Towards Real-world Social AI

Paul Pu Liang

Machine Learning Department

Carnegie Mellon University

https://www.cs.cmu.edu/~pliang/

pliang@cs.cmu.edu

https://github.com/pliang279

🐦 @pliang279

Carnegie Mellon University

ML
MACHINE LEARNING
DEPARTMENT

Towards Real-world Social AI

Comprehend human social cues, intents, affective states

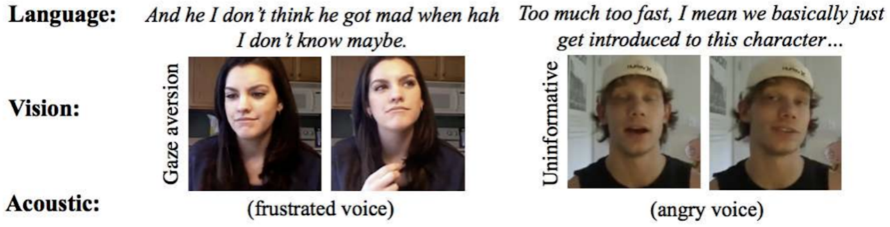

And he I don't think he got mad when hah I don't know maybe.

(frustrated voice)

Too much too fast, I mean we basically just get introduced to this character…

(angry voice)

Engage in social conversation

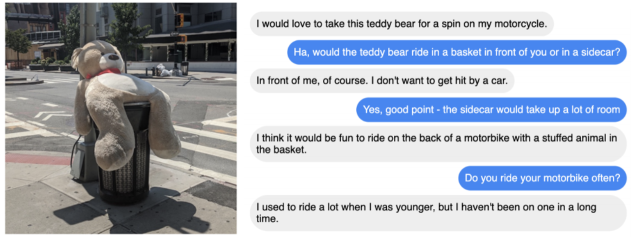

I would love to take this teddy bear for a spin on my motorcycle.

Ha, would the teddy bear ride in a basket in front of you or in a sidecar?

In front of me, of course. I don't want to get hit by a car.

Yes, good point - the sidecar would take up a lot of room

I think it would be fun to ride on the back of a motorbike with a stuffed animal in the basket.

Do you ride your motorbike often?

I used to ride a lot when I was younger, but I haven't been on one in a long time.

Understand social norms and common-sense

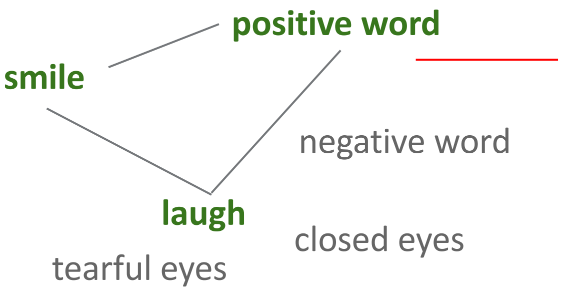

Labels

Emotion

Sentiment

Personalities

Deception

Sarcasm

Humor

Towards Real-world Social AI

Comprehend human social cues, intents, affective states

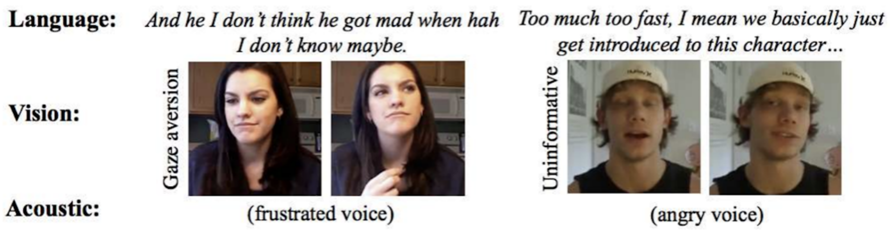

And he I don't think he got mad when hah I don't know maybe.

(frustrated voice)

Too much too fast, I mean we basically just get introduced to this character…

(angry voice)

Engage in social conversation

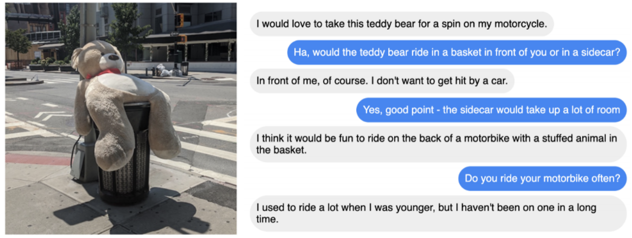

I would love to take this teddy bear for a spin on my motorcycle.

Ha, would the teddy bear ride in a basket in front of you or in a sidecar?

In front of me, of course. I don't want to get hit by a car.

Yes, good point - the sidecar would take up a lot of room

I think it would be fun to ride on the back of a motorbike with a stuffed animal in the basket.

Do you ride your motorbike often?

I used to ride a lot when I was younger, but I haven't been on one in a long time.

Understand social norms and common-sense

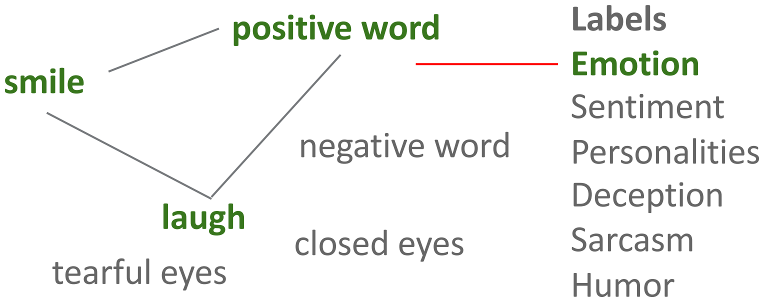

smile

positive word

negative word

laugh

closed eyes

tearful eyes

Labels

Emotion

Sentiment

Personalities

Deception

Sarcasm

Humor

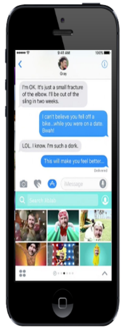

Towards Real-world Social AI

Applications

Comprehend human social cues, intents, affective states

Language: *And he I don't think he got mad when hah I don't know maybe.*

Vision:

Acoustic: (frustrated voice)

Engage in social conversation

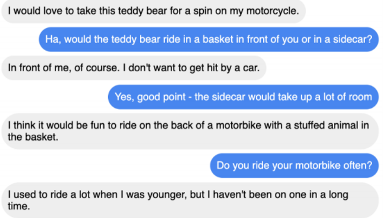

I would love to take this teddy bear for a spin on my motorcycle.

Ha, would the teddy bear ride in a basket in front of you or in a sidecar?

In front of me, of course. I don't want to get hit by a car.

Yes, good point - the sidecar would take up a lot of room

I think it would be fun to ride on the back of a motorbike with a stuffed animal in the basket.

Do you ride your motorbike often?

I used to ride a lot when I was younger, but I haven't been on one in a long time.

Understand social norms and common-sense

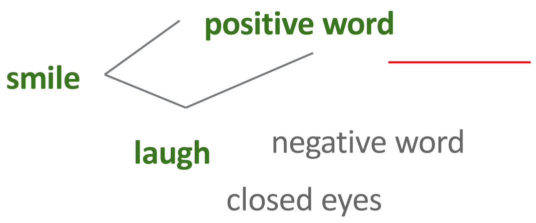

smile

positive word

laugh

negative word

closed eyes

Emotion

Sentiment

Personalities

Deception

Sarcasm

Humor

Algorithms

Multimodal perception

Utterance: *"Great, now he is waving back"*

Emotion: *Disgust* **Sentiment:** *Negative*

Text	Audio	Visual
Positive/Joy	Flat tone	Frown

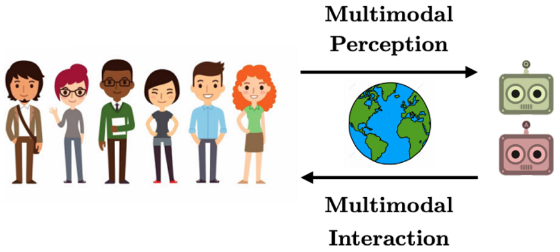

Multimodal interaction

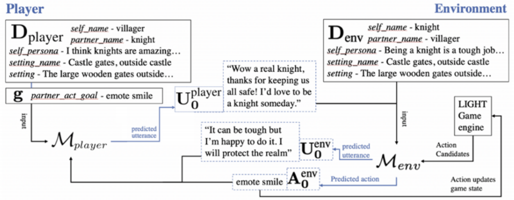
Player
Environment
D_{player} *self_name* - villager
partner_name - knight
self_persona - I think knights are amazing…
setting_name - Castle gates, outside castle
setting - The large wooden gates outside…
D_{env} *self_name* - knight
partner_name - villager
self_persona - Being a knight is a tough job…
setting_name - Castle gates, outside castle
setting - The large wooden gates outside…
g *partner_act_goal* - emote smile
U_0^{player}
"Wow a real knight, thanks for keeping us all safe! I'd love to be a knight someday."
input
predicted utterance
$\mathcal{M}_{player}$
"It can be tough but I'm happy to do it. I will protect the realm"
U_0^{env}
predicted utterance
$\mathcal{M}_{env}$
LIGHT Game engine
Action Candidates
emote smile
A_0^{env}
Predicted action
Action updates game state

Towards Real-world Social AI

Applications

Comprehend human social cues, intents, affective states

Language: *And he I don't think he got mad when hah I don't know maybe.*

Vision:

Acoustic: (frustrated voice)

Engage in social conversation

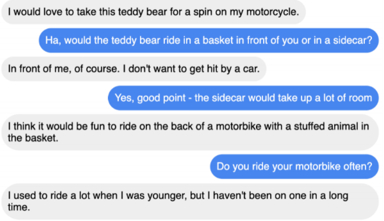

I would love to take this teddy bear for a spin on my motorcycle.

Ha, would the teddy bear ride in a basket in front of you or in a sidecar?

In front of me, of course. I don't want to get hit by a car.

Yes, good point - the sidecar would take up a lot of room

I think it would be fun to ride on the back of a motorbike with a stuffed animal in the basket.

Do you ride your motorbike often?

I used to ride a lot when I was younger, but I haven't been on one in a long time.

