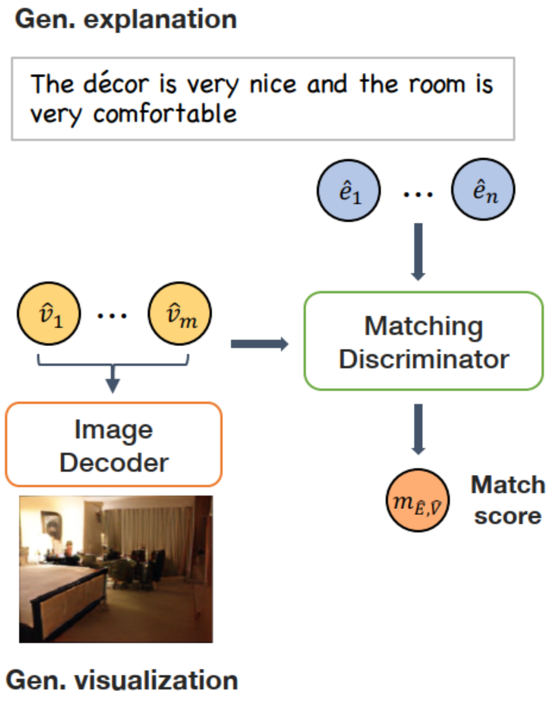

$$\mathcal{L}_{\mathrm{d}} = \mathbb{E}\left[\log\left(D(\mathbf{E},\mathbf{V})\right)\right] + \mathbb{E}\left[\log\left(1 - D\left(\mathbf{E},\hat{\mathbf{V}}\right)\right)\right] + \mathbb{E}\left[\log\left(1 - D\left(\hat{\mathbf{E}},\mathbf{V}\right)\right)\right]$$

Datasets (Li et al., CIKM’20)

- Yelp
  - Restaurant
- TripAdvisor
  - Hotel
- The explanation is a review sentence containing at least one feature
- No images

Dataset	Yelp	TripAdvisor
#users	27,147	9,765
#items	20,266	6,280
#explanations	1,293,247	320,023
#features	7,340	5,069
#images	649,370	331,540

Excellent movie

Reviewed in the United States on January 11, 2012

Verified Purchase

I love Brad Pitt and always watch his movies, and I'm rarely disappointed, and wasn't this time. Moneyball is a great movie based on a true story, you don't have to be into baseball to get the movie but it does help if you know a little. The story was really good and gives the viewer an understanding of how hard it is for small budget major league baseball teams to compete with teams like the Yankees and others where money is no object. I really enjoyed this film, great story and acting.

Text-Image Association

Text Expl.: we also had huevos rancheros and cheese grits from room Service one morning which was great
Assigned image visualization:

Text Expl.: The executive floor was well stocked and snacks where great
Assigned image visualization:

- Use Sentence-BERT (Reimers and Gurevych, EMNLP’19) to compute the embeddings of explanation sentences
- Cluster sentence semantics into different groups representing similar concepts and topics
- Query relevant images through Google Images API with sentences at cluster centers
- Assign each textual explanation the most suitable image with CLIP (Radford et al., ICML’21)

Evaluation Metrics

- Text quality: not equal to explainability (Chen et al., SIGIR'19 Workshop; Li et al., CIKM'20)
  - BLEU (Papineni et al., ACL'02)
  - ROUGE (Lin, ACL'04 Workshop)
- Explainability from the angle of item features (Li et al., CIKM'20)
  - Unique Sentence Ratio (USR)
  - Feature Matching Ratio (FMR)
  - Feature Coverage Ratio (FCR)
  - Feature Diversity (DIV)
- Image consistency
  - CLIPScore (Hessel et al., EMNLP'21)

Quantitative Analysis on Explanations

- Best or comparable performance

Methods	Text Explainability		Text Diversity		Text Quality				Image Consistency
	FMR↑	FCR↑	DIV↓	USR↑	BLUE-1↑	BLUE-4↑	ROUGE-1↑	ROUGE-2↑	CLIPScore↑
TripAdvisor									
Att2Seq	0.06	0.15	4.32	0.17	15.27	1.03	15.92	2.09	-
Transformer	0.04	0.00	10.00	0.00	12.79	0.71	15.88	2.34	-
NETE	0.78	0.27	2.22	**0.57**	22.39	3.66	27.71	7.66	-
PETER	0.89	0.35	1.61	0.25	24.32	4.55	30.49	9.24	-
METER	**0.90**	**0.39**	**1.42**	<u>0.56</u>	**24.57**	**4.76**	**30.77**	**9.41**	0.62
Yelp									
Att2Seq	0.07	0.12	2.41	0.13	10.29	0.58	13.29	1.31	-
Transformer	0.06	0.06	2.46	0.01	7.39	0.42	12.56	1.09	-
NETE	0.80	0.27	1.48	**0.52**	19.31	2.69	25.56	6.63	-
PETER	0.86	0.38	1.08	0.34	20.80	3.43	27.95	7.94	-
METER	**0.88**	**0.43**	**1.02**	<u>0.42</u>	**21.30**	**3.61**	**28.32**	**8.09**	0.59

