Affective Computing and Bandits: Capturing Context in Cold Start Situations

Sebastian Oehme and Linus W. Dietz
Technical University of Munich

Joint Workshop on Interfaces and Human Decision Making for Recommender Systems
ACM Conference on Recommender Systems 2018

GRUNDY (1979)

EARLY DAYS OF COMPUTING

*[...] a computer **cannot help** but be at a loss compared to a human in **quickly sizing up a person** on the basis of superficial characteristics, if for no other reason than that it can **neither see him** (to **determine his age, type of clothing, or sex**) nor hear him […]*

E. Rich, “User modeling via stereotypes,” Cognitive science, vol. 3, no. 4, pp. 329–354, 1979.

COMPARISION

GRUNDY AND NEW TECHNOLOGIES

Grundy

can not see or hear to determine:

- age
- type clothing
- gender
- origine
- self-assurance

Today

- **Powerful chip technology** (Moore's law)
- **Computer Vision** (YOLO, Inception V4, etc.)
- **Vast amount of data** (IBM estimates 2.5 quintillion bytes of every day)

→ Technology has evolved!

HYPOTHESIS

IMPROVING RECOMMENDER SYSTEMS

Contextual recommender systems that incorporate facial classification can outperform traditional systems in cold start situations

AFFECTIVE RECOMMENDER SYSTEMS

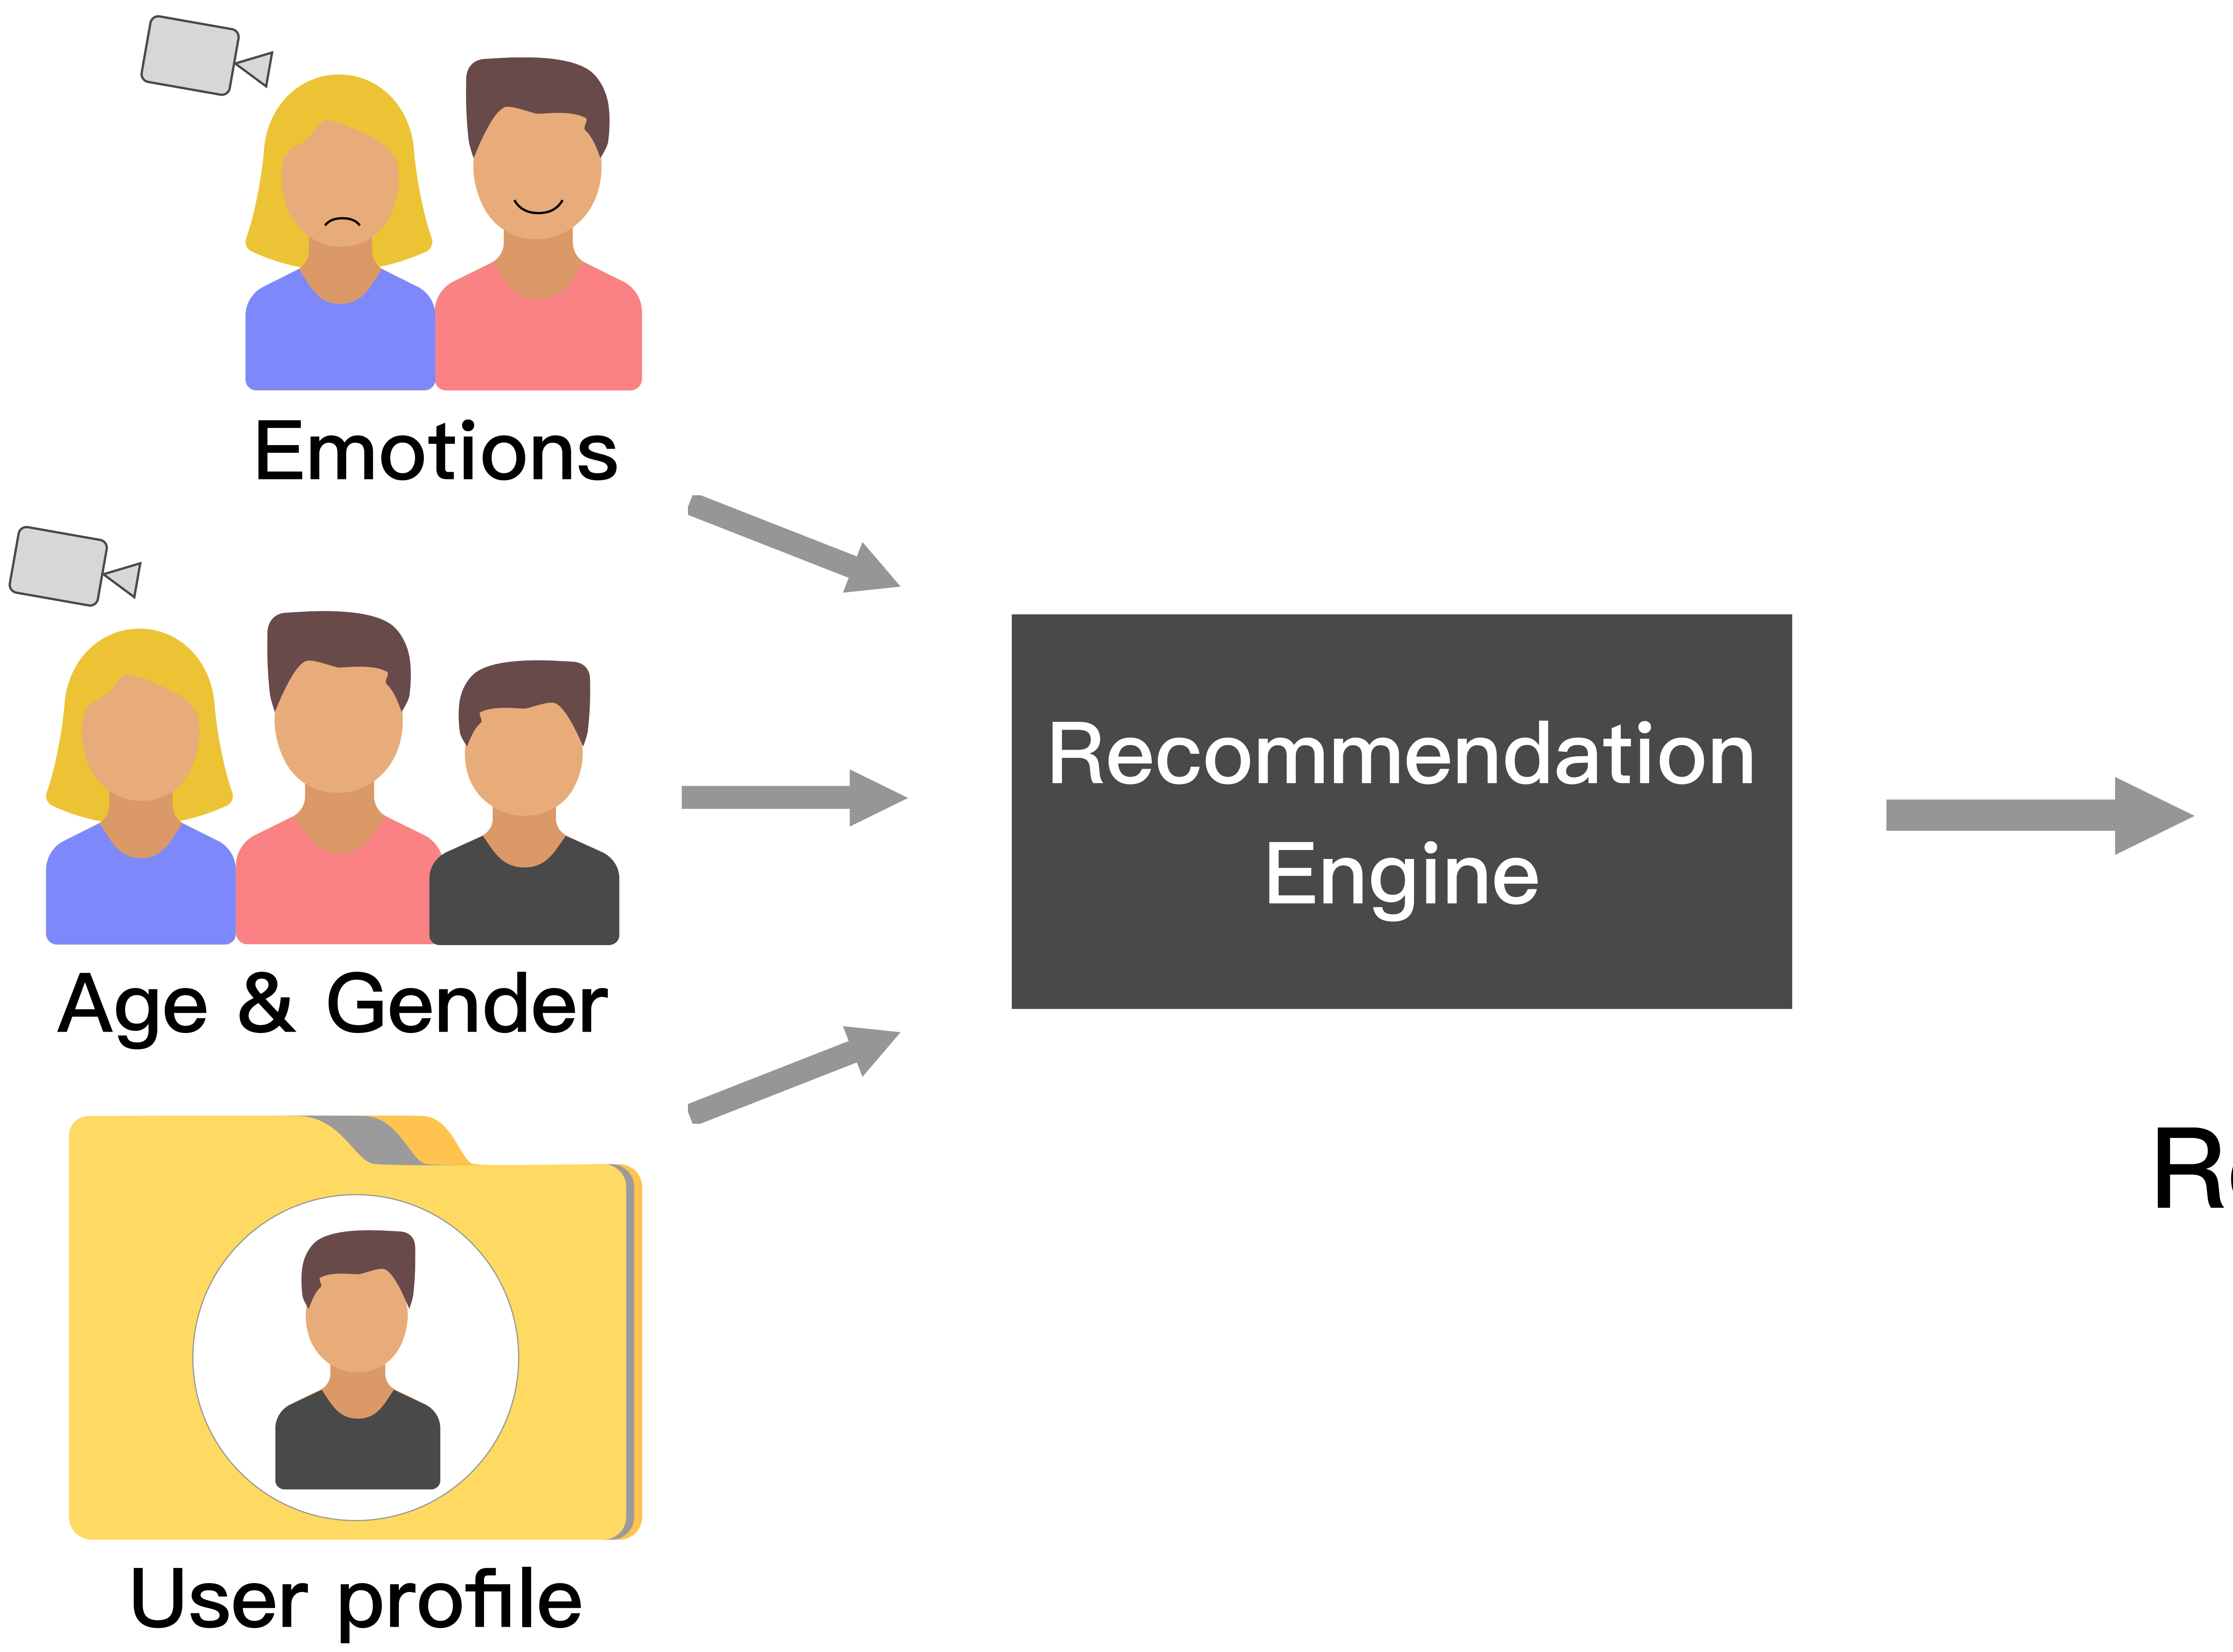

item	score
1	0,1
2	0,7
3	1
4	0,4
...	...