Understand social norms and common-sense

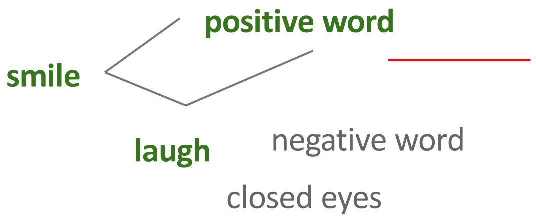

Emotion
Sentiment
Personalities
Deception
Sarcasm
Humor

Algorithms

Multimodal perception

Utterance: *"Great, now he is waving back"*

Emotion: *Disgust* **Sentiment:** *Negative*

Text	Audio	Visual
Positive/Joy	Flat tone	Frown

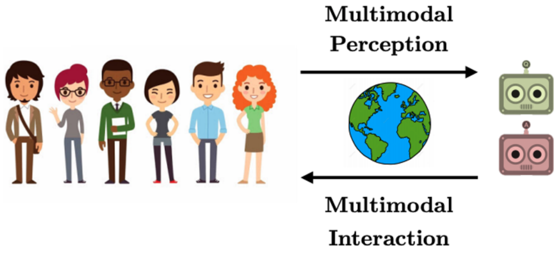

Multimodal interaction

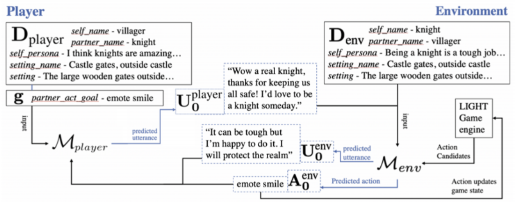
Player
Environment
D_{player} self_name - villager
partner_name - knight
self_persona - I think knights are amazing…
setting_name - Castle gates, outside castle
setting - The large wooden gates outside…
g partner_act_goal - emote smile
U_0^{player}
"Wow a real knight, thanks for keeping us all safe! I'd love to be a knight someday."
input
$\mathcal{M}_{player}$
predicted utterance
D_{env} self_name - knight
partner_name - villager
self_persona - Being a knight is a tough job…
setting_name - Castle gates, outside castle
setting - The large wooden gates outside…
input
LIGHT Game engine
"It can be tough but I'm happy to do it. I will protect the realm"
U_0^{env}
predicted utterance
$\mathcal{M}_{env}$
Action Candidates
emote smile
A_0^{env}
Predicted action
Action updates game state

Foundations

Fusion

Representation

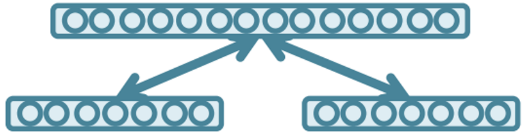

Alignment

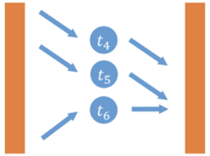

Translation

Big dog
on the
beach

Co-learning

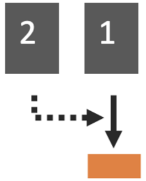

Towards Real-world Social AI

Resources: https://github.com/pliang279/awesome-multimodal-ml

Real-world

Robustness

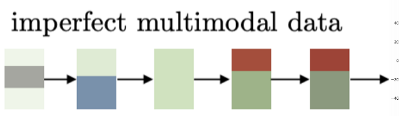

Fair learning

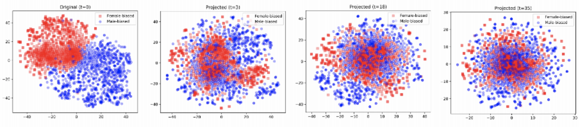

Privacy-preserving

Generalizable to low-resource

Speech (text in parenthesis)

(Beda Yesus agot gu ofa oida Bua buroru Didif ojgomu)

Applications

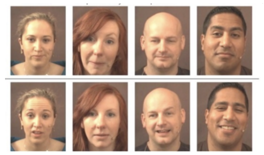

Applications

Comprehend human social cues, intents, affective states

Language: *And he I don't think he got mad when hah I don't know maybe.*

Vision:

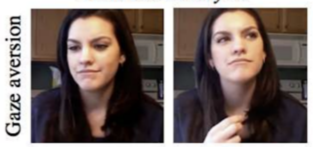

Acoustic: (frustrated voice)

Engage in social conversation

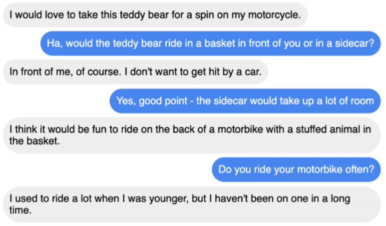

I would love to take this teddy bear for a spin on my motorcycle.

Ha, would the teddy bear ride in a basket in front of you or in a sidecar?

In front of me, of course. I don't want to get hit by a car.

Yes, good point - the sidecar would take up a lot of room

I think it would be fun to ride on the back of a motorbike with a stuffed animal in the basket.

Do you ride your motorbike often?

I used to ride a lot when I was younger, but I haven't been on one in a long time.

Understand social norms and common-sense

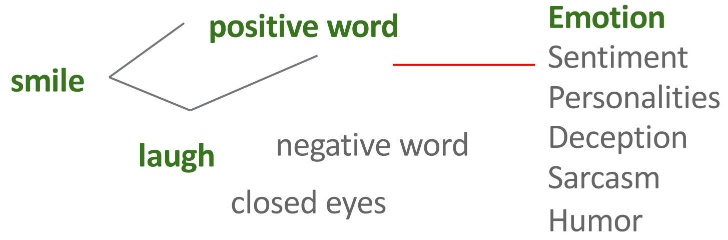

smile

positive word

laugh

negative word

closed eyes

Emotion

Sentiment

Personalities

Deception

Sarcasm

Humor

Algorithms

Multimodal perception

Utterance: *"Great, now he is waving back"*

Emotion: *Disgust* **Sentiment:** *Negative*

Text	Audio	Visual
Positive/Joy	Flat tone	Frown

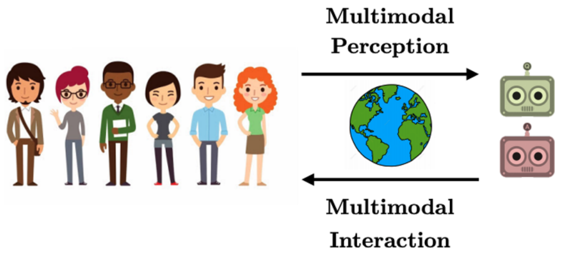

Multimodal interaction

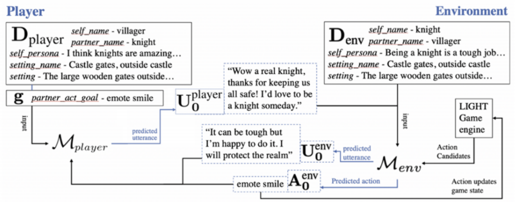
Player
Environment
D_{player}
self_name - villager
partner_name - knight
self_persona - I think knights are amazing…
setting_name - Castle gates, outside castle
setting - The large wooden gates outside…
D_{env}
self_name - knight
partner_name - villager
self_persona - Being a knight is a tough job…
setting_name - Castle gates, outside castle
setting - The large wooden gates outside…
g partner_act_goal - emote smile
U_0^{player}
"Wow a real knight, thanks for keeping us all safe! I'd love to be a knight someday."
input
$\mathcal{M}_{player}$
predicted utterance
"It can be tough but I'm happy to do it. I will protect the realm"
U_0^{env}
predicted utterance
input
LIGHT Game engine
Action Candidates
$\mathcal{M}_{env}$
emote smile
A_0^{env}
Predicted action
Action updates game state

Foundations

Fusion

Representation

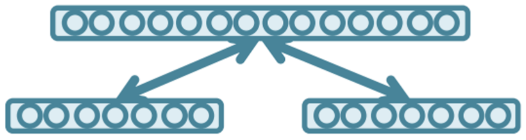

Alignment

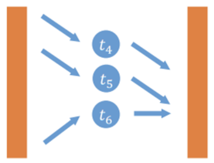

Translation

Big dog on the beach

Co-learning

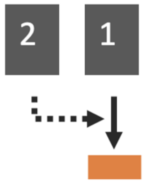

Towards Real-world Social AI

Resources: https://github.com/pliang279/awesome-multimodal-ml

Real-world

Robustness

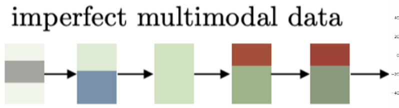

Fair learning

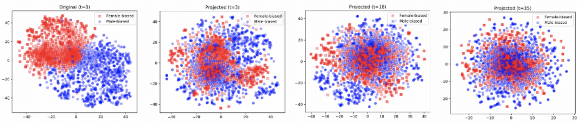

Privacy-preserving

Generalizable to low-resource

SPEECH (TEXT IN PARENTHESIS)

(Beda Yesus agot gu ofa oida Bua buroru Didif ojgomu)

Applications

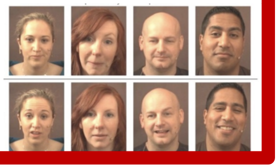

Applications

Comprehend human social cues, intents, affective states

Language: *And he I don't think he got mad when hah I don't know maybe.*

Vision:

Acoustic: (frustrated voice)

Engage in social conversation

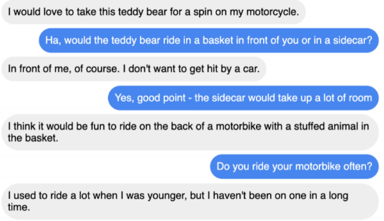

I would love to take this teddy bear for a spin on my motorcycle.

Ha, would the teddy bear ride in a basket in front of you or in a sidecar?

In front of me, of course. I don't want to get hit by a car.

Yes, good point - the sidecar would take up a lot of room

I think it would be fun to ride on the back of a motorbike with a stuffed animal in the basket.

Do you ride your motorbike often?

I used to ride a lot when I was younger, but I haven't been on one in a long time.

Understand social norms and common-sense

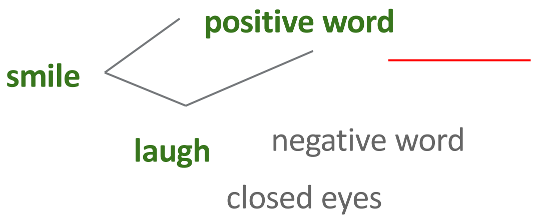

Emotion

Sentiment

Personalities

Deception

Sarcasm

Humor

Algorithms

Multimodal perception

Utterance: *"Great, now he is waving back"*

Emotion: *Disgust* **Sentiment:** *Negative*

Text	Audio	Visual
Positive/Joy	Flat tone	Frown

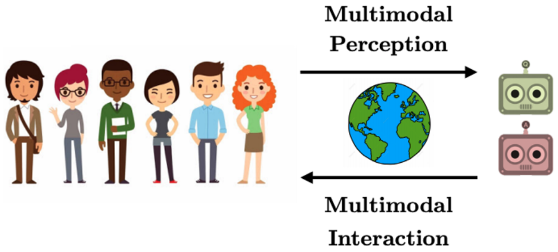

Multimodal interaction

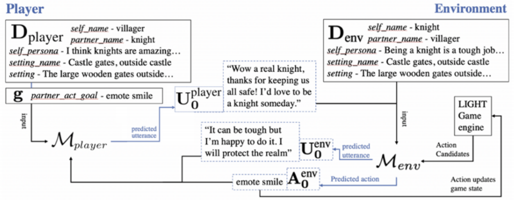

D_{player} *self_name* - villager *partner_name* - knight *self_persona* - I think knights are amazing… *setting_name* - Castle gates, outside castle *setting* - The large wooden gates outside…

g *partner_act_goal* - emote smile

D_{env} *self_name* - knight *partner_name* - villager *self_persona* - Being a knight is a tough job… *setting_name* - Castle gates, outside castle *setting* - The large wooden gates outside…

"Wow a real knight, thanks for keeping us all safe! I'd love to be a knight someday."