Qualitative Case Study on Explanations

- High-quality images aligning with the textual explanations

Feat. word: beds
Pred. rating: 4.51 *GT rating:* 4
Gen. Expl.: the beds are very comfortable and the room is spacious
Ref. Expl.: the beds were super comfortable
CLIPScore: 0.71

Feat. word: pool
Pred. rating: 4.48 *GT rating:* 4
Gen. Expl.: the swimming pool and outside area was really nice
Ref. Expl.: the pool is lovely and the service was excellent
CLIPScore: 0.69

Feat. word: reception
Pred. rating: 4.82 *GT rating:* 5
Gen. Expl.: the reception staff were very helpful and friendly
Ref. Expl.: check in was very professional and fast and same receptionist the next 3 days remembered
CLIPScore: 0.62

Feat. word: floor
Pred. rating: 3.88 *GT rating:* 4
Gen. Expl.: i was on the top floor and had a view of the city
Ref. Expl.: breakfast on the 3rd floor was good also
CLIPScore: 0.53

(a)

Feat. word: sushi
Pred. rating: 4.32 *GT rating:* 5
Gen. Expl.: we ate in the japanese restaurant one night and the food was excellent
Ref. Expl.: we enjoyed the sushi bar and the steakhouse
CLIPScore: 0.67

Feat. word: juice
Pred. rating: 3.59 *GT rating:* 4
Gen. Expl.: i ordered a juice and it was delicious
Ref. Expl.: was pure lemon juice
CLIPScore: 0.55

Feat. word: bar
Pred. rating: 4.48 *GT rating:* 4
Gen. Expl.: the bar is a great place to unwind after a long day of sightseeing
Ref. Expl.: great bar and excellent food
CLIPScore: 0.65

Feat. word: buffet
Pred. rating: 4.51 *GT rating:* 4
Gen. Expl.: the breakfast buffet was excellent with a wide selection of hot and cold items
Ref. Expl.: the breakfast buffet is busy on the weekend with a good selection
CLIPScore: 0.64

(b)

User Study

- Randomly pick 500 samples for each method
- Invite 30 participants to give 5-likert ratings from four aspects
- Baselines
  - PETER for sentence
  - METER without VQ-GAN for image

	Sentence		**Image**	
	Faithfulness	Diversity	Consistency	Quality
Baselines	3.41	2.96	2.54	3.04
Ours	**4.57**	**3.70**	**3.06**	**4.19**

Limitations

- Not every concept can be expressed with images
  - E.g., “the service is good”
- The quality of generated images is good, but not perfect
- Users may not tolerate with flawed product images, in contrast to those in art or design

Conclusion

- Propose the first approach to jointly generate textual explanations and corresponding image visualizations
- Introduce a text-image matching discriminator to encourage sentences with fine-grained and diverse concepts
- Demonstrate with experiments that the model can provide diverse and faithful text explanations, together with image visualizations

Outline

- Explainable Recommendation
- Natural Language Explanation (ACL'21)
- Visual Explanation (ACL'22)
- **Future Work**

Bias in Generated Explanations

- When a new game is recommended, it is unfair to generate short and generic explanations for female users, e.g., “Good game”.