Recommendations

BANDIT STRATEGIES

OVERVIEW

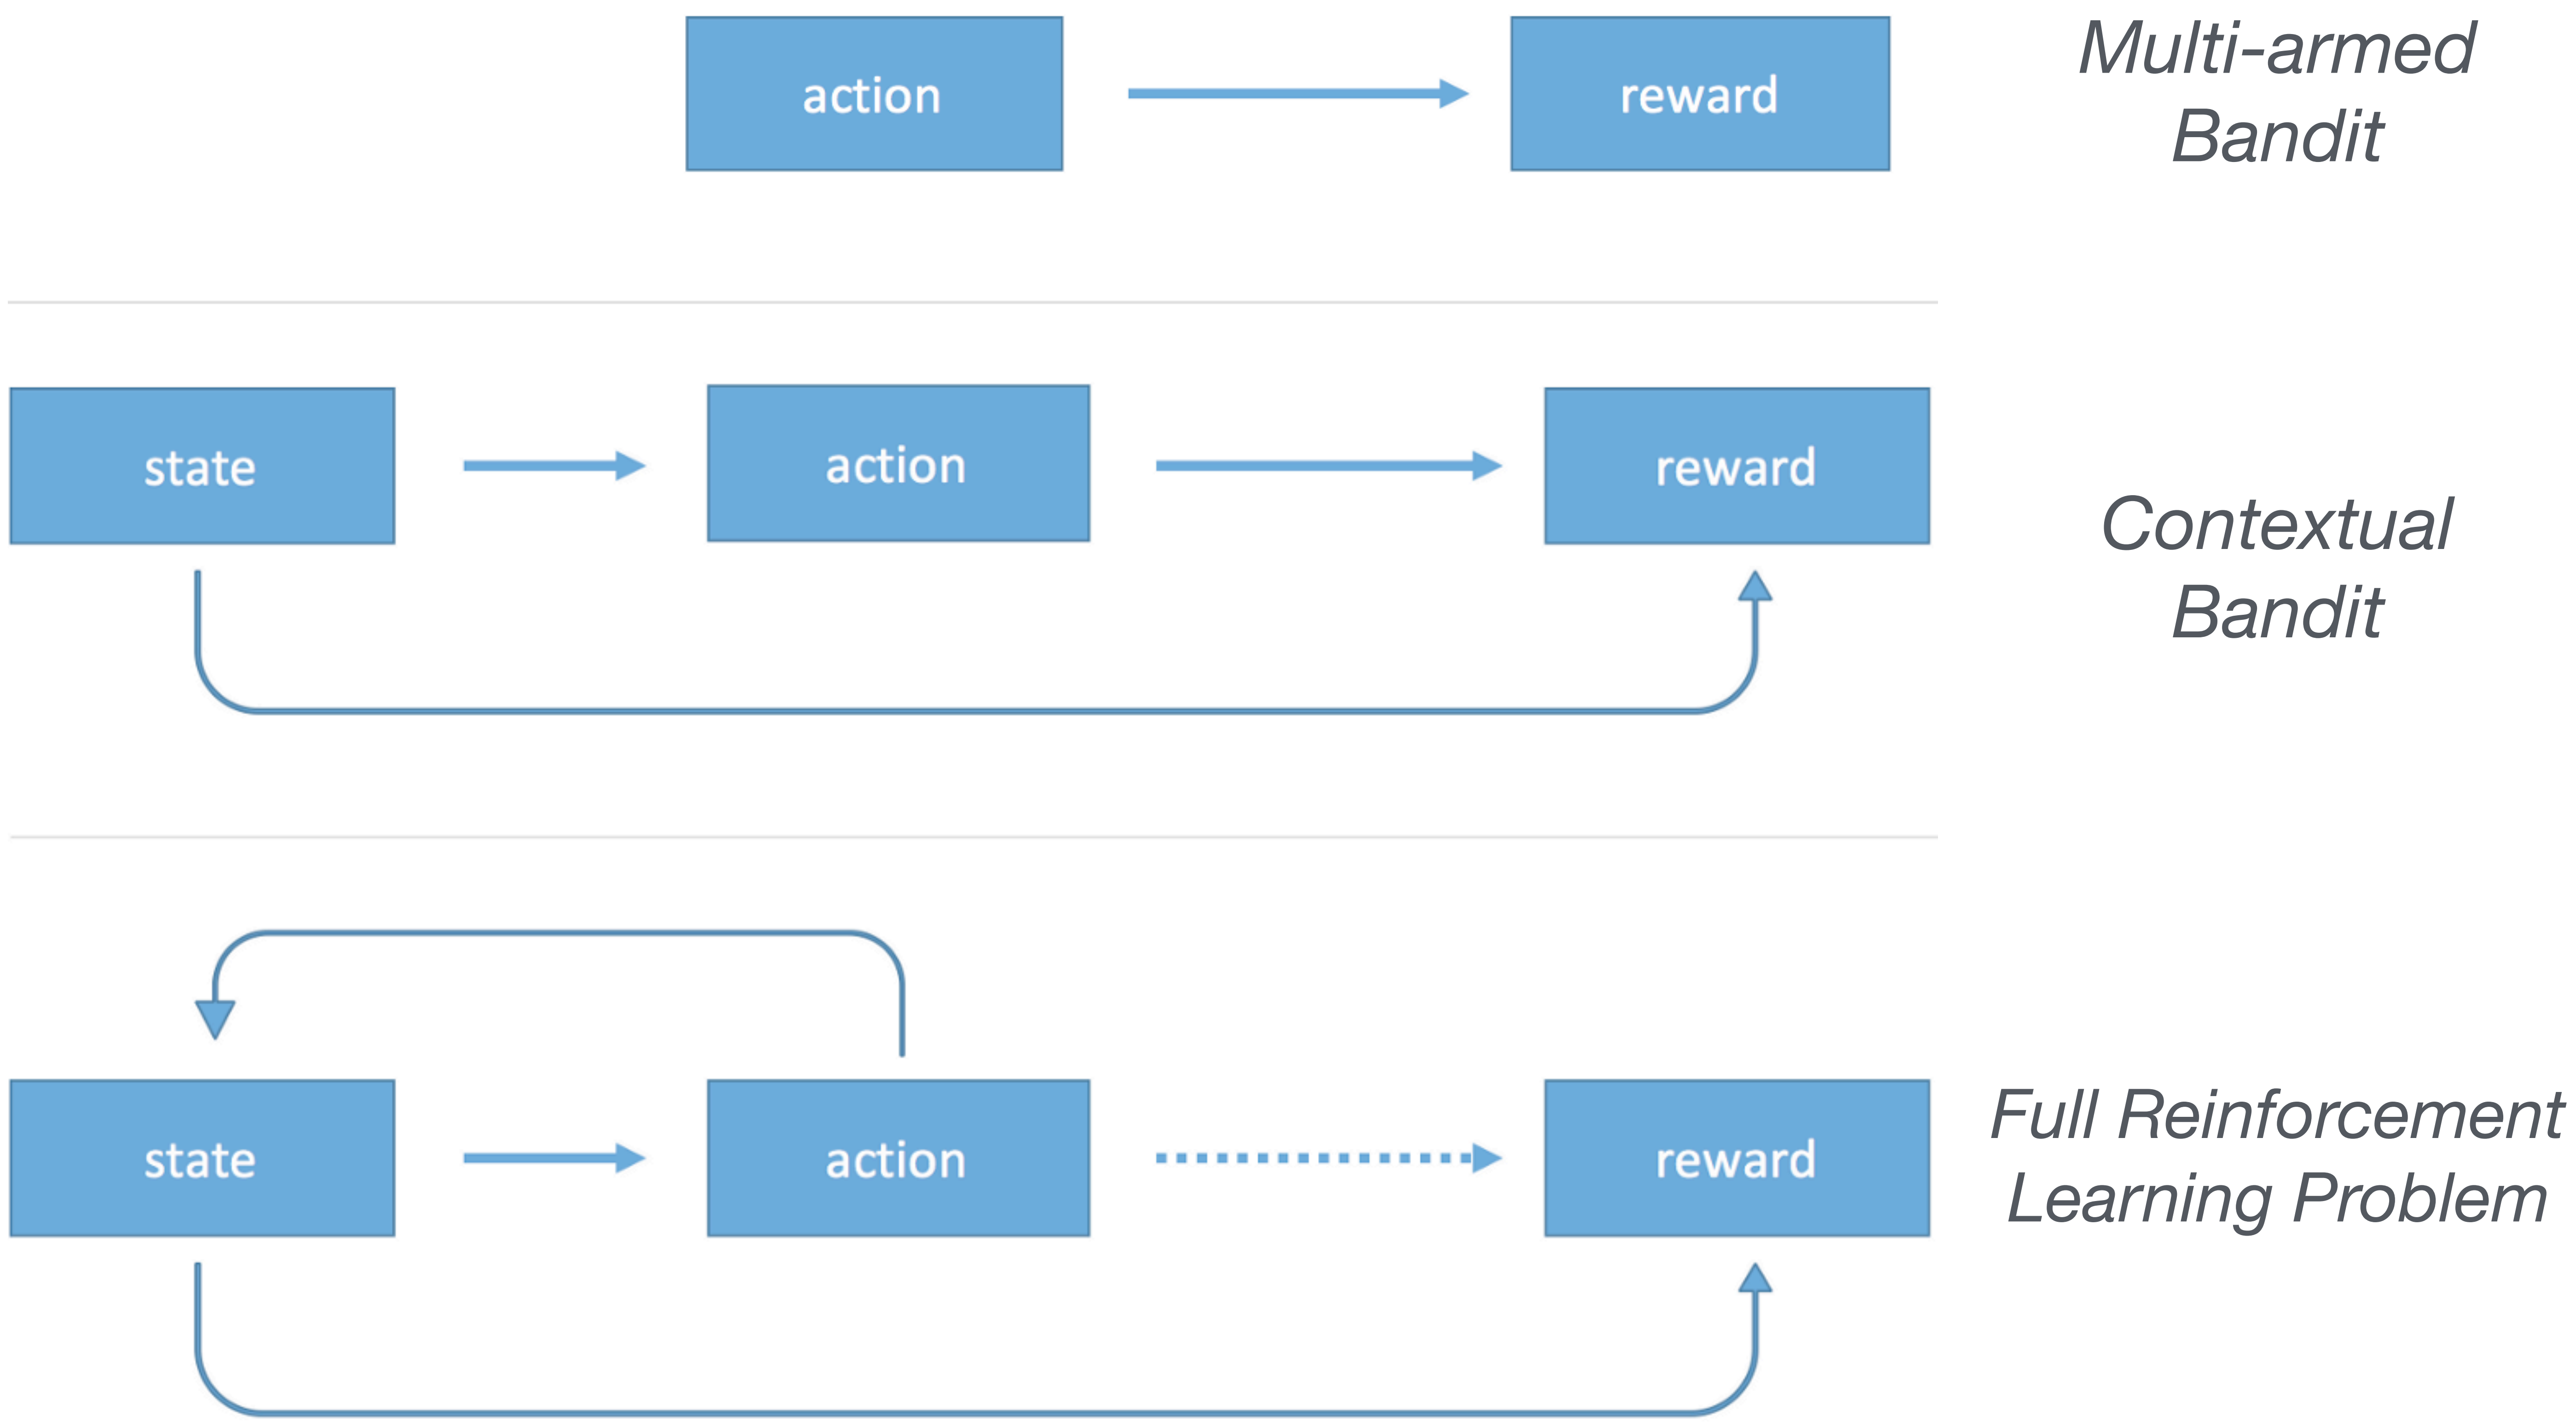

BANDIT STRATEGIES

MULTI-ARMED BANDIT (BASELINE)

Multi-armed Bandit

action → reward

Example:
ε-Greedy strategy

q > ε

Exploitation:
Content with max. expected reward

q ≤ ε

Exploration:
Random content

Update reward

BANDIT STRATEGIES

CONTEXTUAL BANDIT

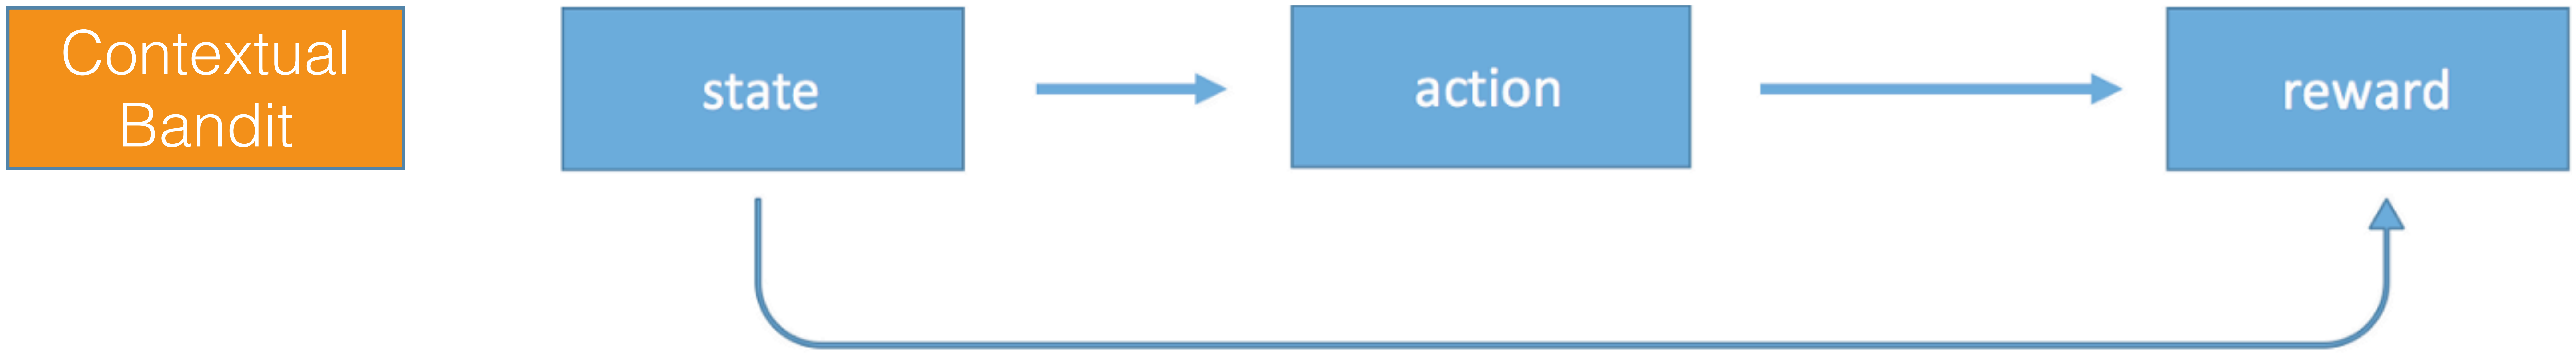

Contextual-ε-Greedy

D. Bouneffouf et al. 2012

ε dependent on context

$q > \varepsilon$ **Exploitation:** Content with max. expected reward with similar context

$q \leq \varepsilon$ **Exploration:** Random content

Update reward

MODEL

COMPUTER VISION-BASED CONTEXTUAL BANDIT

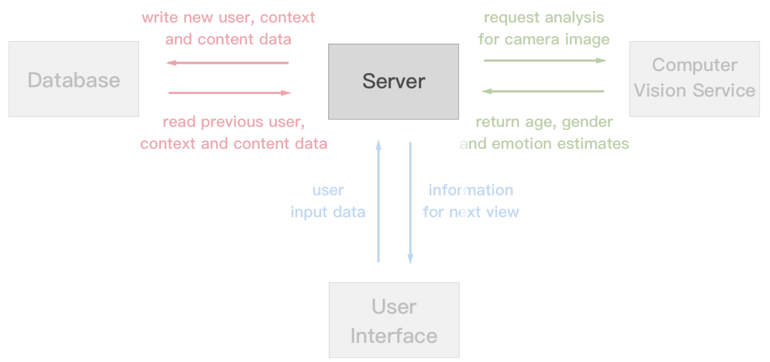

Exploitation confidence

$$\varepsilon = 1 - \operatorname{argmax}(sim(U^t, U^c))$$

User Similarity

Weights

$$sim(U^t, U^c) = \alpha \cdot sim(a^t, a^c) + \beta \cdot sim(g^t, g^c) + \gamma \cdot EF$$

Age

Gender

Emotional Feedback

$$EF = \frac{\sum_k sim_k(f_k^t, f_k^c) \cdot \left(1 + sim_k(e_k^t, e_k^c)\right)}{2i}$$

EXPERIMENT

SETUP AND EVALUATION

Comparing two recommendation strategies

Independent variable

- Bandit strategy ∈ {Normal,Contextual}

Dependent variables

- User’s ratings (positive or negative)
- Measurement of facial emotions (Microsoft Face)
- Feedback from the questionnaire

Details

Dataset

- 3000 filtered memes
- gathered from 9gag.com
- 24.01.2018 - 9.02.2018.