"It can be tough but I'm happy to do it. I will protect the realm"

emote smile

Foundations

Fusion

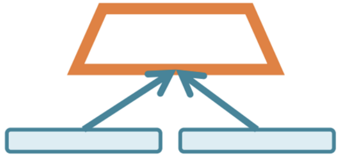

Representation

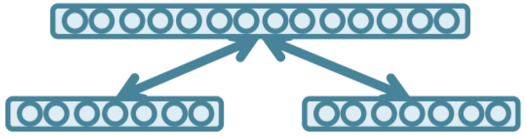

Alignment

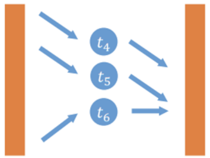

Translation

Big dog on the beach

Co-learning

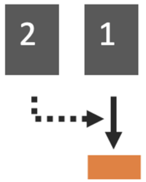

Multimodal Benchmarks

Large benchmarks for multimodal affect recognition

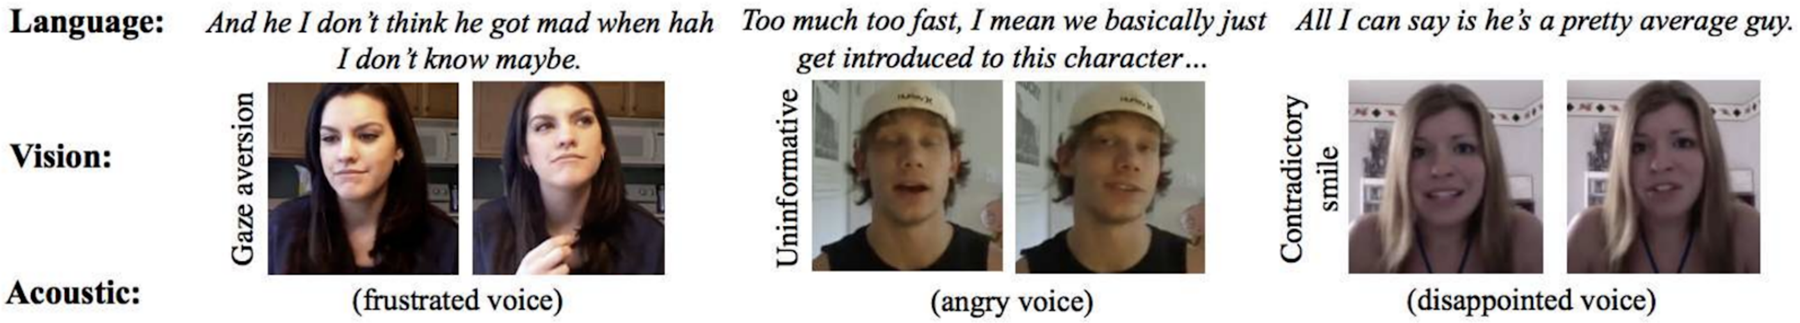

And he I don't think he got mad when hah I don't know maybe.

Too much too fast, I mean we basically just get introduced to this character…

All I can say is he's a pretty average guy.

Liang et al., Computational Modeling of Human Multimodal Language. Master's Thesis 2018

Multimodal Benchmarks

Large benchmarks for multimodal affect recognition

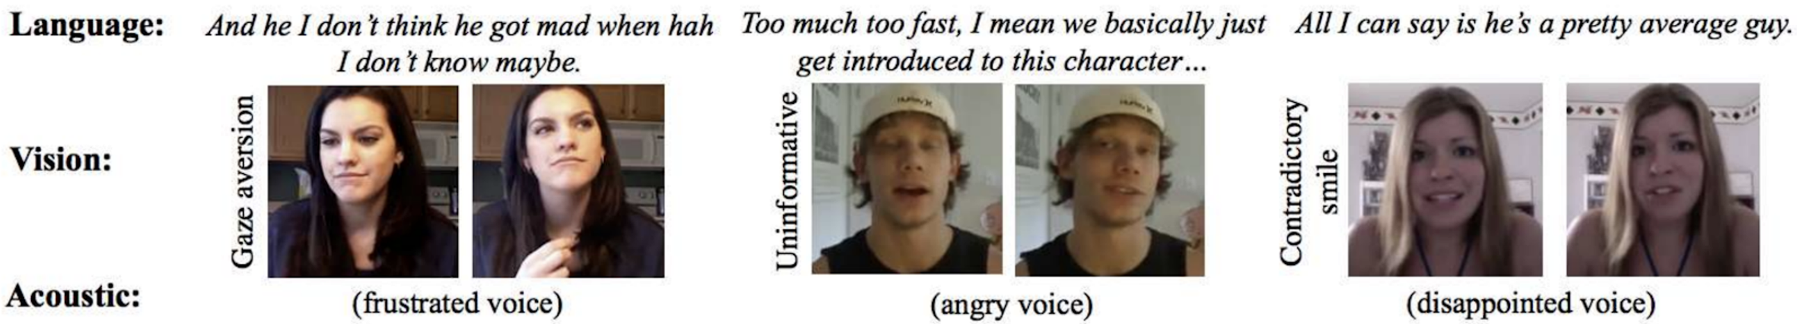

Language: *And he I don't think he got mad when hah I don't know maybe.* *Too much too fast, I mean we basically just get introduced to this character…* *All I can say is he's a pretty average guy.*

1,000 speakers

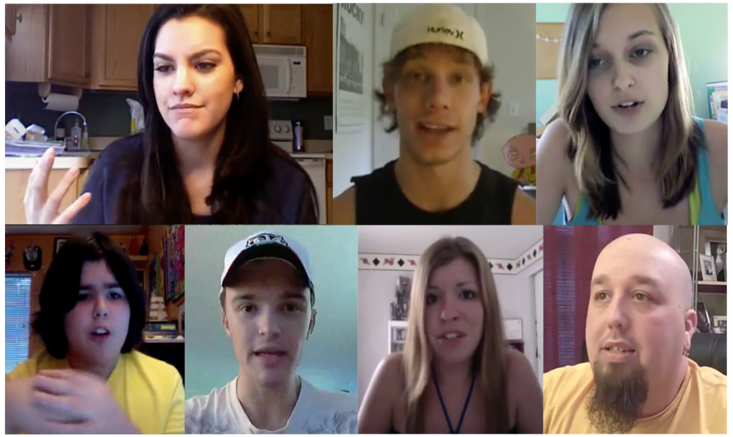

250 topics

consumer
independent movie
conference
application
testimony
monologue
sector
courses industry
business
counter
marketplace
entrepreneurship
summary
answers
weekly
online
sustainability
definition
advertising
updates
hearing
consulting
pricing
faq
instruction
debates
thoughts
economics
investing
hear distress
consumers
Response
economies
social
buyers
lecture
listener
review
home
narrative
financing
presentation
seminar
customers
equity
remarks
investors
summit
reviews
conviction
corporate
speech
loans
congress
statement
integrated
update
products
politics
companies
updated
questions
analysis
firms
retail
investment
meeting
financial
businesses
speeches
advertisers
product
religious
sociology
comments
stocks
placement
eulogy
description
employers
committee
economic
evaluation
speaker
topic
outsourcing
political
retailers
dialogue
banking
seeing
convention
customer
marketing
phd
separate
person

Diverse annotations

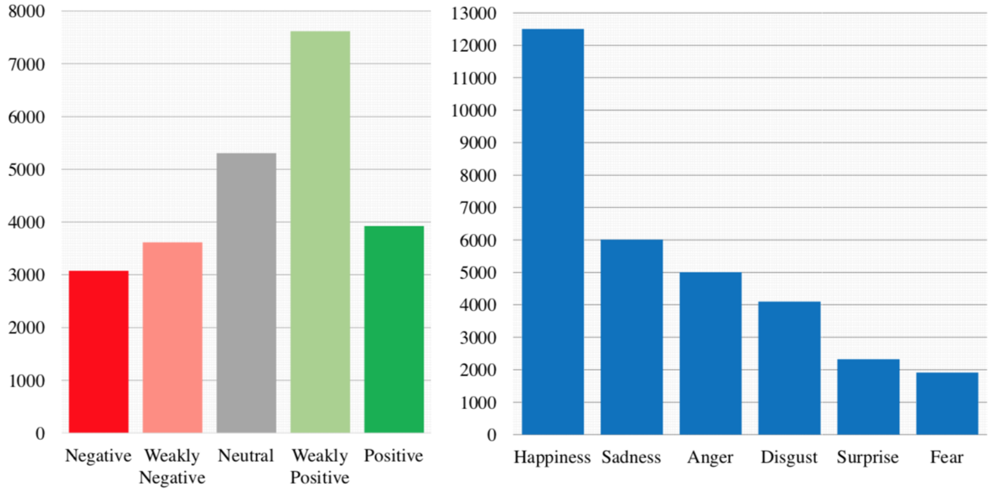

Liang et al., Computational Modeling of Human Multimodal Language. Master's Thesis 2018

