

CMU CoALA Lab

Co- {Augmentation, Learning, & AI}



Overview of Current Research

CMU CoALA Lab

Co- {Augmentation, Learning, & AI}

How can humans + AI systems augment each other's abilities and learn from each other to support more effective, responsible forms of human-AI collaboration?

CMU CoALA Lab

Co- {Augmentation, Learning, & AI}

How can humans + AI systems **augment each other's abilities and **learn from each other** to support more effective, responsible forms of human-AI collaboration?**

1) Studying and supporting human-AI complementarity

2) Participatory design & prototyping for AI

3) Supporting AI fairness in practice

Studying and Supporting Human-AI Complementarity in Practice



theCoALALab.com

How are AI tools augmenting worker practices today?

(especially in frontline and high-touch work)

How can we foster more effective forms of human-AI collaboration?

(both in general and in specific work contexts)

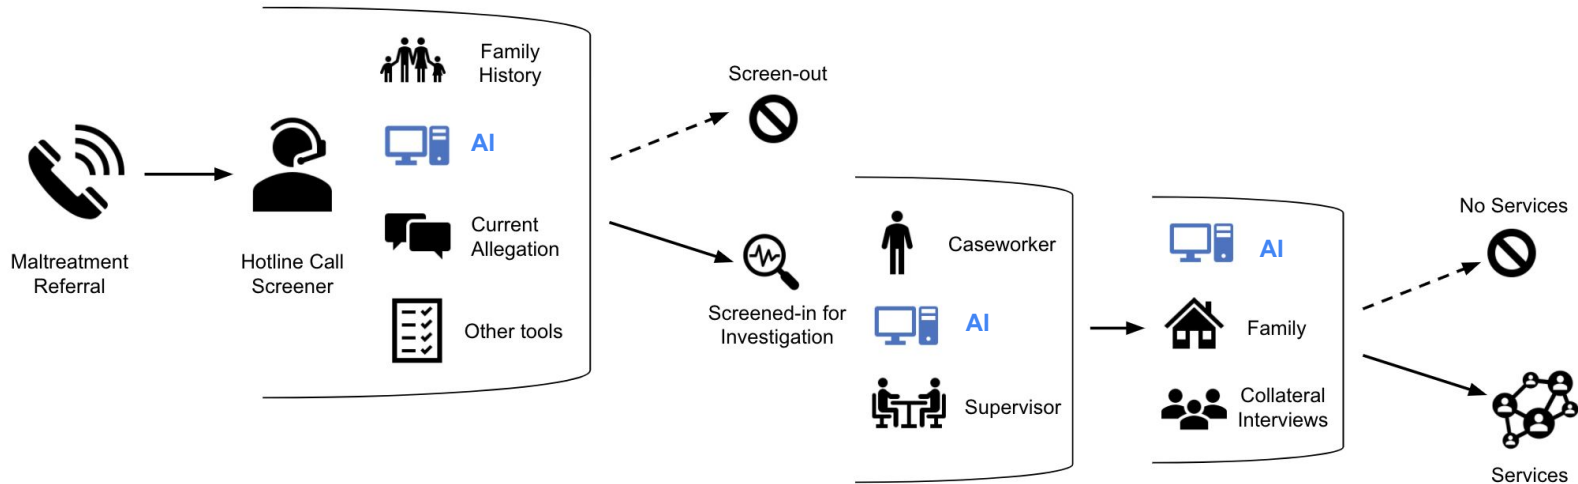
How do we identify opportunities for human-AI collaboration?

Studying and Supporting Human-AI Complementarity in Practice



theCoALALab.com

Case study: AI systems are being used to augment worker decision-making in child welfare today



with Haiyi Zhu, Steven Wu, Alex Chouldechova, Adam Perer, Anna Kawakami, Hao-Fei Cheng, Venkat Sivaraman, Logan Stapleton, & Luke Guerdan

Studying and Supporting Human-AI Complementarity in Practice



theCoALALab.com

**Training for human-AI
collaborative work**

How might we help workers learn how to work with AI tools more effectively?

**Interfaces to improve
human-AI collaboration**

How might we design interfaces that help workers make more effective use of AI in practice?

**Designing AI models to
complement humans**

How might we develop AI-based decision support tools that better complement human judgement?

**Exploring new ways AI
can support workers**

Beyond existing, deployed AI tools, how might new technologies help address worker needs?

with Haiyi Zhu, Steven Wu, Alex Chouldechova, Adam Perer, Anna Kawakami,
Hao-Fei Cheng, Venkat Sivaraman, Logan Stapleton, & Luke Guerdan

Studying and Supporting Human-AI Complementarity in Practice



theCoALALab.com

Training for human-AI collaborative work

How might we help workers learn how to work with AI tools more effectively?