Users

- 21 participants
- Age: 19 years to 31 years
- 10 male and 11 female
- Rated 60 images per strategy

PROTOTYPE

HIGH LEVEL DESIGN

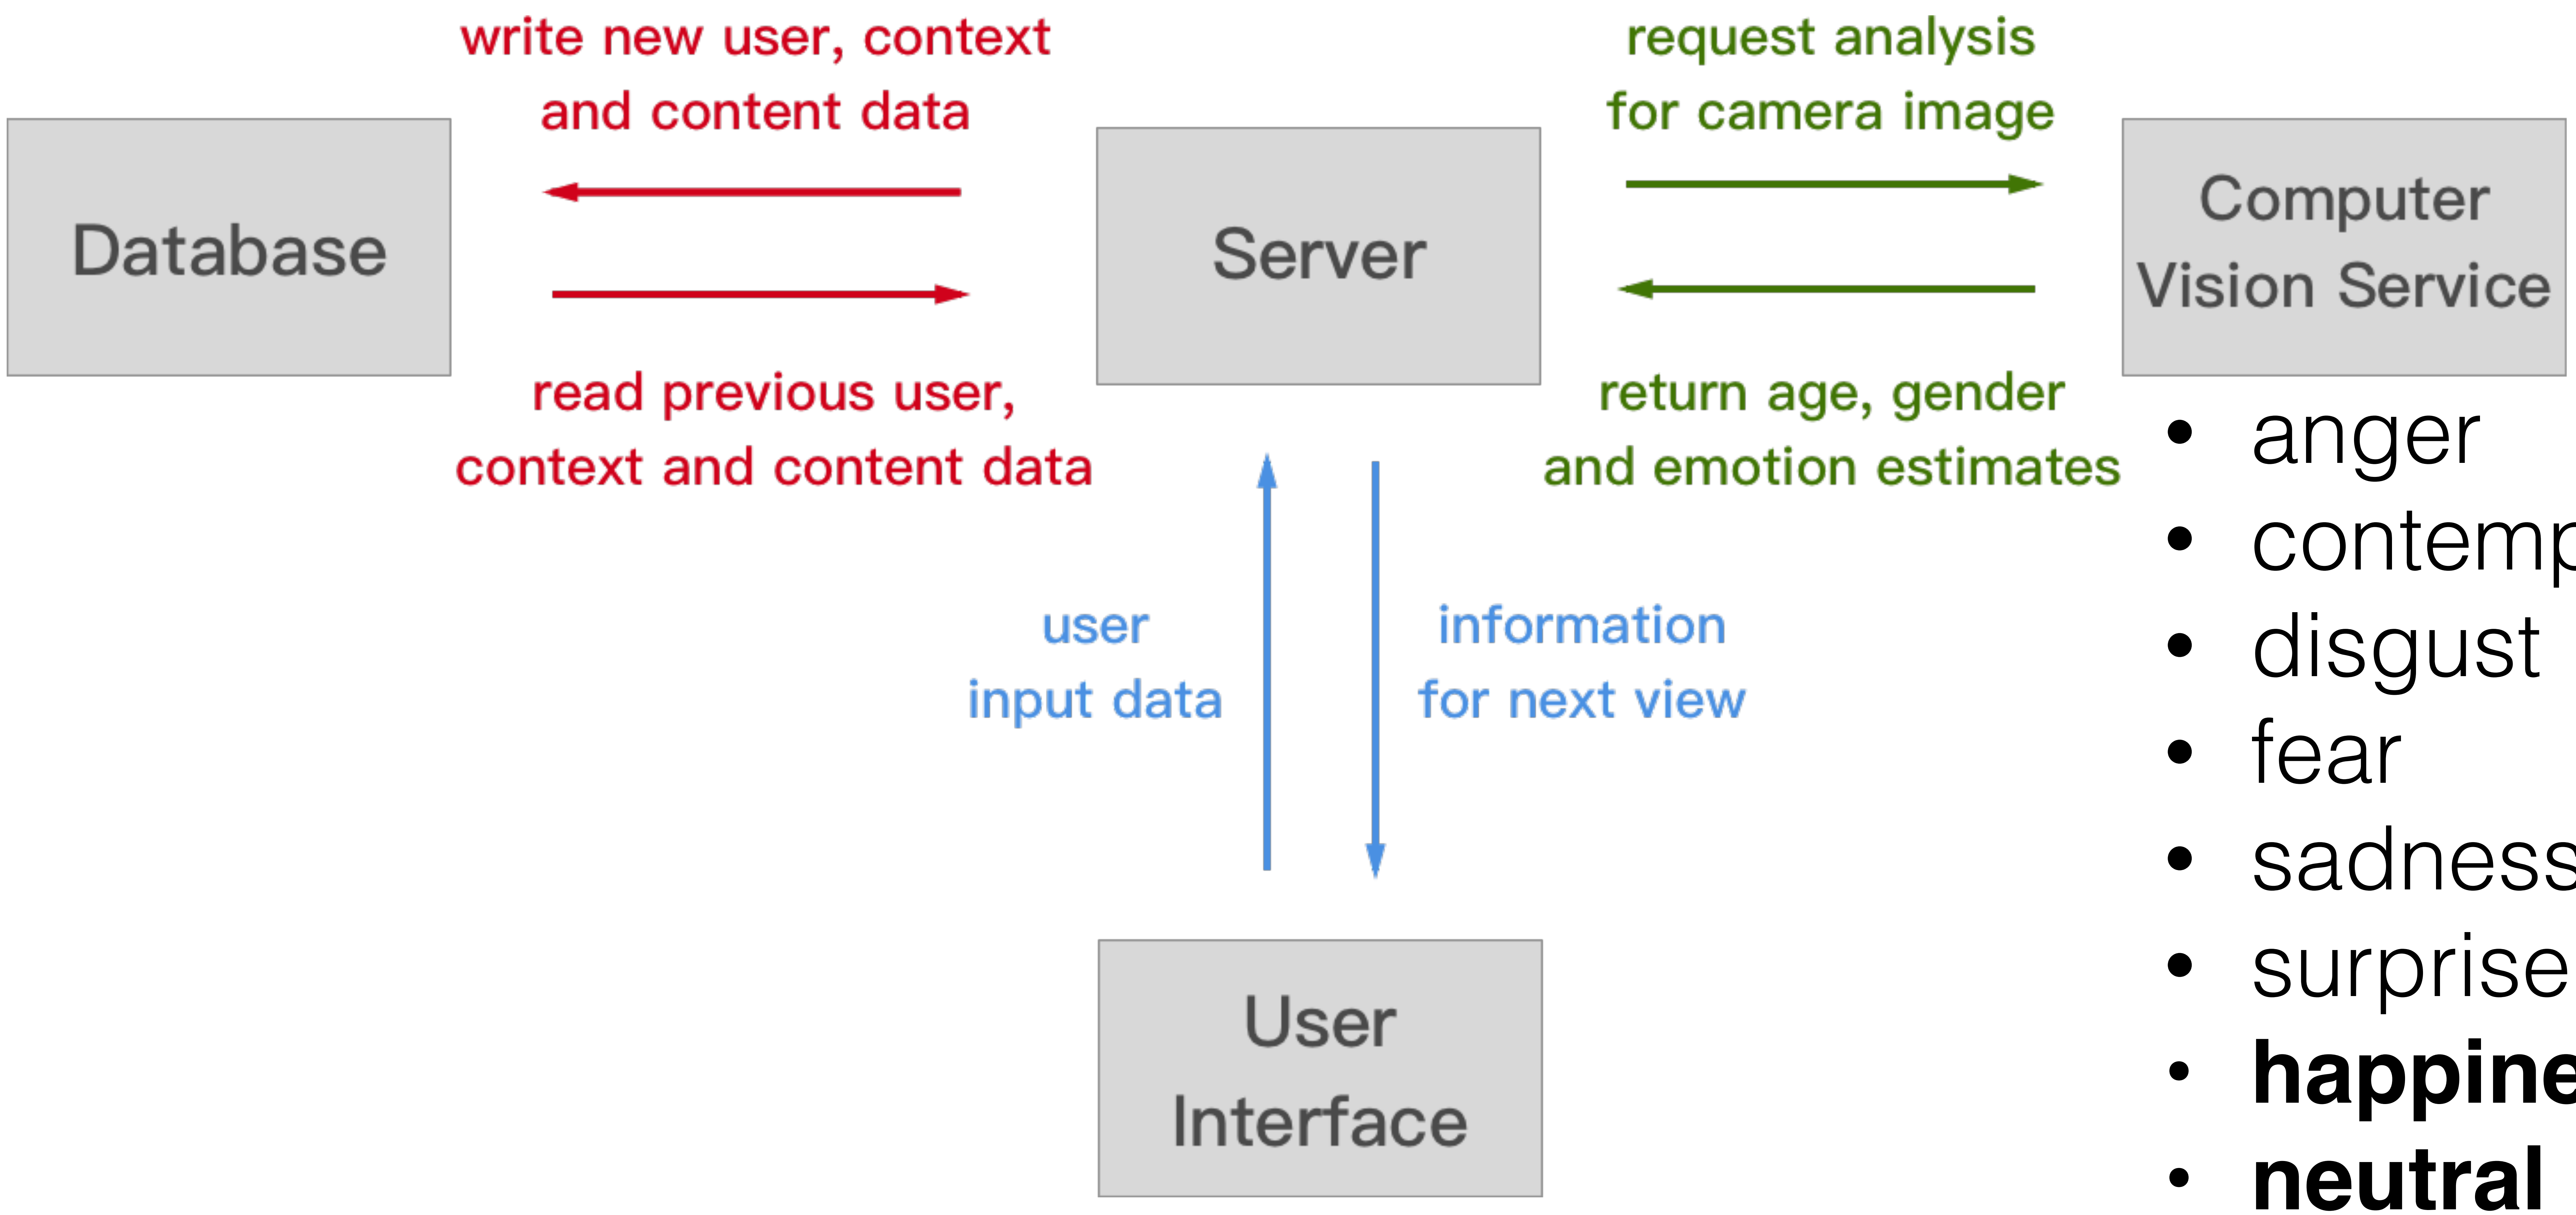

- anger
- contempt
- disgust
- fear
- sadness
- surprise
- **happiness**
- **neutral**