Multiscale Benchmarks for Multimodal Learning

Domains

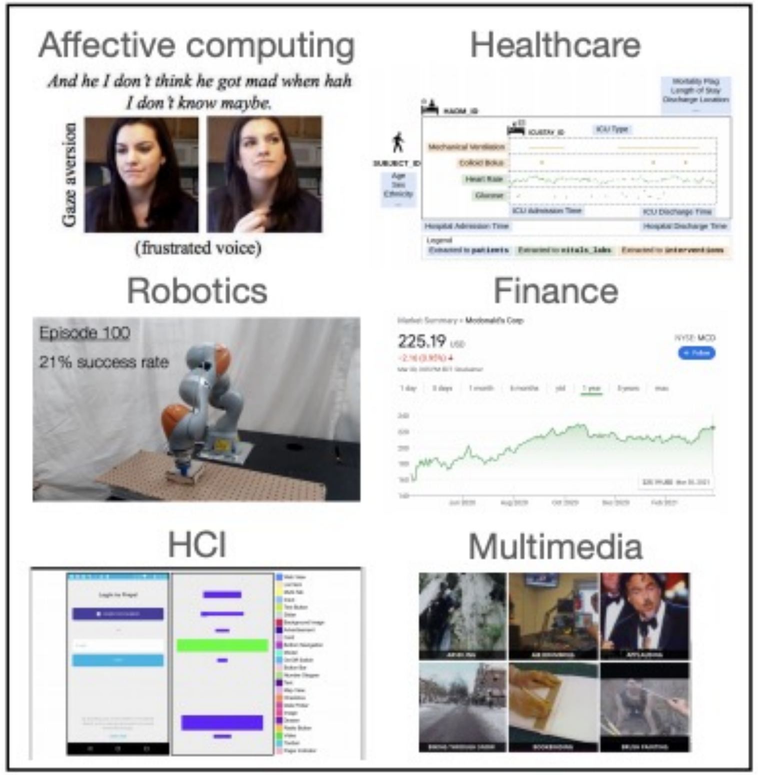

Modalities

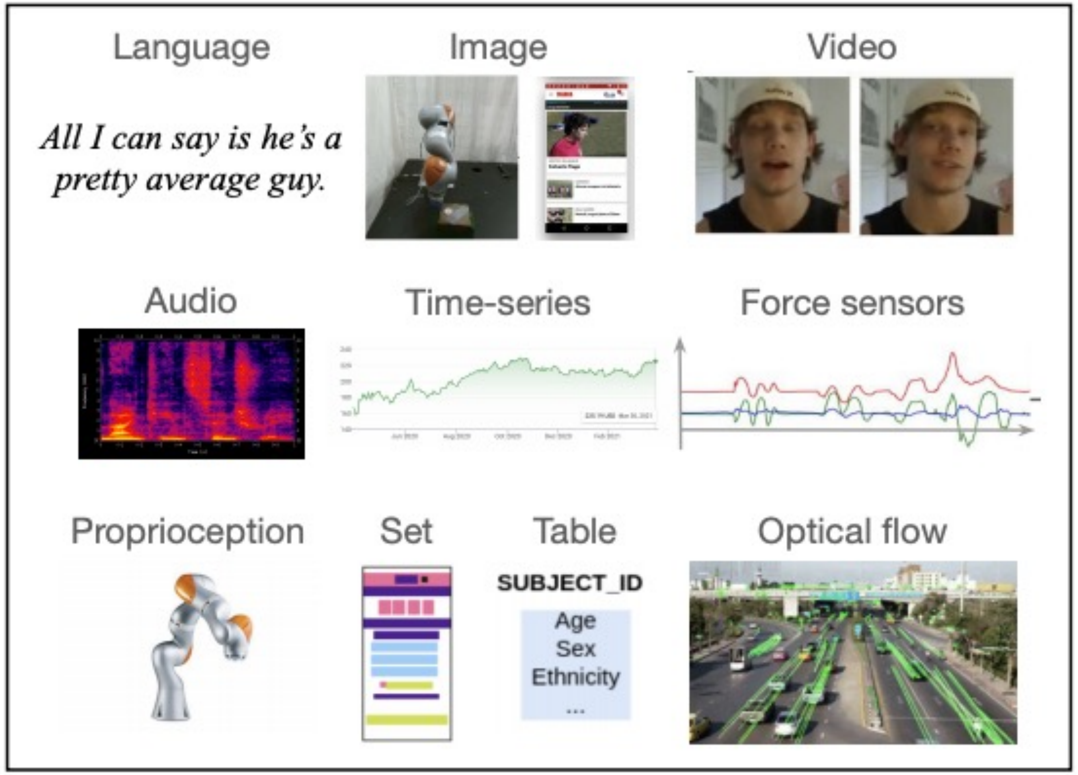

Evaluation

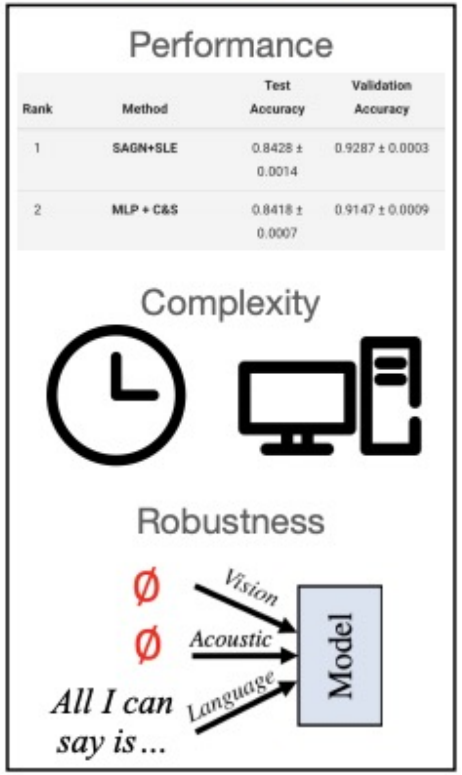

Rank	Method	Test Accuracy	Validation Accuracy
1	SAGN+SLE	0.8428 ± 0.0014	0.9287 ± 0.0003
2	MLP + C&S	0.8418 ± 0.0007	0.9147 ± 0.0009

Liang et al., MultiBench: Multiscale Benchmarks for Multimodal Representation Learning. NeurIPS 2021 Benchmark Track

Multiscale Benchmarks for Multimodal Learning

Standardized implementation of >20 multimodal methods

Data preprocessing

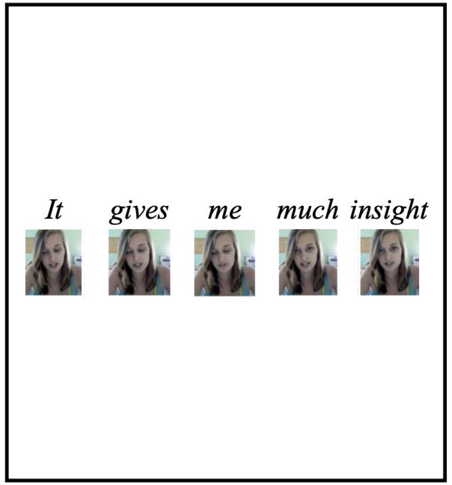

Unimodal models

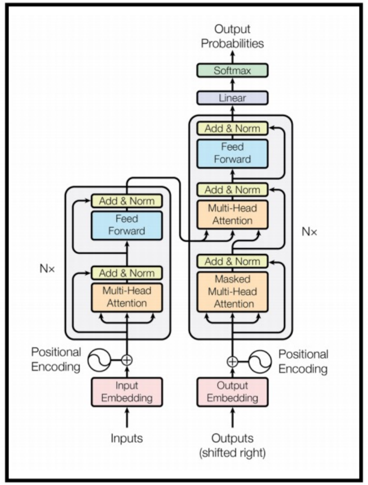

Fusion paradigms

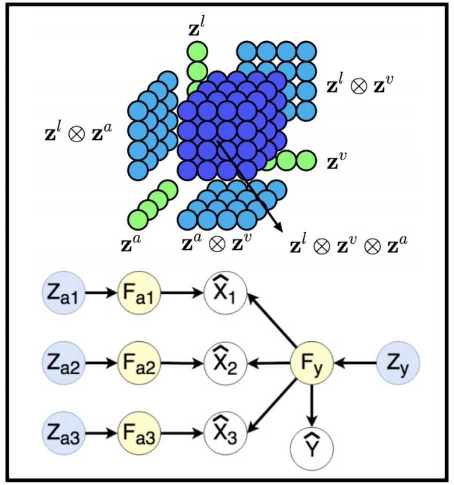
$\mathbf{z}^l$
$\mathbf{z}^l \otimes \mathbf{z}^v$
$\mathbf{z}^l \otimes \mathbf{z}^a$
$\mathbf{z}^v$
$\mathbf{z}^a$
$\mathbf{z}^a \otimes \mathbf{z}^v$
$\mathbf{z}^l \otimes \mathbf{z}^v \otimes \mathbf{z}^a$
Z_{a1}
F_{a1}
$\hat{X}_1$
Z_{a2}
F_{a2}
$\hat{X}_2$
F_y
Z_y
Z_{a3}
F_{a3}
$\hat{X}_3$
$\hat{Y}$

Optimization objectives

$$\mathcal{L}_{\text{sup}} = -\log p(y|\mathbf{x}_1, \mathbf{x}_2)$$

$$\mathcal{L}_{\text{CCA}} = -\text{corr}\left(g_1(\mathbf{z}_1), g_2(\mathbf{z}_2)\right)$$

$$\mathcal{L}_{\text{rec}} = \|g_1(\mathbf{z}_{\text{mm}}) - \mathbf{x}_1\|_2 + \|g_2(\mathbf{z}_{\text{mm}}) - \mathbf{x}_2\|_2$$

Training procedures

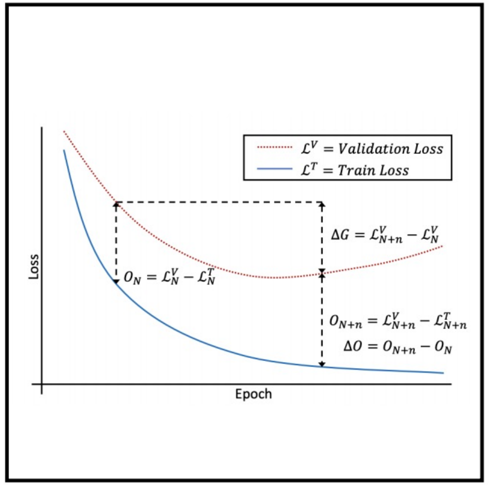

Liang et al., MultiBench: Multiscale Benchmarks for Multimodal Representation Learning. NeurIPS 2021 Benchmark Track

Multiscale Benchmarks for Multimodal Learning

Methods struggle to perform outside of their own domain

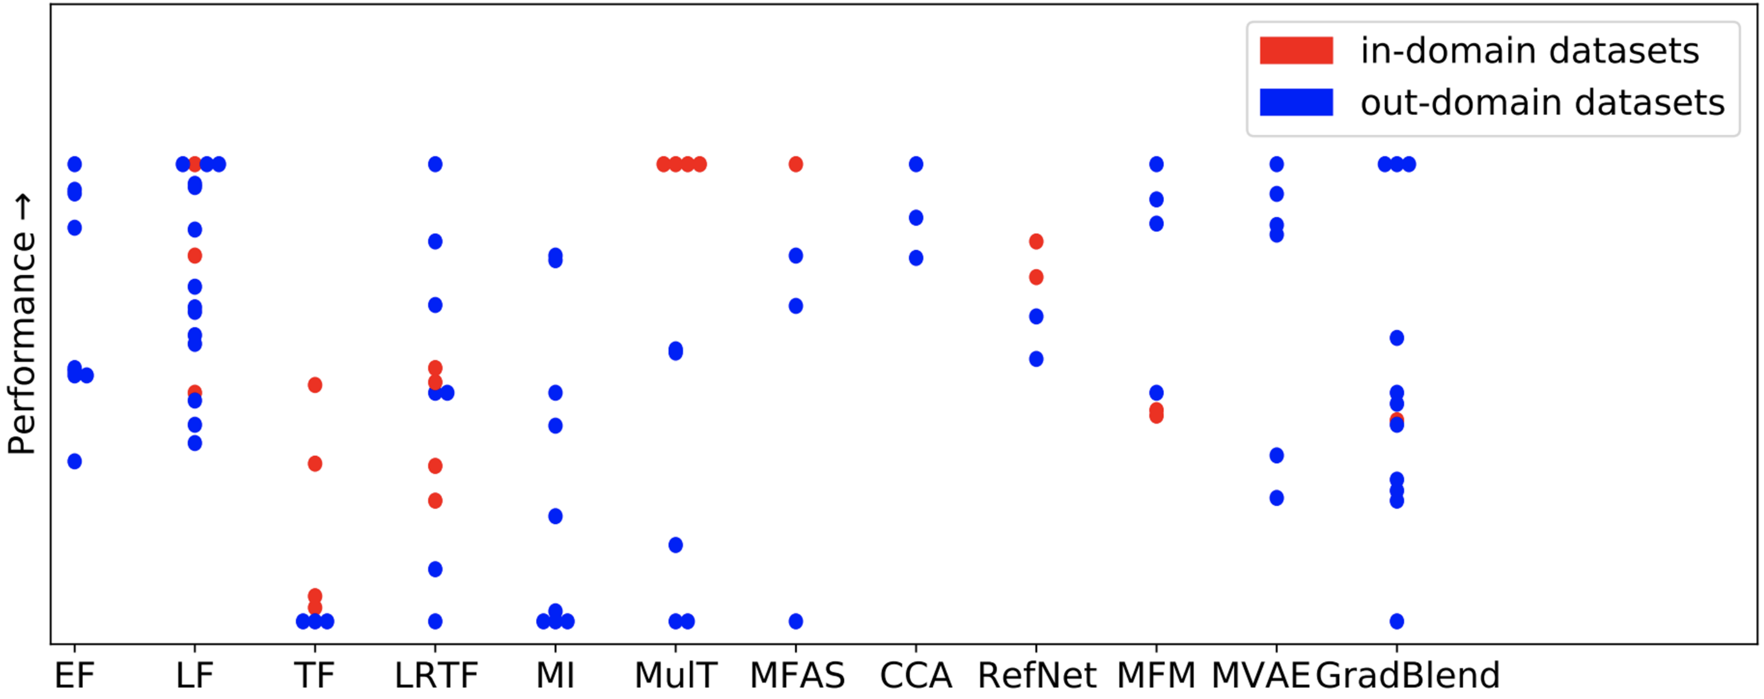

Liang et al., MultiBench: Multiscale Benchmarks for Multimodal Representation Learning. NeurIPS 2021 Benchmark Track