Per RS (school staff at high school), Tracy (age 17) reported that her Mother has taken her house key and looked her outside of the house without a coat until Mother comes home. Child is outside anywhere between 15 minutes and 2 hours, according to Child. This is not the first time this has happened, there have been previous incidents involving the police. Last week, Mother smacked Child in the head. There are no injuries, child is reporting no pain or impairment. Child has asthma. Child is well fed, cared for generally. Mother is employed and there are no known substance abuse, mental health or domestic violence concerns. Mother "may get angry, hit her once or twice and tell her to get out." There was a prior referral seven years ago and a prior referral six months ago. The referral six months prior was regarding parent/child conflict. Mother has sole custody and court records show filings for custody. Per Client View, Mother and child has services through DHS.

Based on the case description, what is your initial assessment of the case?

Safety rating No safety risk Present danger Impending danger

Risk rating Low risk Moderate risk High risk

Screening recommendation Screen in for investigation Screen out for investigation

This is the AFST score:  0 10 15 16 20

Please indicate whether you agree or disagree with the AFST score: Agree Disagree

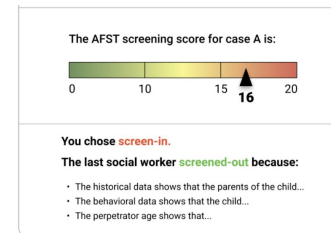
Why do you agree with the AFST score?

The initial assessment aligns with the AFST score.

I changed my assessment after seeing the AFST score because ...

← Previous Case Next Case →

Submit



with Anna Kawakami, Haiyi Zhu, Luke Guerdan, Scott Carter, Matt Lee, Anita Sun, Alison Hu, Matthew Ok

Studying and Supporting Human-AI Complementarity in Practice



theCoALALab.com

**Interfaces to improve
human-AI collaboration**

How might we design interfaces that help workers make more effective use of AI in practice?



with Haiyi Zhu, Maria De-Arteaga, Adam Perer, Anna
Kawakami, Lakshmi Tumati, & Yang Cheng

Studying and Supporting Human-AI Complementarity in Practice



theCoALALab.com

Interfaces to improve human-AI collaboration

How might we design interfaces that help workers make more effective use of AI in practice?



In case you missed it... these factors **played a large role** in determining the AFST score:

Child age: The child is less than 13 years old.

Number of referrals: The father has 4 past referrals.

[Learn more](#)

The AFST did **not** consider the following factors, which you may have taken into account:

Private records: The father has private health records.

Past allegations: 3 of 4 past referrals were for *Alcohol Use/Abuse - Caretaker*.

[Learn more](#)

with Haiyi Zhu, Maria De-Arteaga, Anna Kawakami,
Lakshmi Tumati, & Yang Cheng

Studying and Supporting Human-AI Complementarity in Practice



theCoALALab.com

Interfaces to improve human-AI collaboration

How might we design interfaces that help workers make more effective use of AI in practice?

# Full Baths	2
# Half Baths	0
Heating	Gas Air
Paved Drive	Yes
Rating of Material and Finish (out of 10)	9
Rating of Overall Condition (out of 10)	5
Year Built	2006
Zoning Classification	Residential Low Density


The model does *not* have access to the information highlighted in yellow:

- # Full Baths
- Heating
- Rating of Material and Finish

You may want to take this into consideration as you make your own prediction.

Model Prediction: **\$200,749**

with Maria De-Arteaga, Meghna Sudhakar, Lakshmi Tumati, & Yang Cheng



Supporting Designers in Learning to Co-create with AI

CMU CoALA Lab

Co- {Augmentation, Learning, & AI}

How can humans + AI systems **augment each other's abilities and **learn from each other** to support more effective, responsible forms of human-AI collaboration?**

1) Studying and supporting human-AI complementarity

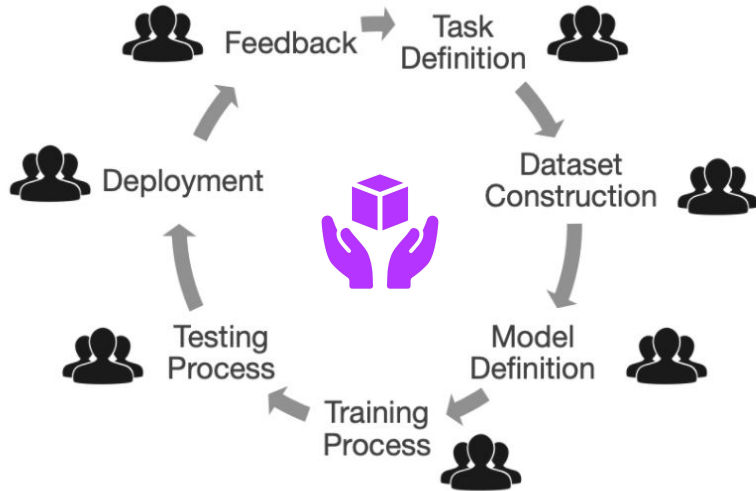
2) Participatory design & prototyping for AI