Experiment

USER INTERFACE

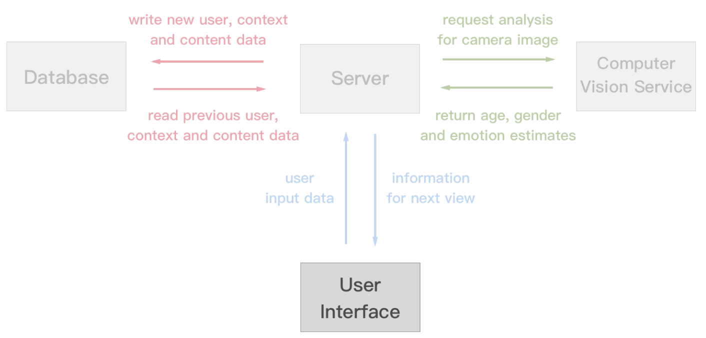

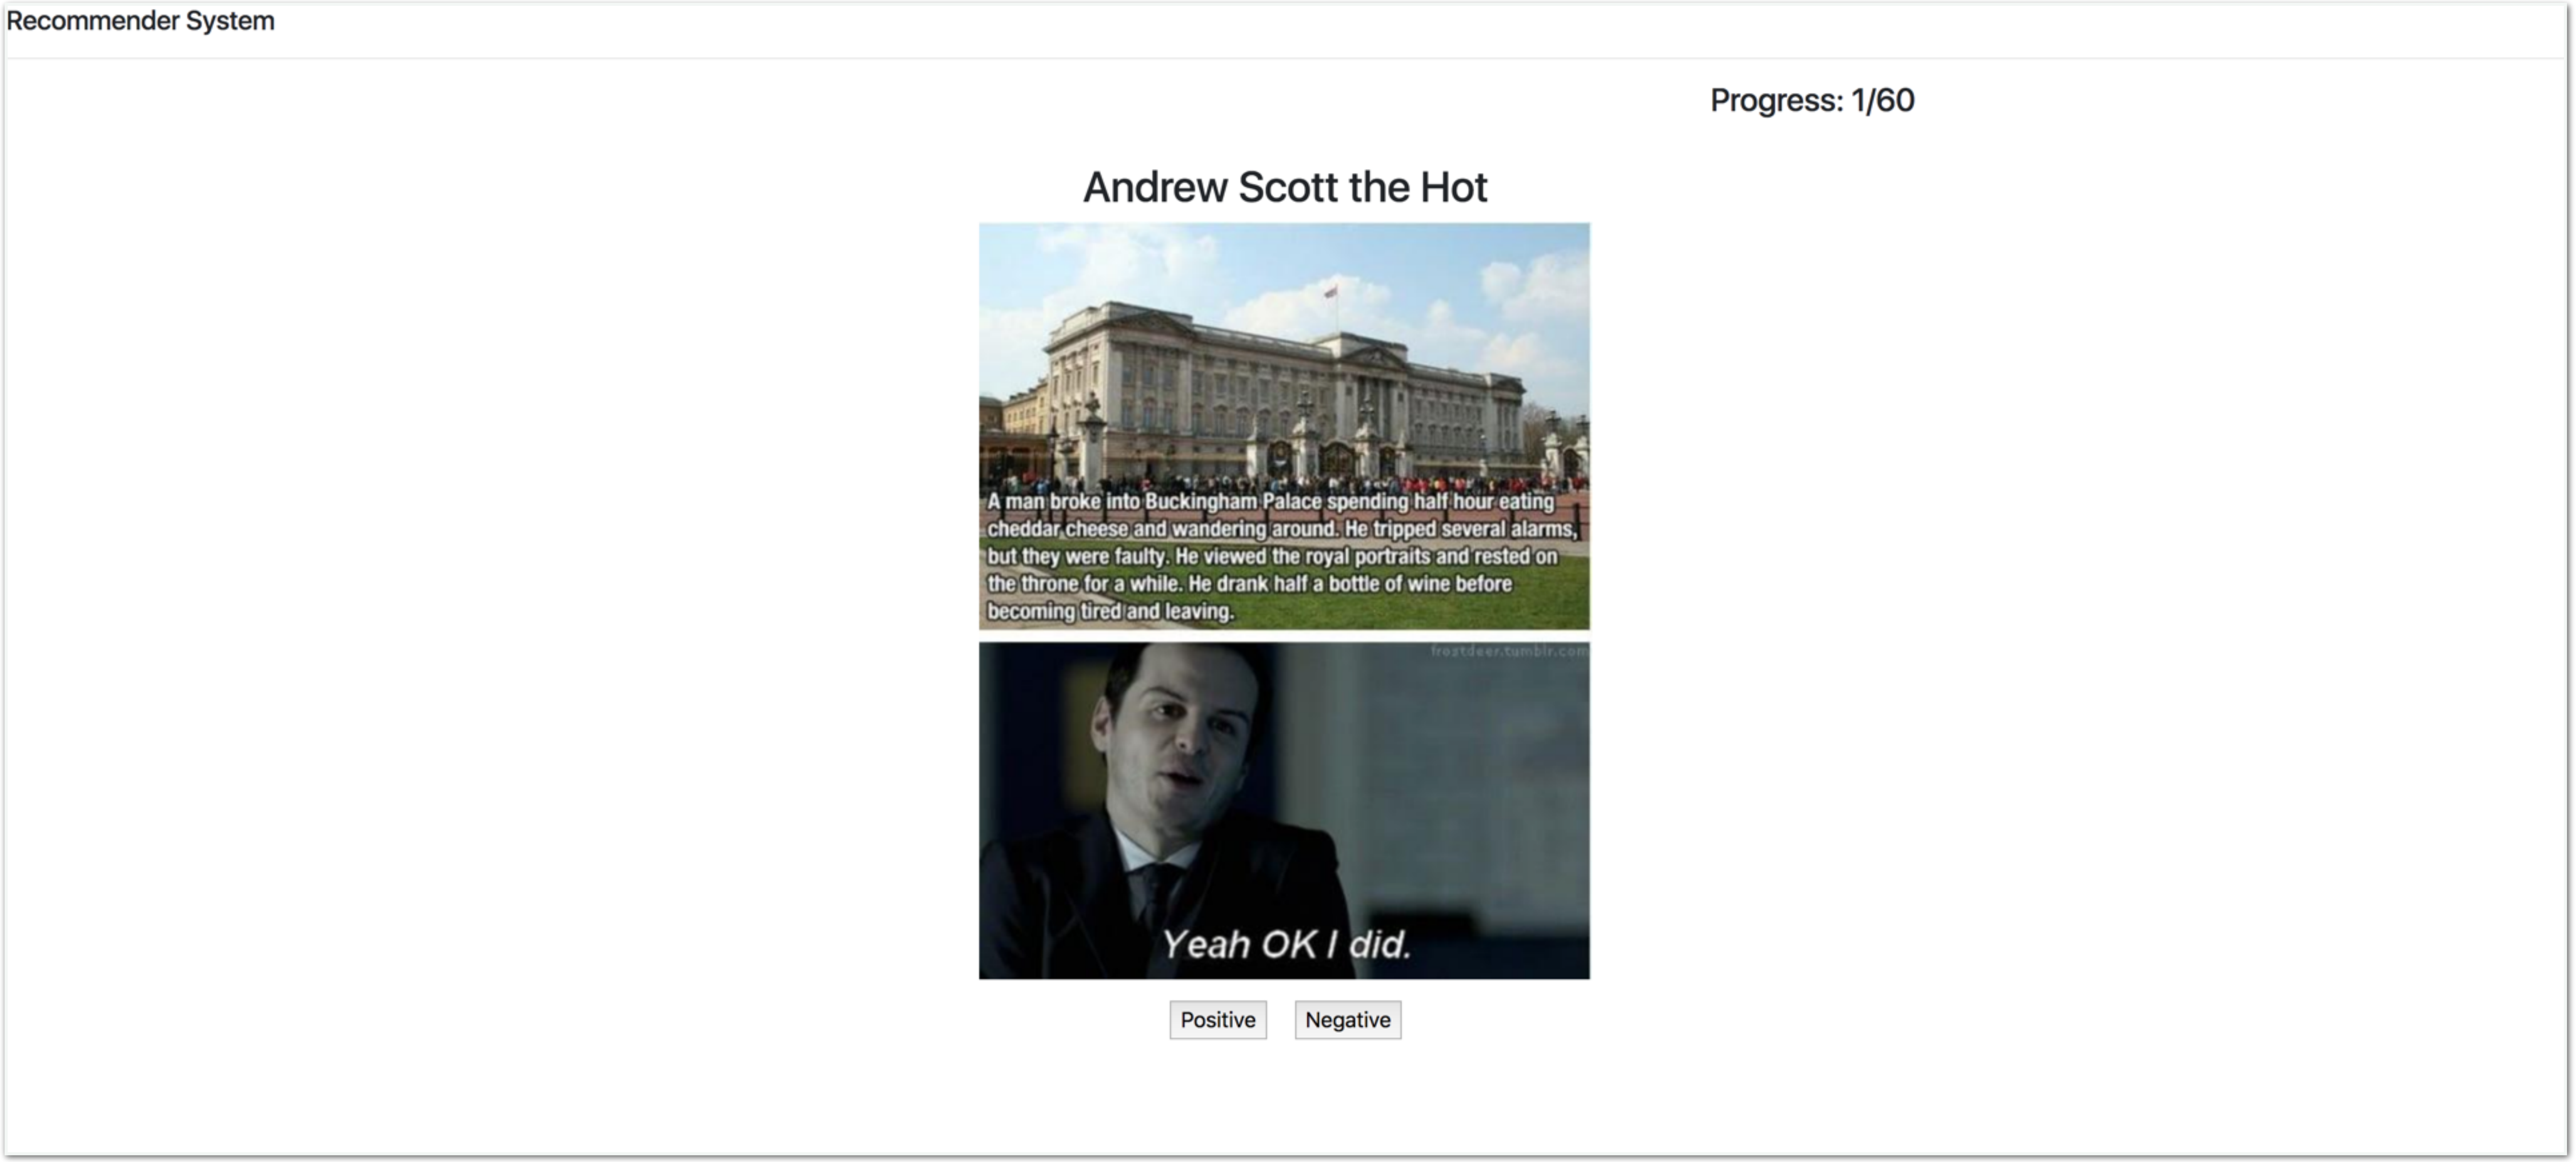

RESULTS

LINEAR MODEL POSITIVE RATINGS

Normal

0.48

0.58

Contextual

0.46

0.61

RESULTS

EXPERIMENTAL EVALUATION

Age			20	26	19	25	23	29	31	26	24	21	25	21	28	24	24	21	21	26	23	25	24
	Gender		M	F	M	M	M	M	F	M	M	F	F	F	M	F	F	F	M	F	F	M	M
		ID	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**	**16**	**17**	**18**	**19**	**20**	**21**
20	M	**1**		0.793	0.400	0.498	0.408	0.598	0.932	0.563	0.455	0.690	0.763	0.719	0.581	0.755	0.771	0.708	0.393	0.850	0.713	0.501	0.500
26	F	**2**			0.824	0.658	0.712	0.710	0.498	0.718	0.703	0.451	0.312	0.501	0.707	0.378	0.399	0.494	0.806	0.418	0.404	0.705	0.690
19	M	**3**				0.368	0.414	0.620	0.883	0.643	0.520	0.693	0.800	0.703	0.643	0.798	0.812	0.691	0.407	0.872	0.775	0.510	0.495
25	M	**4**					0.323	0.451	0.807	0.475	0.403	0.757	0.671	0.749	0.454	0.703	0.709	0.780	0.463	0.719	0.697	0.392	0.392
23	M	**5**						0.480	0.847	0.441	0.348	0.671	0.695	0.673	0.456	0.671	0.717	0.714	0.318	0.726	0.683	0.358	0.361
29	M	**6**							0.689	0.360	0.386	0.855	0.757	0.868	0.305	0.769	0.818	0.868	0.502	0.722	0.824	0.424	0.420
31	F	**7**								0.810	0.886	0.561	0.448	0.629	0.714	0.563	0.568	0.607	0.924	0.516	0.605	0.822	0.874