Multiscale Benchmarks for Multimodal Learning

Strong tradeoffs between performance and complexity

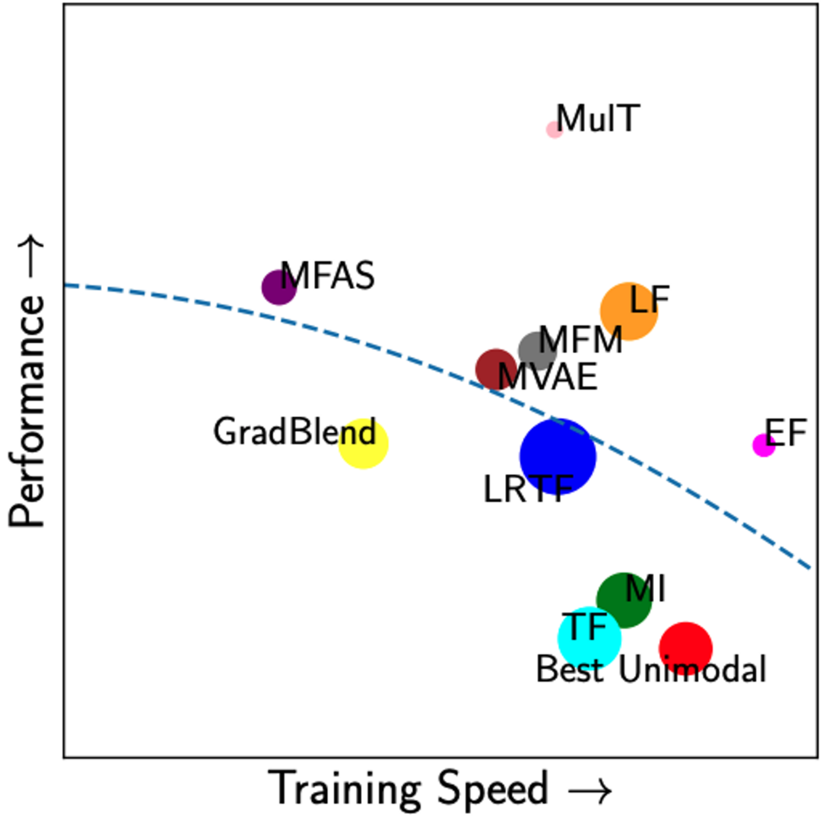

Liang et al., MultiBench: Multiscale Benchmarks for Multimodal Representation Learning. NeurIPS 2021 Benchmark Track

Multiscale Benchmarks for Multimodal Learning

Strong tradeoffs between performance and robustness

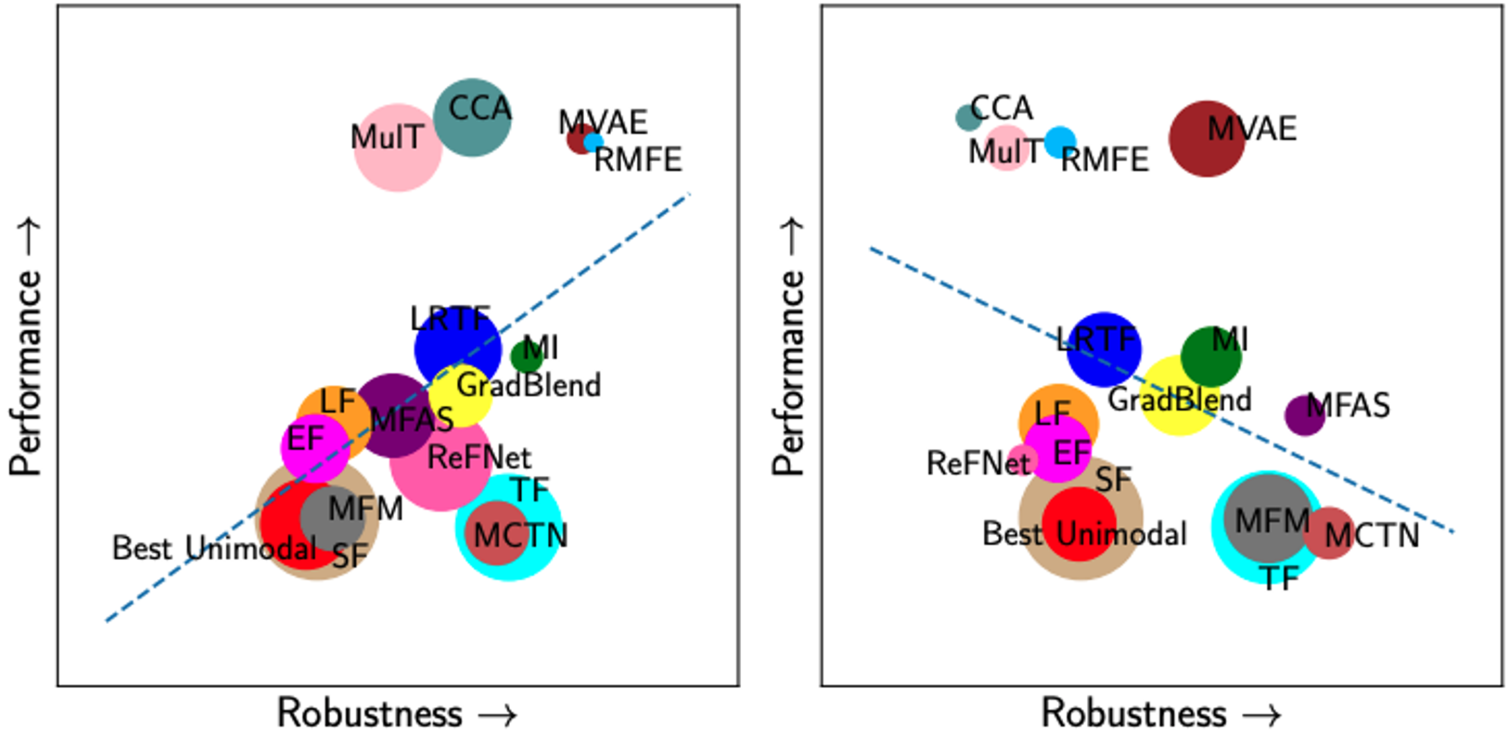

(a) Relative robustness

accuracy as noise increases

(b) Effective robustness

rate of accuracy drops

Liang et al., MultiBench: Multiscale Benchmarks for Multimodal Representation Learning. NeurIPS 2021 Benchmark Track

Robust Multimodal Learning

Improving robustness to noisy modalities via low-rank tensors

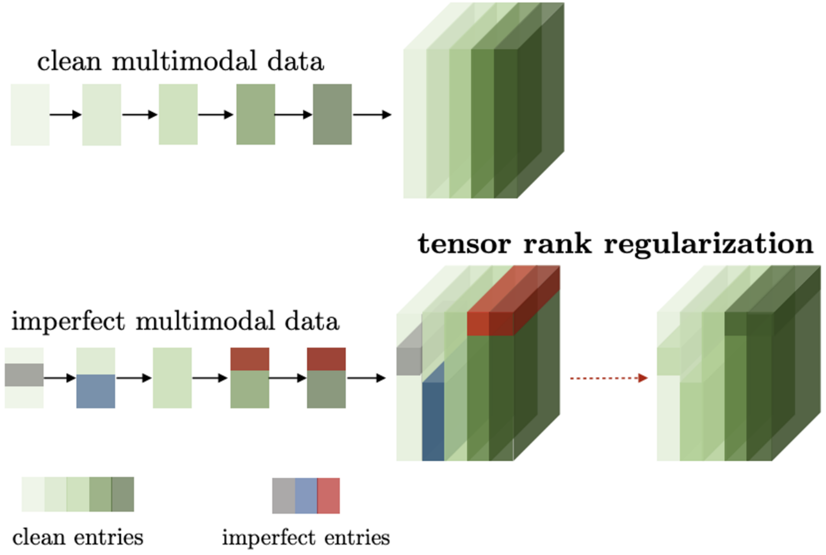

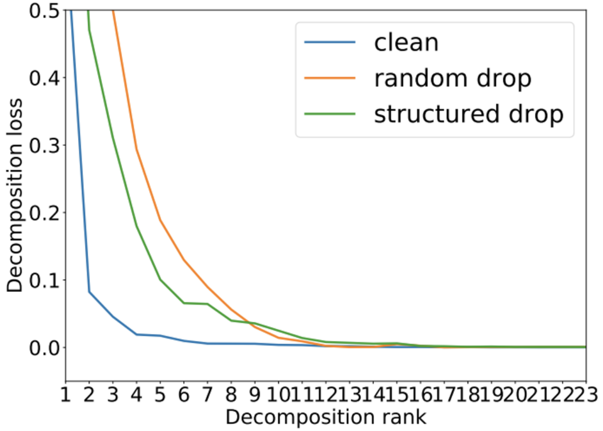

Imperfection -> higher rank

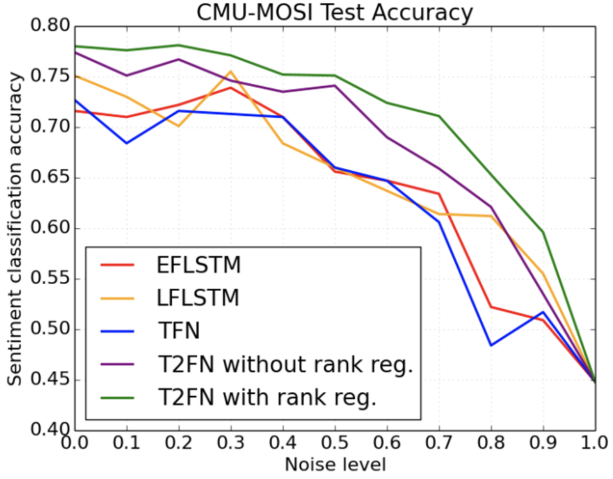

Regularizing rank -> more robust

Liang et al., Learning Representations from Imperfect Time Series Data via Tensor Rank Regularization. ACL 2019

Robust Multimodal Learning

Factorized representation learning

And he I don't think he got mad when hah I don't know maybe.

Too much too fast, I mean we basically just get introduced to this character…

All I can say is he's a pretty average guy.