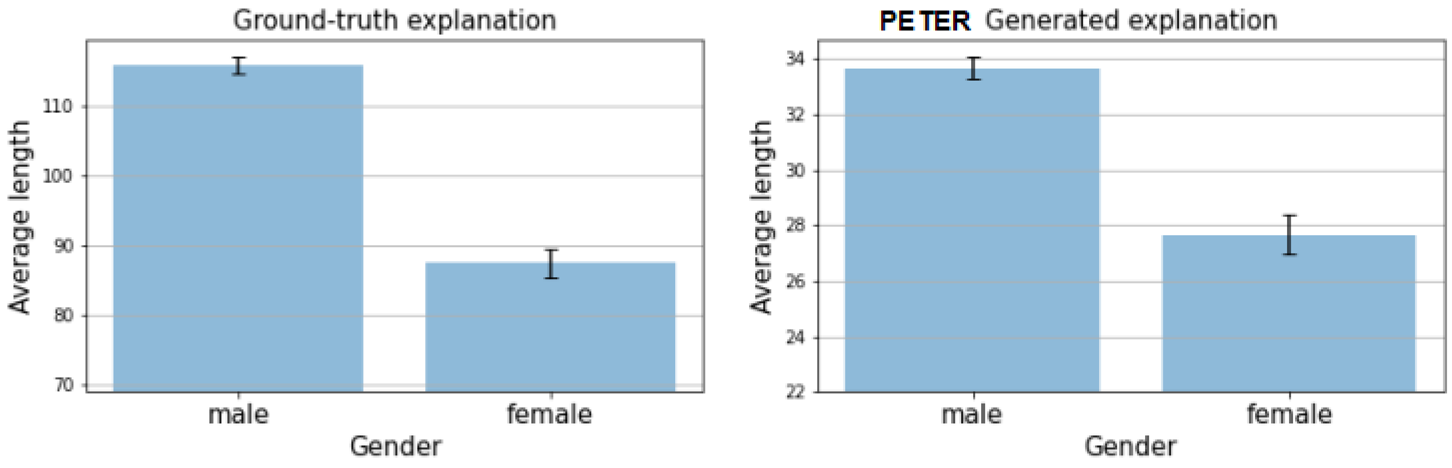

Courtesy image from COFFEE (Wang et al., arXiv’22)

Sentiment of Generated Explanations

- There are various types of sentiment in generated explanations.
- Would it be different for users with different demographic attributes?

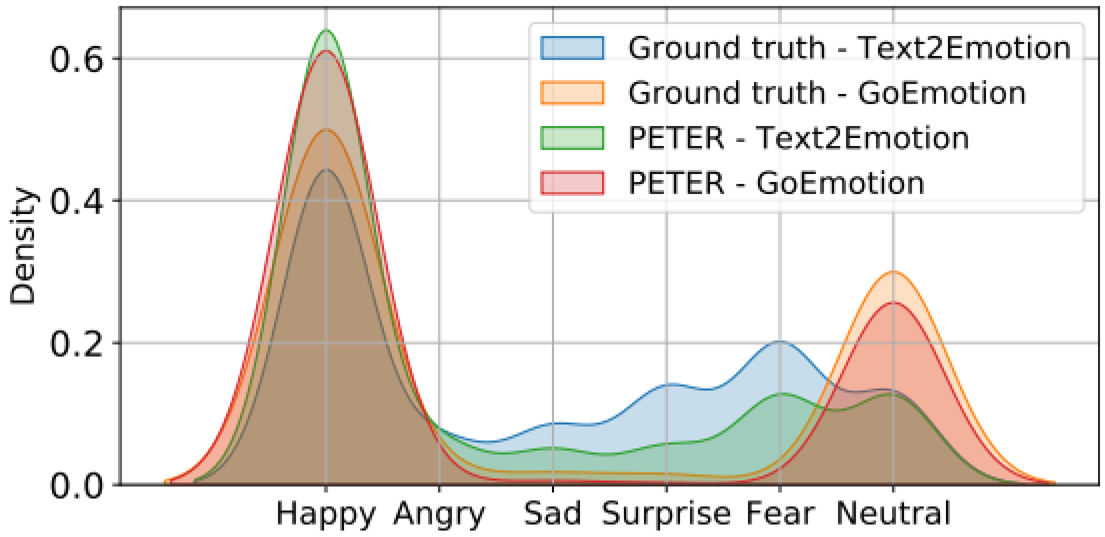

Courtesy image from EmoTER (Wen et al., arXiv'22)

Gender Bias in Pre-trained Models

- Gender bias in GPT-2 (Liang et al., ICML'21)

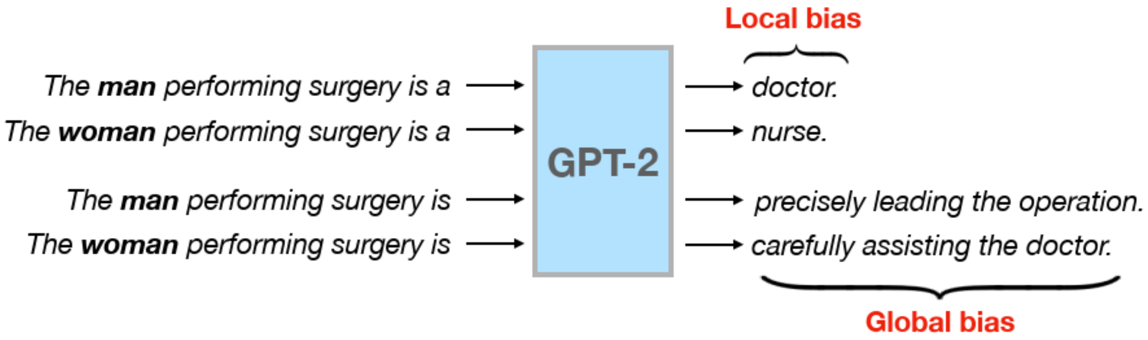

- Does the bias exist, when the models are adapted to recommendation explanation generation?
  - GPT-2 (Li et al., arXiv'22)
  - T5 (Geng et al., RecSys'22)
  - BERT (Ni et al., EMNLP'19)