26	M	**8**									0.326	0.814	0.712	0.817	0.350	0.684	0.709	0.792	0.475	0.698	0.735	0.409	0.405
24	M	**9**										0.733	0.711	0.739	0.412	0.676	0.709	0.717	0.313	0.696	0.693	0.355	0.383
21	F	**10**											0.408	0.345	0.852	0.346	0.371	0.376	0.695	0.494	0.358	0.784	0.700
25	F	**11**												0.442	0.719	0.383	0.391	0.462	0.765	0.389	0.319	0.705	0.712
21	F	**12**													0.851	0.364	0.403	0.392	0.698	0.465	0.337	0.796	0.721
28	M	**13**														0.741	0.773	0.889	0.465	0.711	0.770	0.394	0.419
24	F	**14**															0.236	0.283	0.723	0.303	0.390	0.709	0.683
24	F	**15**																0.258	0.738	0.336	0.398	0.709	0.720
21	F	**16**																	0.715	0.404	0.412	0.776	0.747
21	M	**17**																		0.814	0.693	0.429	0.391
26	F	**18**																			0.439	0.711	0.711
23	F	**19**																				0.719	0.709
25	M	**20**																					0.327
24	M	**21**																					

Color scale for epsilon values:

0.000	0.100	0.200	0.300	0.400	0.500	0.600	0.700	0.800	0.900	1.000

exploitation … exploration

RESULTS

OVERFLOWING EMOTIONS

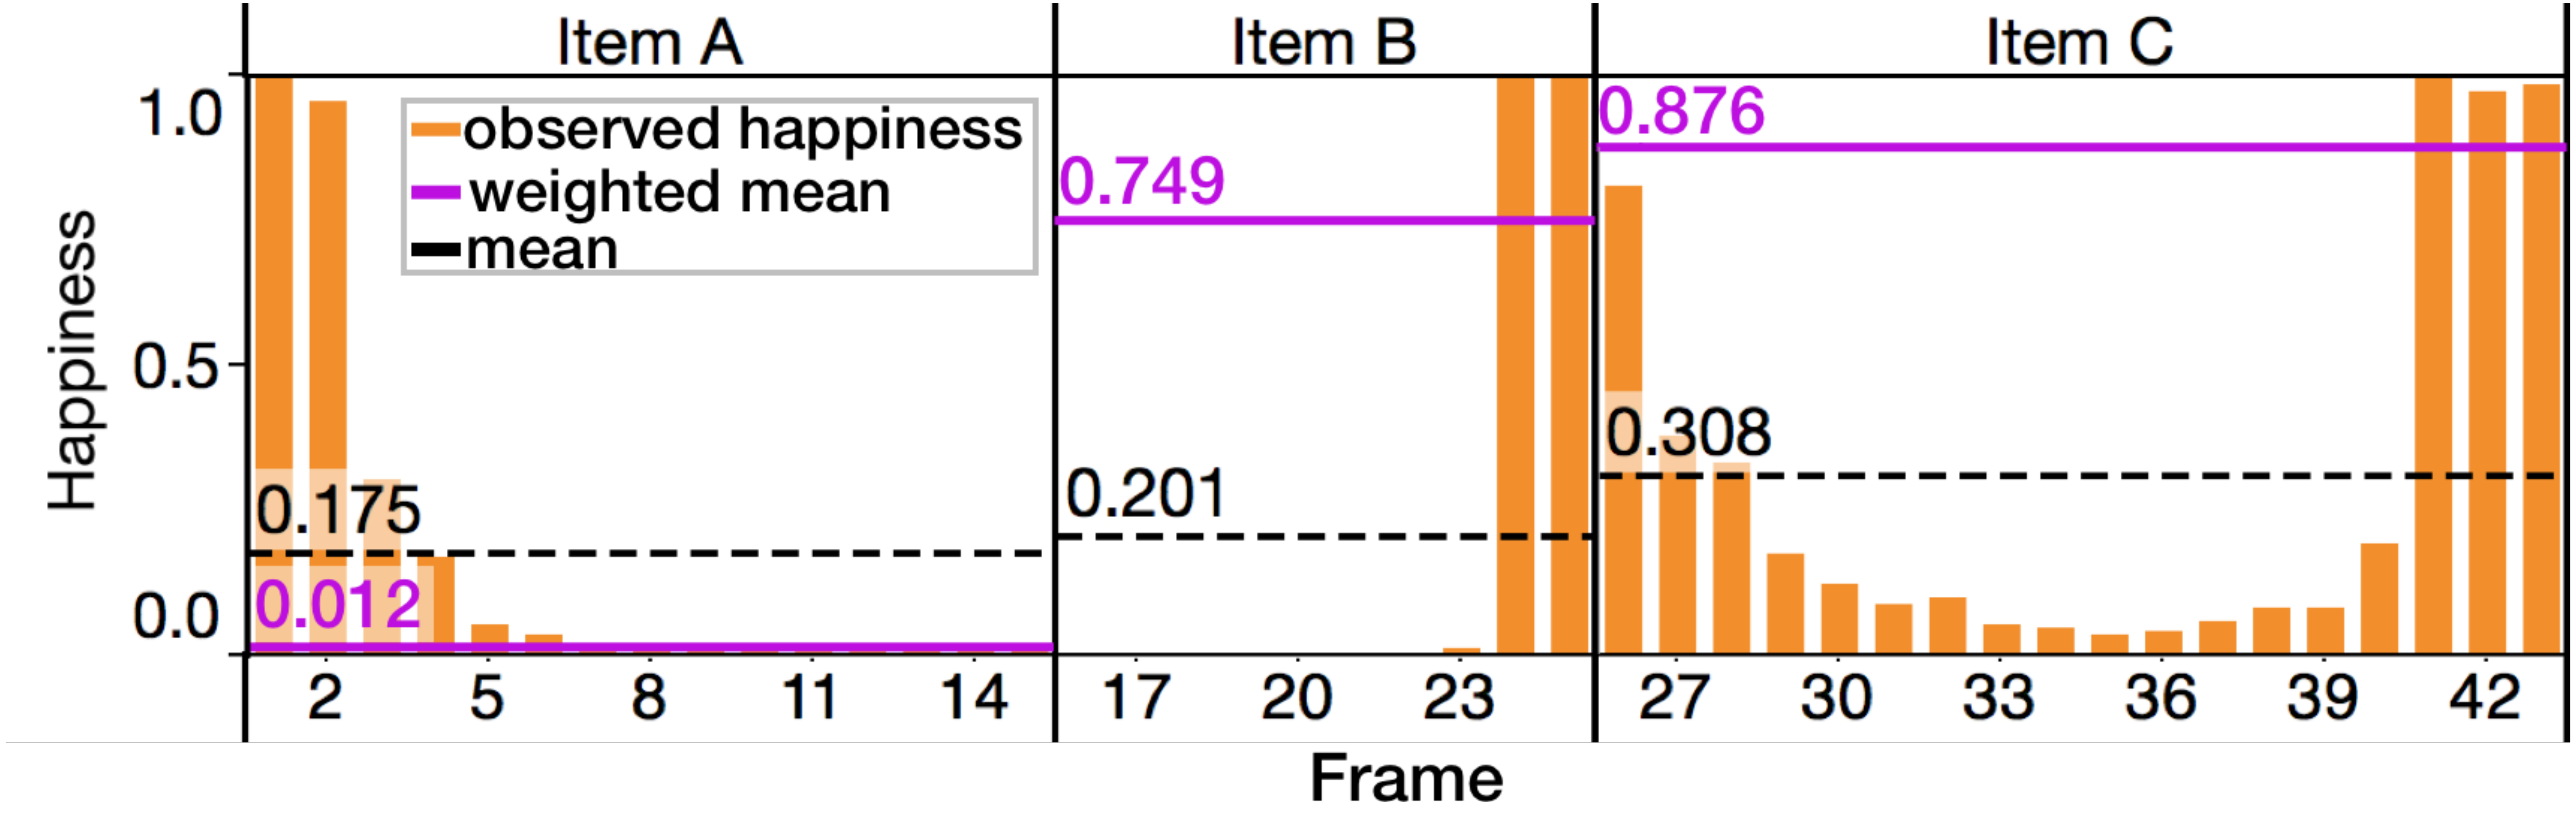

RESULTS

EXPERIMENTAL EVALUATION

I often show my emotions (e.g., laugh out loud, cry ...) while watching movies.