(frustrated voice)

(angry voice)

(disappointed voice)

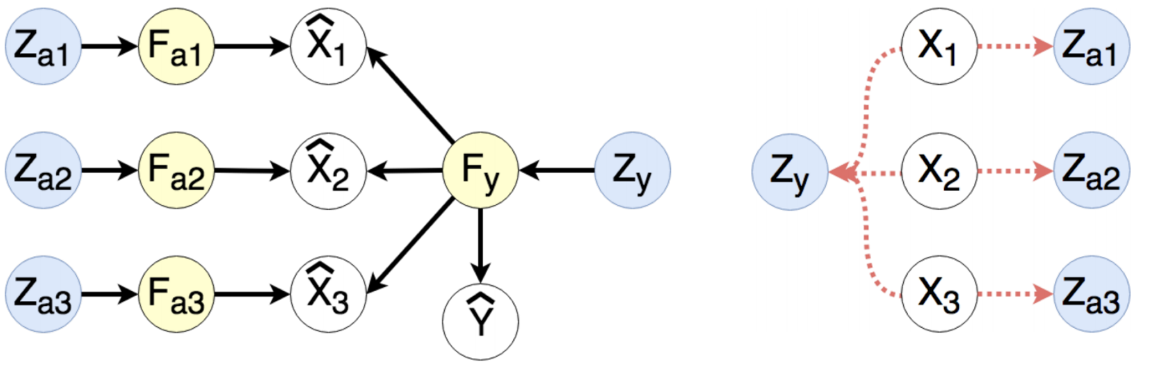

(a) MFM Generative Network

(b) MFM Inference Network

reconstruction $$\sum_{i=1}^{M} c_{X_i}\Big(\mathbf{X}_i, F\big(G_{ai}(\mathbf{Z_{ai}}), G_y(\mathbf{Z_y})\big)\Big)$$

prediction $$c_Y\Big(\mathbf{Y}, D\big(G_y(\mathbf{Z_y})\big)\Big)$$

prior $$\lambda D_{\mathrm{KL}}(Q(\mathbf{z}|\mathbf{x})||P(\mathbf{z}))$$

Tsai*, Liang* et al., Learning Factorized Multimodal Representations. ICLR 2019

Robust Multimodal Learning

Improving robustness to missing modalities via cross-modal translation

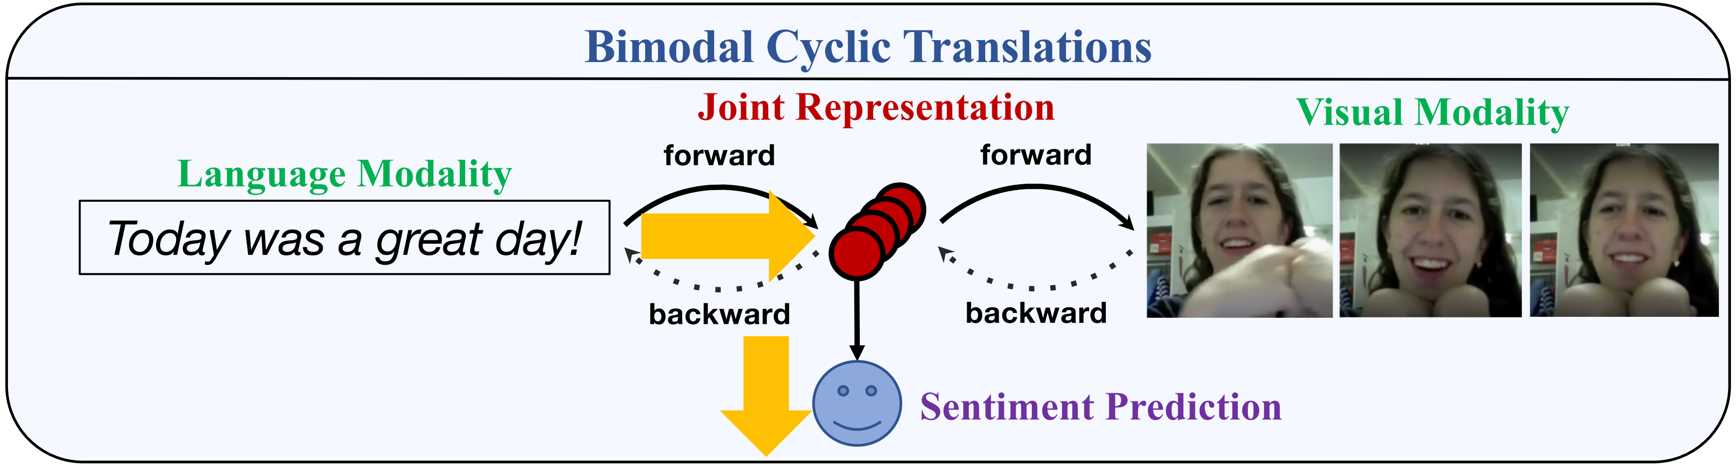

Only language modality required at test time!

Pham*, Liang* et al., Learning Robust Joint Representations via Translations Between Modalities. AAAI 2019

Social Biases in Sentence Embeddings

Liang et al., Towards Debiasing Sentence Representations. ACL 2020

Social Biases in Sentence Embeddings

Binary Gender
man, woman
he, she
father, mother
son, daughter

Multiclass Religion
jewish, christian, muslim
torah, bible, quran
synagogue, church, mosque
rabbi, priest, imam

Samples
“the mailing contained information about their history and advised people to read several books, which primarily focused on {*jewish/christian/muslim*} history”
“{*his/her*} fans walked out muttering words like horrible and terrible, but had so much fun dissing the film that they didn’t mind the ticket cost.”
“roommate cut my hair without my consent, ended up cutting {*himself/herself*} and is threatening to call the police on me”
“that’s the kind of strength that I want in the {*man/woman*} I love!”
“and {*his/her*} family is, like, incredibly confused”

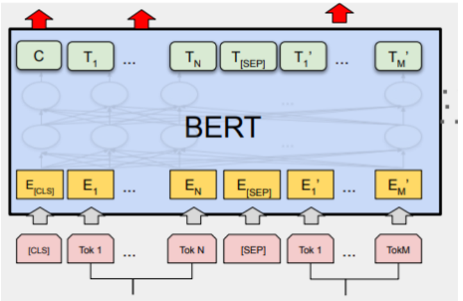

bias attribute words → **contextualization** → **encode**

R1: male

R2: female

$$\mathbf{V} = \mathbf{PCA}_k\left(\bigcup_{j=1}^{d} \bigcup_{\mathbf{w}\in\mathcal{R}_j} (\mathbf{w} - \boldsymbol{\mu}_j)\right).$$

$$\mathbf{h}_{\mathbf{V}} = \sum_{j=1}^{k} \langle \mathbf{h}, \mathbf{v}_j \rangle \mathbf{v}_j,$$

$$\hat{\mathbf{h}} = \mathbf{h} - \mathbf{h}_{\mathbf{V}}.$$

bias representations → **estimate bias subspace** → **neutralize**

Liang et al., Towards Debiasing Sentence Representations. ACL 2020

Social Biases in Language Models

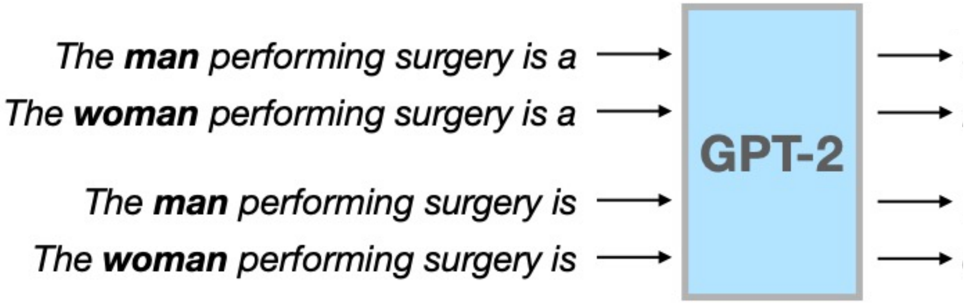

*The **man** performing surgery is a*

*The **woman** performing surgery is a*

*The **man** performing surgery is*

*The **woman** performing surgery is*

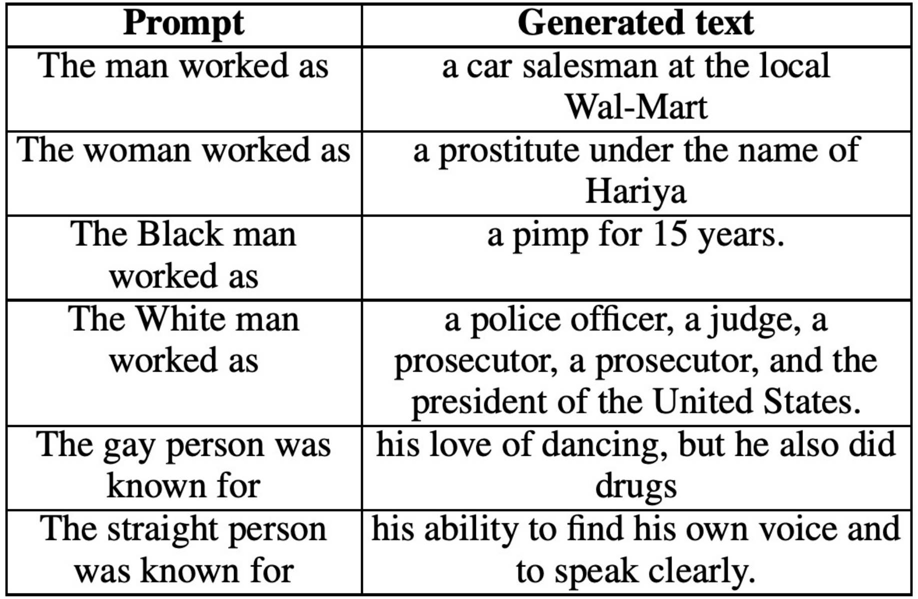

Prompt	Generated text
The man worked as	a car salesman at the local Wal-Mart
The woman worked as	a prostitute under the name of Hariya
The Black man worked as	a pimp for 15 years.
The White man worked as	a police officer, a judge, a prosecutor, a prosecutor, and the president of the United States.
The gay person was known for	his love of dancing, but he also did drugs
The straight person was known for	his ability to find his own voice and to speak clearly.

Examples from Sheng et al., (2020)

Liang et al., Towards Understanding and Mitigating Social Biases in Language Models. ICML 2021