3) Supporting AI fairness in practice

Participatory Design & Prototyping for Data-Driven AI Systems



theCoALALab.com




with Haiyi Zhu, Motahhare Eslami, Sarah Fox, Steven Wu, Hong Shen, Alexandra Chouldechova, Hong Shen, Wesley Deng, Anna Kawakami, Tzu-Sheng Kuo, Logan Stapleton

Participatory Design & Prototyping for Data-Driven AI Systems



theCoALALab.com



How can we better engage relevant stakeholders in formative stages of the AI development lifecycle?
(e.g., problem formulation, task definition, & dataset construction)

How can we construct meaningful quality measures and KPIs for human-AI collaborative decision-making?

How can we help product teams effectively prototype new AI systems?

Scaffolding AI Fairness Efforts in Practice

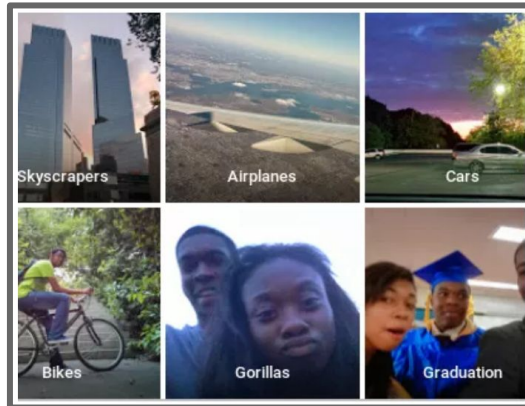


theCoALALab.com

Gender Classifier	Darker Male	Darker Female	Lighter Male	Lighter Female	Largest Gap
Microsoft	94.0%	79.2%	100%	98.3%	20.8%
FACE*	99.3%	65.5%	99.2%	94.0%	33.8%
IBM	88.0%	65.3%	99.7%	92.9%	34.4%

*Facial Recognition Is Accurate,
if You're a White Guy*

By Steve Lohr



with Motahhare Eslami, Jason Hong, Michael Madaio, Adam Perer, Nihar Shah, Hong Shen, Haiyi Zhu, Steven Wu, Alicia DeVos, Wesley Deng, Bill Guo, Nur Yildirim, & Charvi Rastogi



Get in touch!

 theCoALALab.com

 [@CoALALab](https://twitter.com/CoALALab)

 kjholste@cs.cmu.edu
(Ken Holstein)

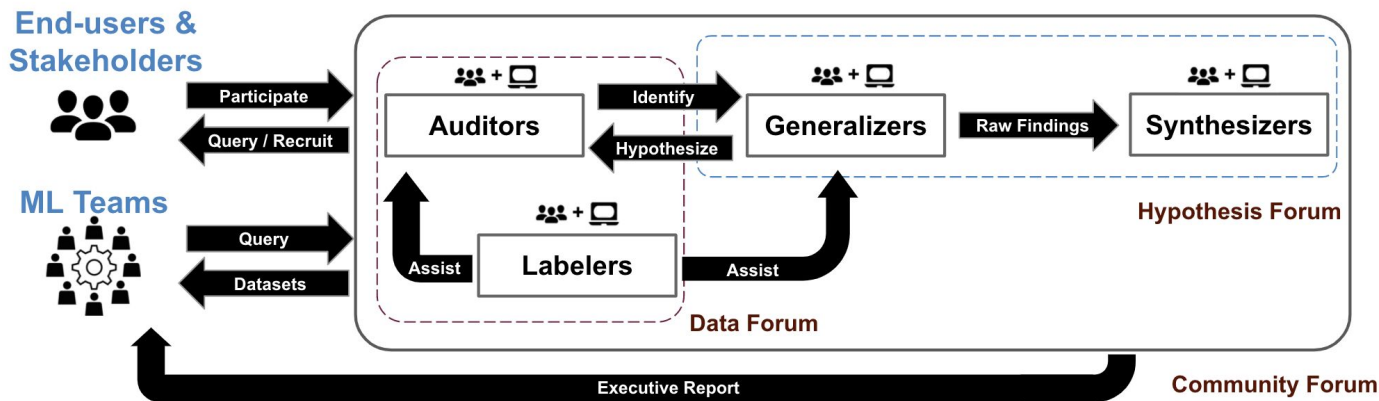
CMU CoALA Lab

Co- {Augmentation, Learning, & AI}

Scaffolding AI Fairness Efforts in Practice



theCoALALab.com



with Motahhare Eslami, Jason Hong, Adam Perer, Nihar Shah, Hong Shen,
Alicia DeVos, Wesley Deng, Bill Guo, & Charvi Rastogi