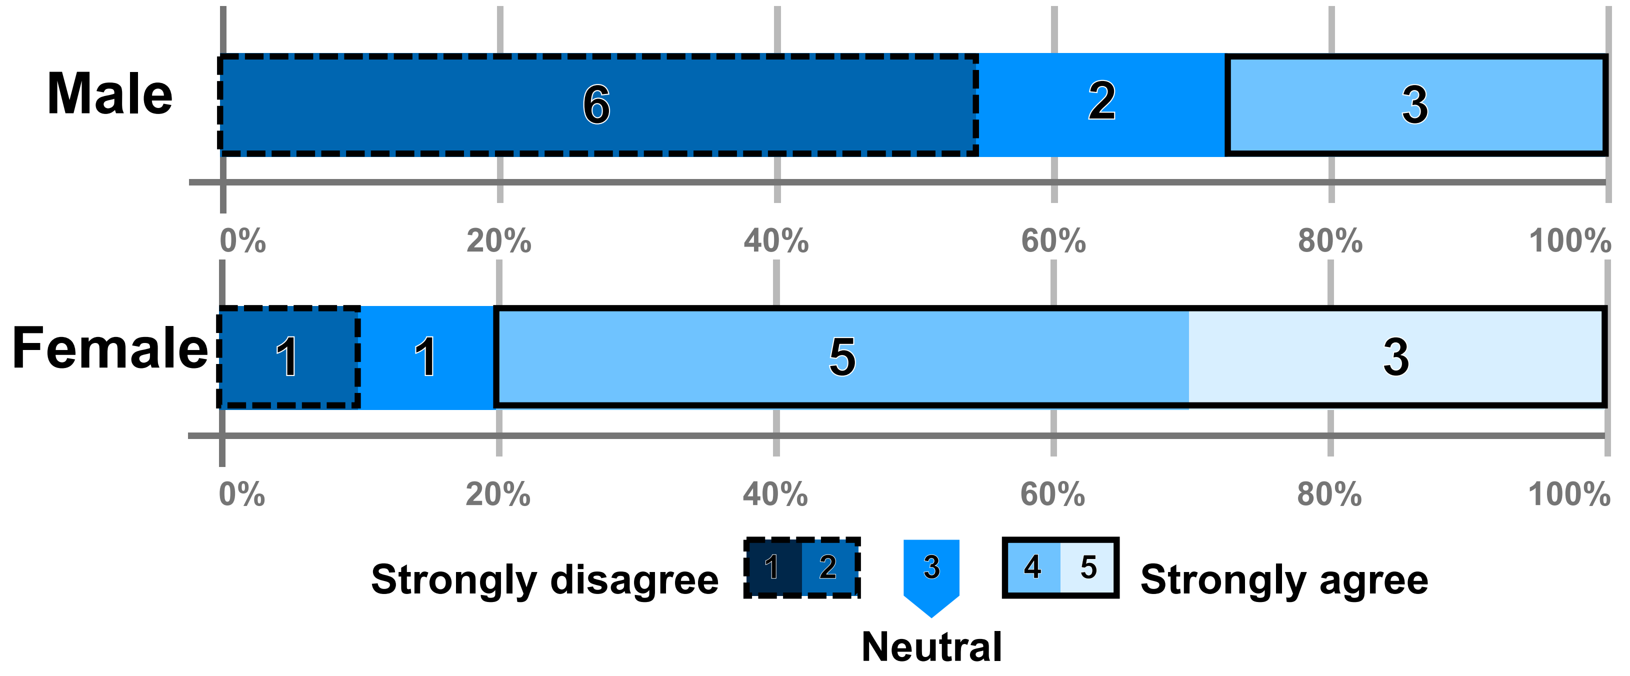

Table 2: Correlation of Emotions with Rating Feedback

Feedback	happiness	neutral	other	n
positive	25.06%	68.90%	6.04%	680
negative	7.24%	86.04%	6.72%	580