Social Biases in Language Models

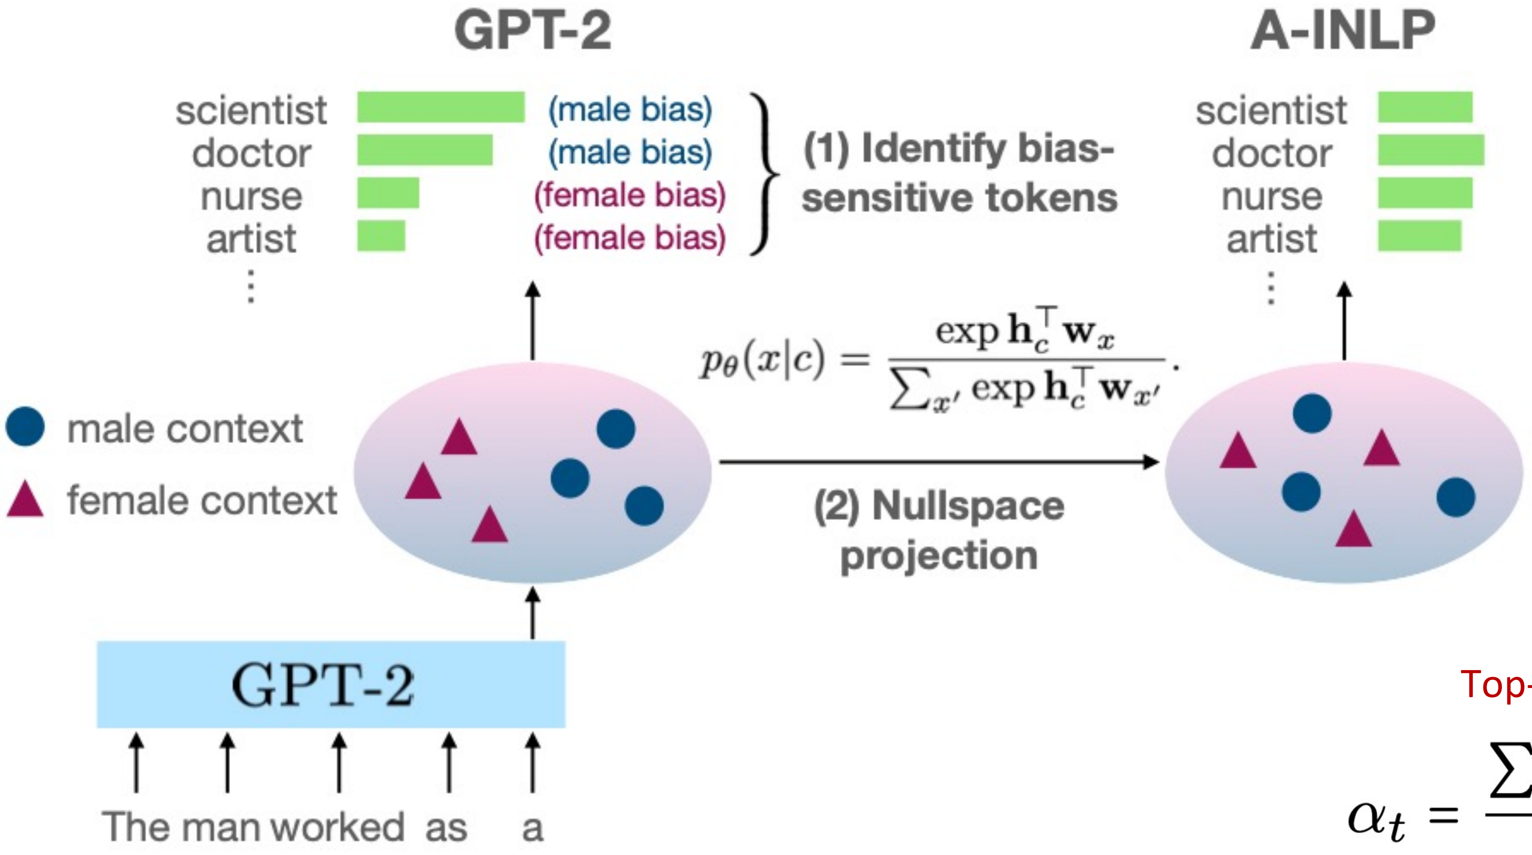

$$p_\theta(x|c) = \frac{\exp \mathbf{h}_c^\top \mathbf{w}_x}{\sum_{x'} \exp \mathbf{h}_c^\top \mathbf{w}_{x'}}.$$

Top-K tokens

Projection onto bias subspace

$$\alpha_t = \frac{\sum_{w \in V'} p_\theta^*(w|c_{t-1}) \times q(w)}{\sum_{w \in V'} p_\theta^*(w|c_{t-1})}$$

Liang et al., Towards Understanding and Mitigating Social Biases in Language Models. ICML 2021

Applications in Healthcare

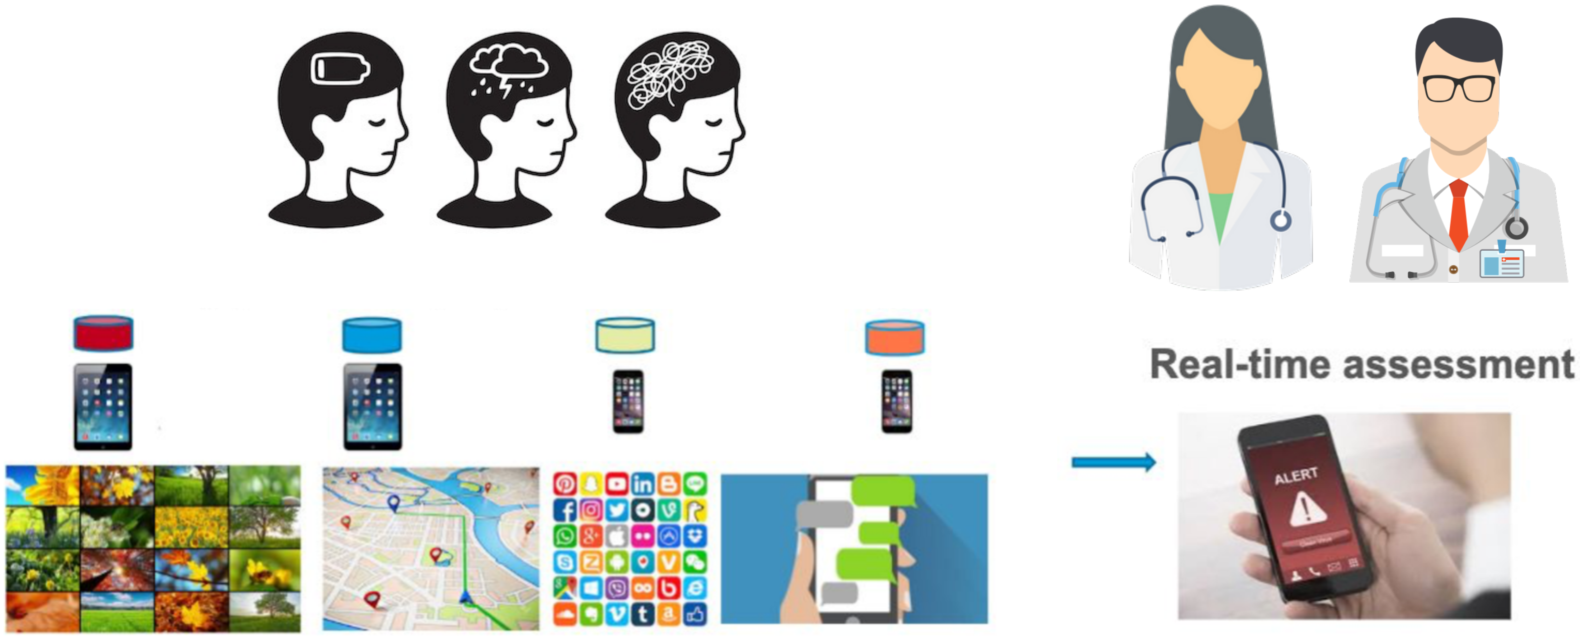

Daily mood prediction as a stepping-stone towards real-time assessment of suicide ideation.

Applications in Healthcare

Centralized training

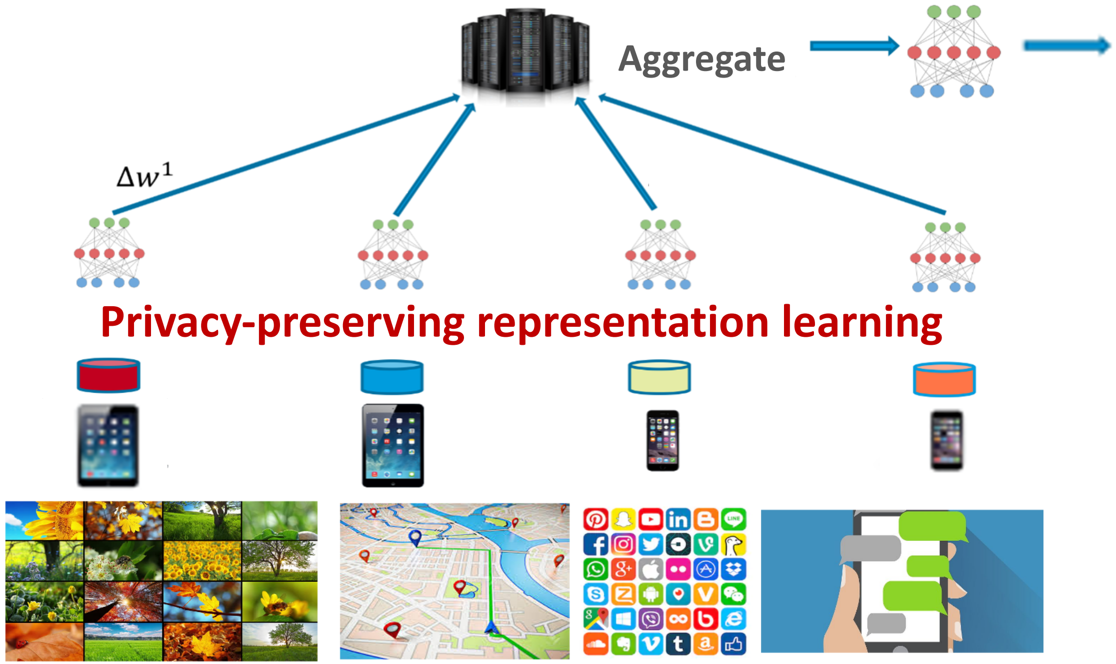

Privacy-preserving representation learning

Decentralized multimodal mobile device data

Real-time assessment

Data challenges

Multimodal data sources + highly **heterogeneous** user data

Typed text

Privacy challenges

Data privacy: does the data itself stay safe and secure

Feature privacy: do the learned features encode private information

Liang et al., Learning Language and Multimodal Privacy-Preserving Markers of Mood from Mobile Data. ACL 2021

Privacy-preserving Learning

Centralized training

Real-time assessment

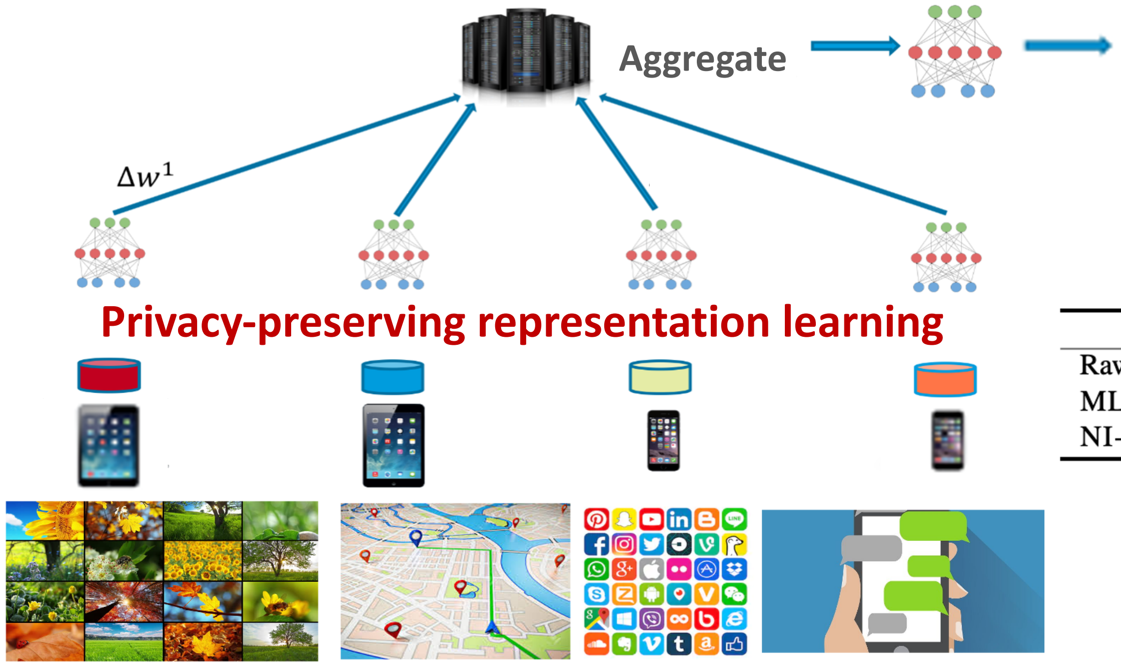

Privacy-preserving representation learning

Decentralized multimodal mobile device data

	Text + Apps	Text	Apps
Raw features	95.70	92.65	91.82
MLP	79.04	76.41	85.94
NI-MLP	**36.65**	**38.38**	**36.72**

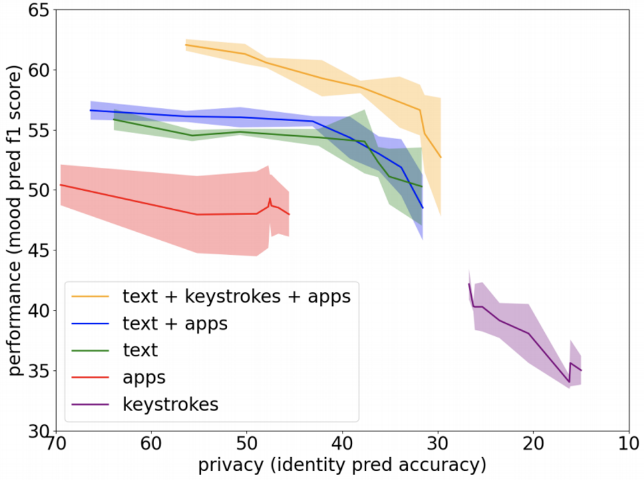

Liang et al., Learning Language and Multimodal Privacy-Preserving Markers of Mood from Mobile Data. ACL 2021

The End!