References (1)

- Tintarev, Nava, and Judith Masthoff. "Explaining recommendations: Design and evaluation." Recommender systems handbook. 2015.
- He, Xiangnan, et al. "Trirank: Review-aware explainable recommendation by modeling aspects." CIKM’15.
- Wang, Xiang, et al. "Tem: Tree-enhanced embedding model for explainable recommendation." WWW’18.
- Zhang, Yongfeng, et al. "Explicit factor models for explainable recommendation based on phrase-level sentiment analysis." SIGIR’14.
- Wang, Nan, et al. "Explainable Recommendation via Multi-Task Learning in Opinionated Text Data." SIGIR’18.
- Tan, Juntao, et al. "Counterfactual explainable recommendation." CIKM’21.
- Gao, Jingyue, et al. "Explainable Recommendation Through Attentive Multi-View Learning." AAAI’19.

References (2)

- Wang, Xiang, et al. "Explainable Reasoning over Knowledge Graphs for Recommendation." AAAI’19.
- Seo, Sungyong, et al. "Interpretable convolutional neural networks with dual local and global attention for review rating prediction." RecSys’17.
- Lu, Yichao, et al. "Coevolutionary Recommendation Model: Mutual Learning between Ratings and Reviews." WWW’18.
- Chen, Xu, et al. "Personalized fashion recommendation with visual explanations based on multimodal attention network: Towards visually explainable recommendation." SIGIR’19.
- Chen, Chong, et al. "Neural attentional rating regression with review-level explanations." WWW’18.
- Catherine, Rose, and William Cohen. "Transnets: Learning to transform for recommendation." RecSys’17.
- Wang, Xiting, et al. "A Reinforcement Learning Framework for Explainable Recommendation." ICDM’18.

References (3)

- Chen, Xu, et al. "Dynamic Explainable Recommendation based on Neural Attentive Models." AAAI’19.
- Lu, Yichao, et al. "Why I like it: multi-task learning for recommendation and explanation." RecSys’18.
- Li, Lei, et al. "Generate neural template explanations for recommendation." CIKM’20.
- Li, Lei, et al. "Personalized transformer for explainable recommendation." ACL’21.
- Li, Lei, et al. "Personalized prompt learning for explainable recommendation." arXiv’22.
- Vaswani, Ashish, et al. "Attention is all you need." NIPS’17.
- Chen, Zhongxia, et al. "Towards Explainable Conversational Recommendation." IJCAI’20.
- Oord, Van Den, et al. "Neural discrete representation learning." NIPS’17.
- Esser, Patrick, et al. "Taming transformers for high-resolution image synthesis." CVPR’21.

References (4)

- Geng, Shijie, et al. "Improving Personalized Explanation Generation through Visualization." ACL’22.
- Reimers, Nils, and Iryna Gurevych. "Sentence-bert: Sentence embeddings using siamese bert-networks." EMNLP’19.
- Radford, Alec, et al. "Learning transferable visual models from natural language supervision." ICML’21.
- Chen, Hanxiong, et al. "Generate natural language explanations for recommendation." SIGIR’19 Workshop.
- Papineni, Kishore, et al. "BLEU: a method for automatic evaluation of machine translation." ACL’02.
- Lin, Chin-Yew. "ROUGE: A Package for Automatic Evaluation of Summaries." ACL’04 Workshop.
- Hessel, Jack, et al. "Clipscore: A reference-free evaluation metric for image captioning." EMNLP’21.

References (5)

- Wang, Nan, et al. "COFFEE: Counterfactual Fairness for Personalized Text Generation in Explainable Recommendation." arXiv'22.
- Wen, Bingbing, et al. "Towards Generating Robust, Fair, and Emotion-Aware Explanations for Recommender Systems." arXiv'22.
- Liang, Paul Pu, et al. "Towards understanding and mitigating social biases in language models." ICML'21.
- Geng, Shijie, et al. "Recommendation as Language Processing (RLP): A Unified Pretrain, Personalized Prompt & Predict Paradigm (P5)." RecSys'22.
- Ni, Jianmo, et al. "Justifying recommendations using distantly-labeled reviews and fine-grained aspects." EMNLP'19.

Q&A

Thank you!

csleili@comp.hkbu.edu.hk

lileipisces.github.io