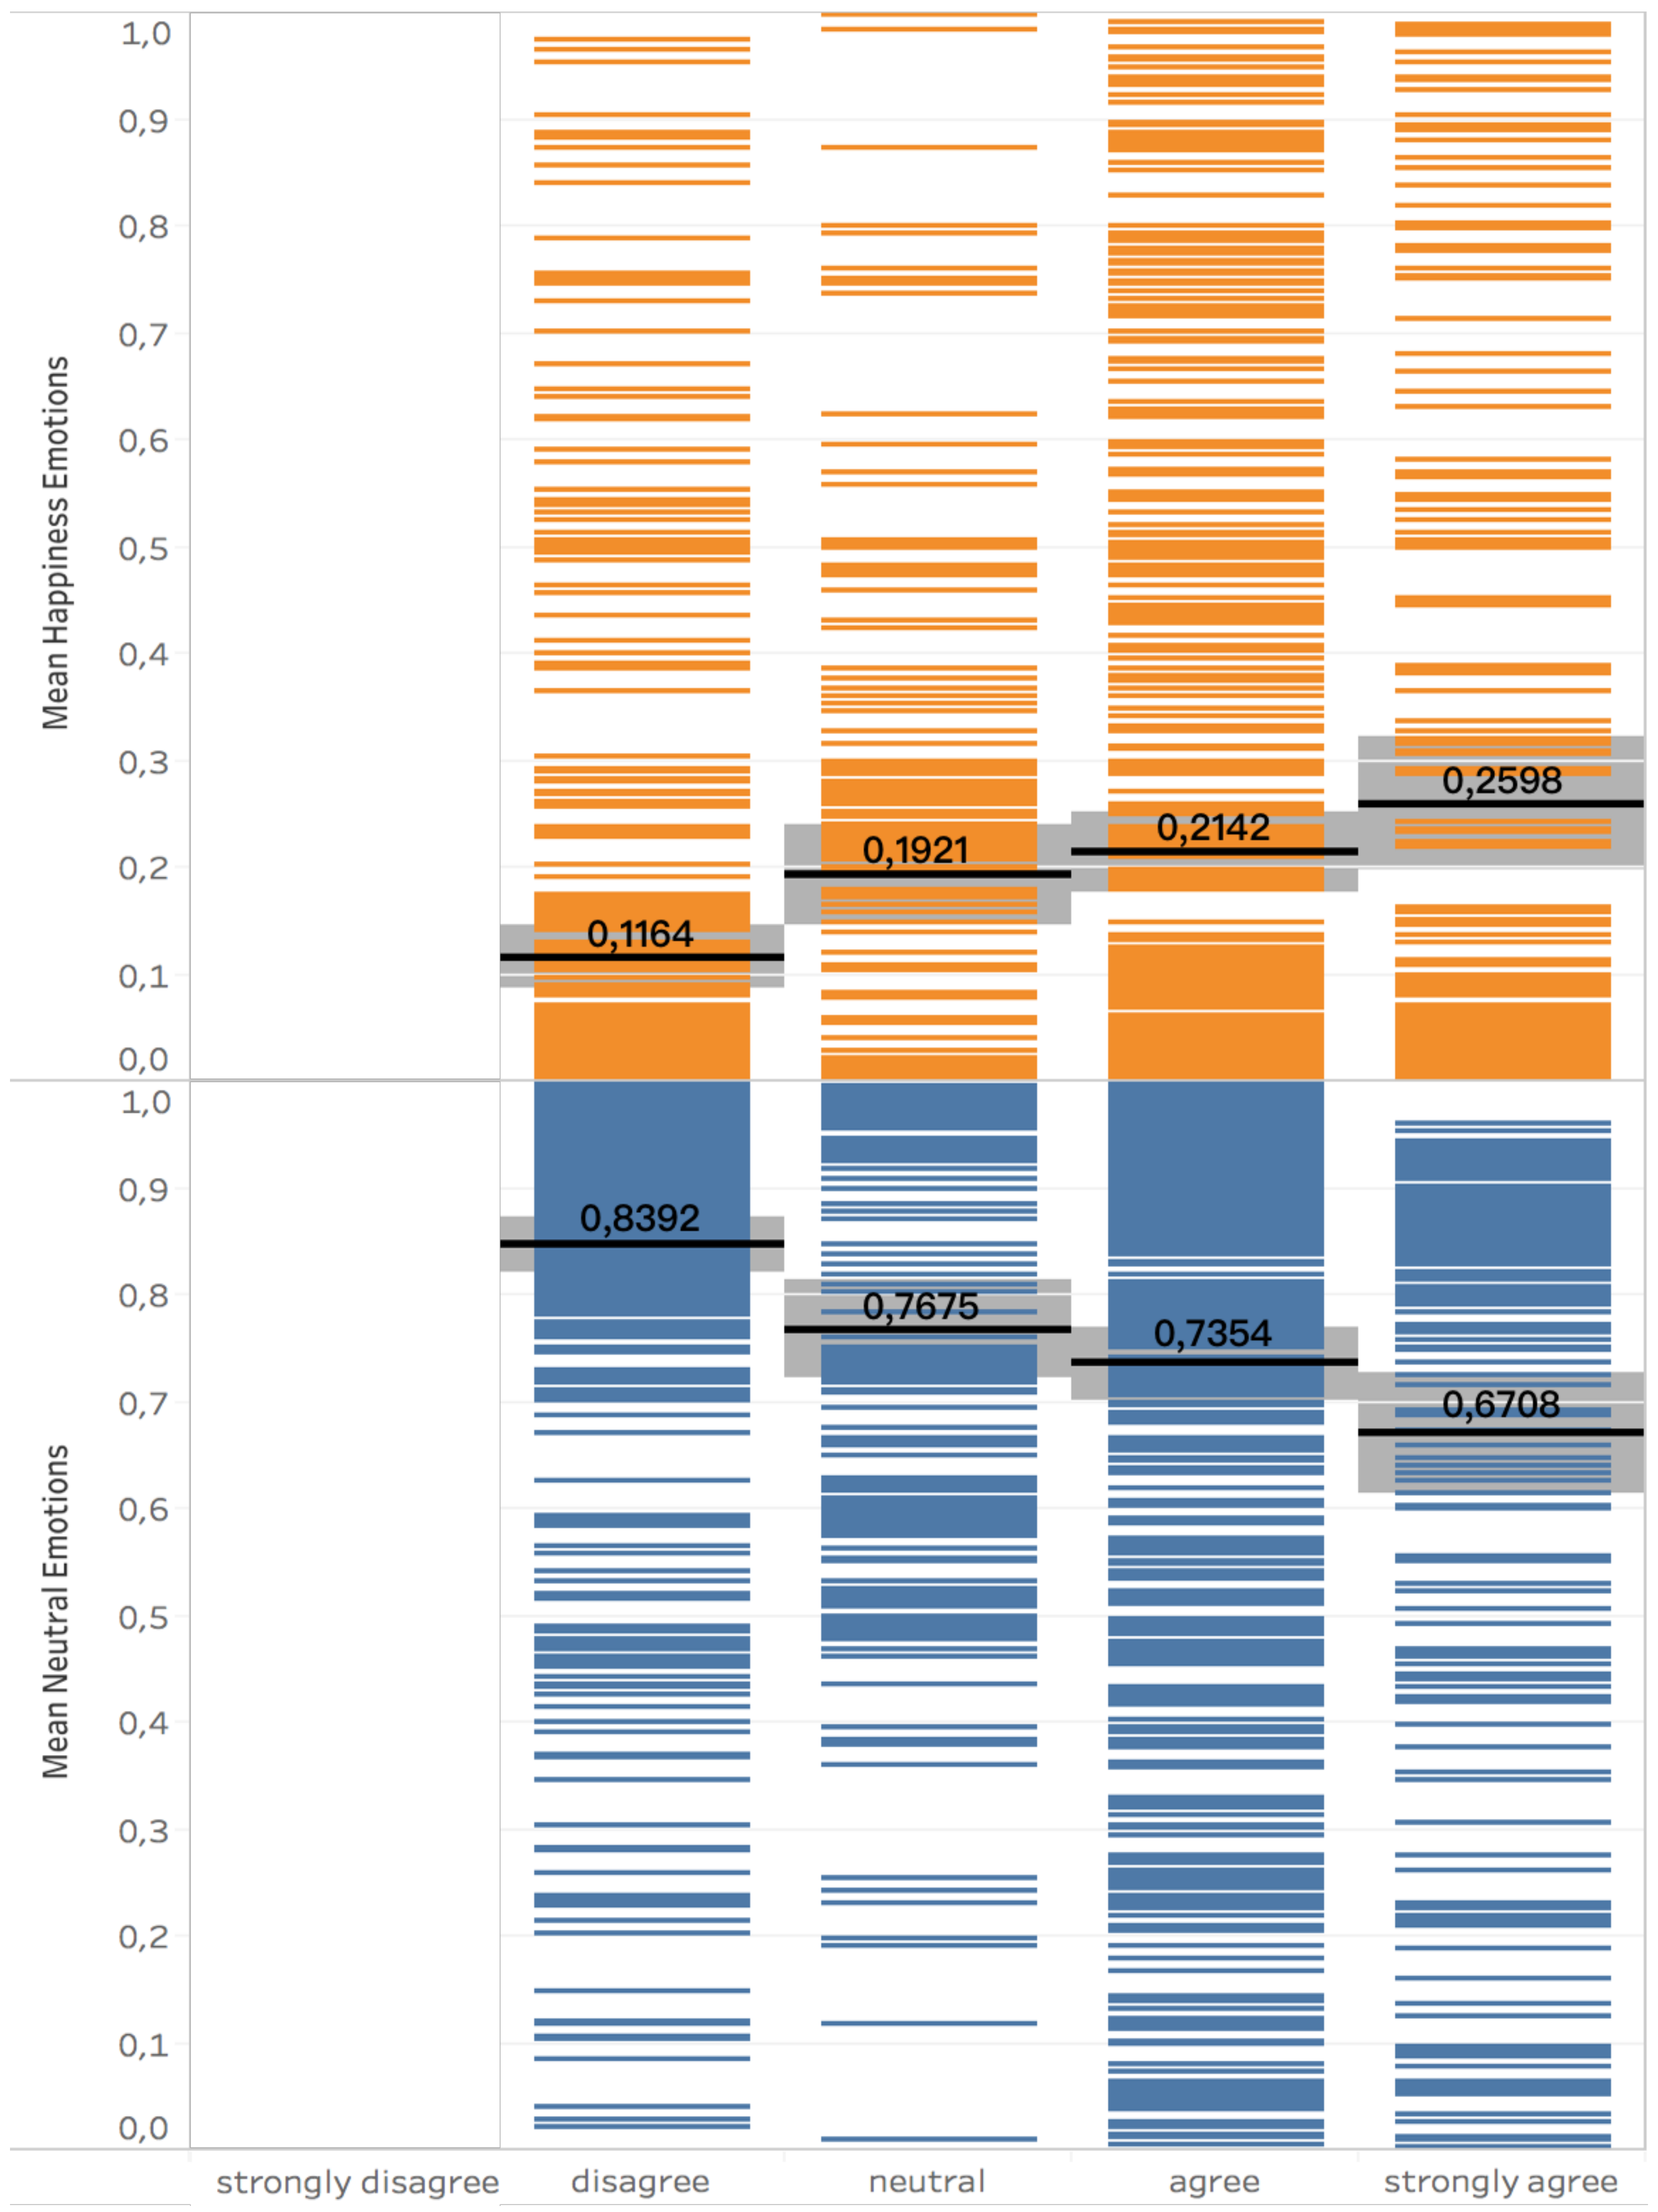

CONCLUSIONS

SUMMARY AND FUTURE WORK

Summary

- Prototype implementation of an affective computing bandit recommender
- Contextual information from computer vision may be helpful
- Affective Recommender Systems have much potential

→ **Not significantly better than the baseline algorithm.**

Future Work

- Future computer vision algorithms may have a better detection of emotions
- Detect more features, e.g., clothing style, location and weather, etc.
- Add further hybrid-modules to improve
- Analyze long-term convergence
- Investigate privacy concerns

THANK YOU

Sebastian Oehme

sebastian.oehme@tum.de

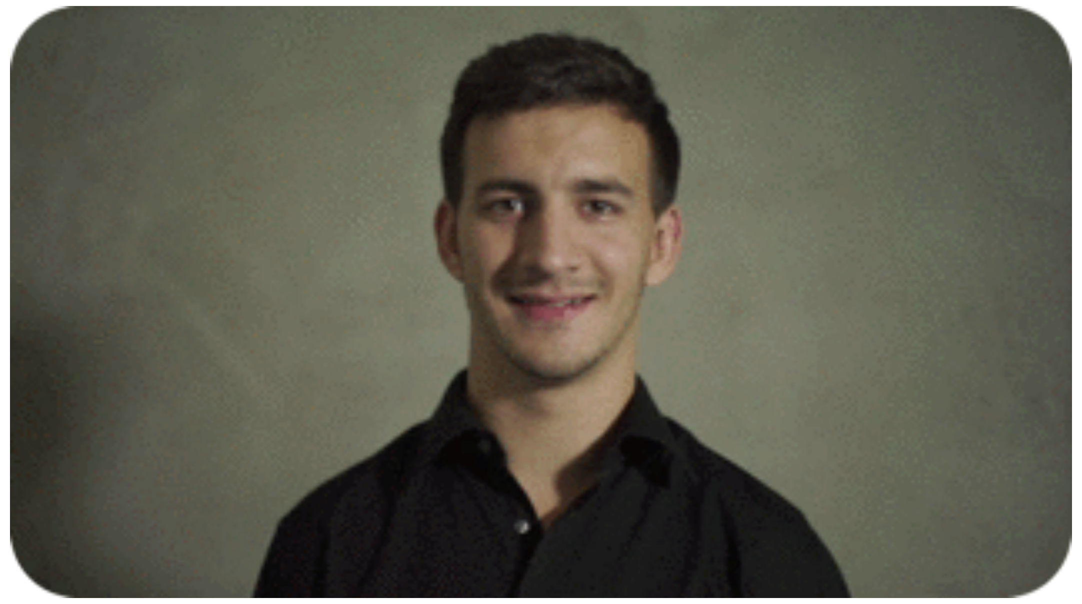

Linus W. Dietz

linus.dietz@tum.de

Paper Download

- Download the full paper (open access): http://ceur-ws.org/Vol-2225/paper1.pdf
- Workshop proceedings available on CEUR-WS: http://ceur-ws.org/Vol-2225

PICTURES

- Wikipedia, User:Colin, CC BY-SA 4.0, https://commons.wikimedia.org/wiki/File:Apple_II_IMG_4212.jpg, checked 24.5.2018
- Medium, Arthur Juliani, https://medium.com/emergent-future/simple-reinforcement-learning-with-tensorflow-part-0-q-learning-with-tables-and-neural-networks-d195264329d0, checked 28.5.2018
- Icons made by Freepik from www.flaticon.com