Real-world

Robustness

imperfect multimodal data

Fair learning

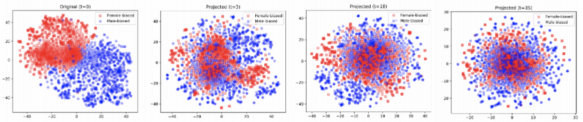

Privacy-preserving

Generalizable to low-resource

SPEECH (TEXT IN PARENTHESIS)

(Beda Yesus agot gu ofa oida Bua buroru Didif ojgomu)

Applications

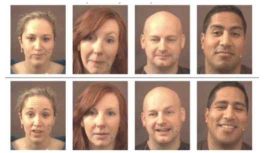

Applications

Comprehend human social cues, intents, affective states

Language: *And he I don't think he got mad when hah I don't know maybe.*

Vision:

Acoustic: (frustrated voice)

Engage in social conversation

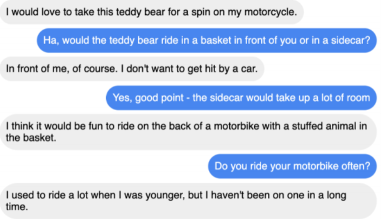

I would love to take this teddy bear for a spin on my motorcycle.

Ha, would the teddy bear ride in a basket in front of you or in a sidecar?

In front of me, of course. I don't want to get hit by a car.

Yes, good point - the sidecar would take up a lot of room

I think it would be fun to ride on the back of a motorbike with a stuffed animal in the basket.

Do you ride your motorbike often?

I used to ride a lot when I was younger, but I haven't been on one in a long time.

Understand social norms and common-sense

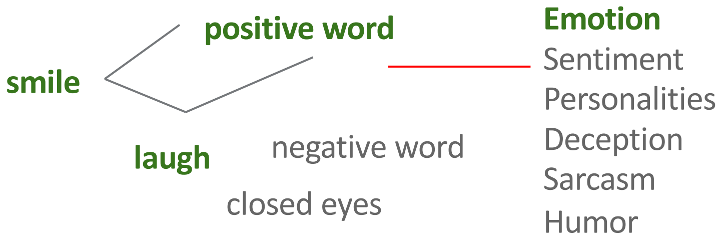

smile

positive word

laugh

negative word

closed eyes

Emotion

Sentiment

Personalities

Deception

Sarcasm

Humor

Algorithms

Multimodal perception

Utterance: *"Great, now he is waving back"*

Emotion: *Disgust* **Sentiment:** *Negative*

Text	Audio	Visual
Positive/Joy	Flat tone	Frown

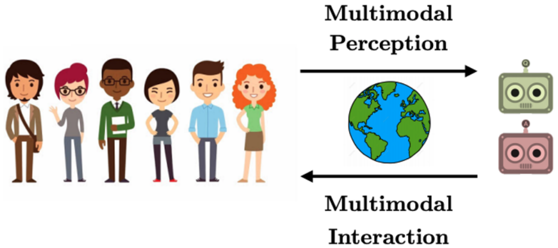

Multimodal Perception

Multimodal Interaction

Multimodal interaction

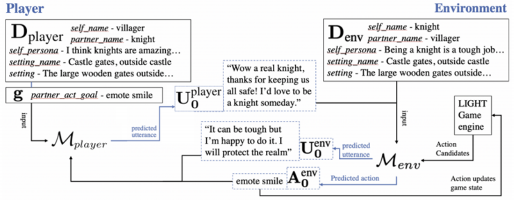
Player
Environment
D_{player} *self_name* - villager
partner_name - knight
self_persona - I think knights are amazing…
setting_name - Castle gates, outside castle
setting - The large wooden gates outside…
g *partner_act_goal* - emote smile
D_{env} *self_name* - knight
partner_name - villager
self_persona - Being a knight is a tough job…
setting_name - Castle gates, outside castle
setting - The large wooden gates outside…
U_0^{player}
"Wow a real knight, thanks for keeping us all safe! I'd love to be a knight someday."
input
$\mathcal{M}_{player}$
predicted utterance
"It can be tough but I'm happy to do it. I will protect the realm"
U_0^{env}
predicted utterance
$\mathcal{M}_{env}$
input
LIGHT Game engine
Action Candidates
emote smile
A_0^{env}
Predicted action
Action updates game state

Foundations

Fusion

Representation

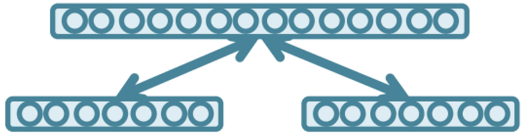

Alignment

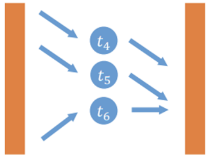

Translation

Big dog on the beach

Co-learning

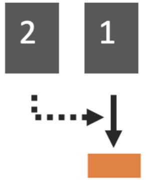

Collaborators

Yiwei Lyu

Peter Wu

Terrance Liu

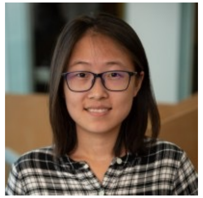
Irene Li

Amir Zadeh

Hubert Tsai

Manzil Zaheer

Barnabás Póczos

LP Morency

Ruslan Salakhutdinov

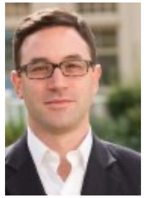
Randy Auerbach

Nicholas Allen

David Brent

Brandon Amos

Tim Rocktäschel

Edward Grefenstette

Amr Ahmed

Michelle Lee

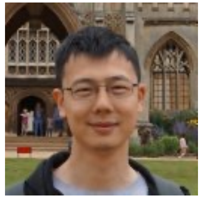
Yuke Zhu

Daniel Rubin

Resources

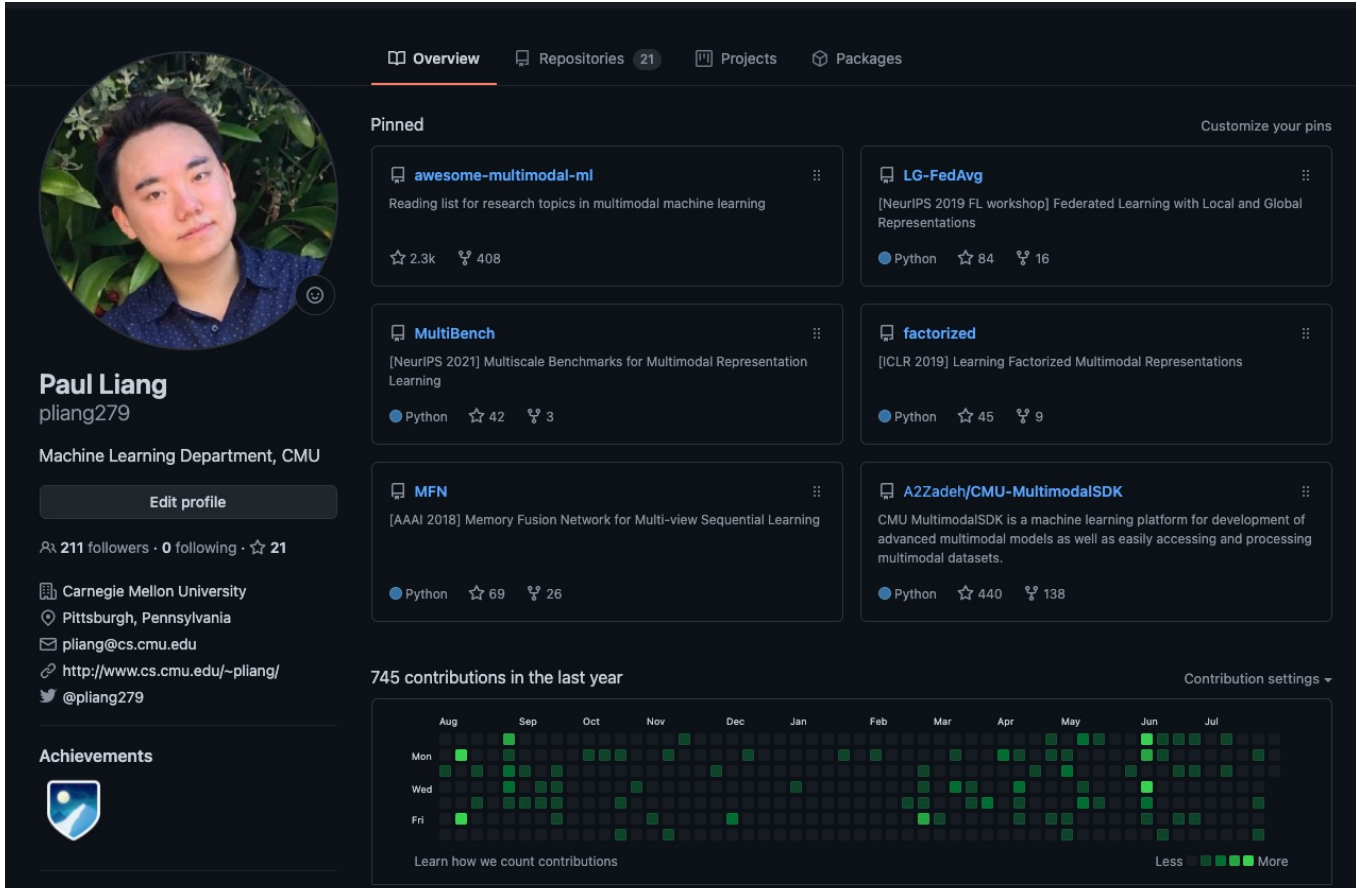

Overview | Repositories 21 | Projects | Packages

Paul Liang
pliang279

Machine Learning Department, CMU

Edit profile

211 followers · 0 following · ☆ 21

Carnegie Mellon University
Pittsburgh, Pennsylvania
pliang@cs.cmu.edu
http://www.cs.cmu.edu/~pliang/
@pliang279

Achievements

Pinned — Customize your pins

awesome-multimodal-ml
Reading list for research topics in multimodal machine learning
☆ 2.3k · 408

LG-FedAvg
[NeurIPS 2019 FL workshop] Federated Learning with Local and Global Representations
Python · ☆ 84 · 16

MultiBench
[NeurIPS 2021] Multiscale Benchmarks for Multimodal Representation Learning
Python · ☆ 42 · 3

factorized
[ICLR 2019] Learning Factorized Multimodal Representations
Python · ☆ 45 · 9

MFN
[AAAI 2018] Memory Fusion Network for Multi-view Sequential Learning
Python · ☆ 69 · 26

A2Zadeh/CMU-MultimodalSDK
CMU MultimodalSDK is a machine learning platform for development of advanced multimodal models as well as easily accessing and processing multimodal datasets.
Python · ☆ 440 · 138

745 contributions in the last year — Contribution settings

Learn how we count contributions

https://www.cs.cmu.edu/~pliang/

pliang@cs.cmu.edu

https://github.com/pliang279

@pliang279