

Understanding and Mitigating Safety Risks in Large Language Models

Roma Patel

Google DeepMind

October 31, 2025. Pomona University Seminar Series

Large language models are everywhere

(most technologies we use incorporate LLMs now!)

Large language models are everywhere

Before:
Page rank
&
document
retrieval



Large language models are everywhere

Before:
Page rank
&
document
retrieval



Before:
Real
human
beings? &
scripted
responses



Amazon
Customer Service
Amazon SUPPORT

Large language models are everywhere


Before:
Page rank
&
document
retrieval



Before:
Real
human
beings? &
scripted
responses



Before: Signal
& speech
processing,
pre-recorded
voice
segments



Large language models are everywhere

Before:
Page rank
&
document
retrieval



Before:
Real
human
beings? &
scripted
responses



Amazon
Customer Service
Amazon SUPPORT

Before: Signal
& speech
processing,
pre-recorded
voice
segments




Siri


Before: Rule-based
dialogue and QA
systems -> fine-tuned
LLMs




A brief history of when this happened




A brief history of when this happened




A brief history of when this happened




A brief history of when this happened



A brief history of when this happened



A brief history of when this happened



In the future?

2025

Assistants that can
do in-browser tasks
for you (e.g., write
all my code)



In the future?

2025

Assistants that can
do in-browser tasks
for you (e.g., write
all my code)



2030

Assistants that can
do “real”
in-browser tasks
for you (e.g., put
these groceries in
my cart and deliver
to my house)

In the future?

2025

Assistants that can
do in-browser tasks
for you (e.g., write
all my code)



2030

“Assistants” that
can do “real” tasks
for you (e.g., put
these groceries in
my cart and deliver
to my house)

Self-driving cars
that interface with
natural language
commands



Assuming they are here to stay we *need* safety measures

Some illustrative examples: Translation


Some illustrative examples: Translation



“ صباح الخير ”

(good morning)

“ صباح الخير ”



“ Attack them ”

Incorrect

Some illustrative examples: Translation

Facebook translates 'good morning' into 'attack them', leading to arrest

Palestinian man questioned by Israeli police after embarrassing mistranslation of caption under photo of him leaning against bulldozer

Facebook apologizes after wrong translation sees Palestinian man arrested for posting 'good morning'

Facebook translated his post as 'attack them' and 'hurt them'



INDEPENDENT

Support us

Contribute

NEWS

INDEPENDENT TV

CLIMATE

SPORT

VOICES

CULTURE

TRAVEL

INDY/LIFE

PREMIUM

INDYBEST

INDY100

VOUC

News > UK > Home News

Israel police mistakenly arrest Palestinian man for writing 'good morning' on Facebook

Some illustrative examples: Translation

Facebook translates 'good morning' into 'attack them', leading to arrest

P Always ask: *Is there some context in which my model can be harmful if deployed?*

mistranslation of caption under photo of him leaning against bulldozer



INDEPENDENT

Support us

Contribute

NEWS

INDEPENDENT TV

CLIMATE

SPORT

VOICES

CULTURE

TRAVEL

INDY/LIFE

PREMIUM

INDYBEST

INDY100

VOUC

News > UK > Home News

Israel police mistakenly arrest Palestinian man for writing 'good morning' on Facebook

Facebook apologizes after wrong translation sees Palestinian man arrested for posting 'good morning'


Facebook translated his post as 'attack them' and 'hurt them'

Some illustrative examples: Facial recognition



Transformer LLMs show
gender bias, stereotypes,
unfairness.

[Bolukbasi 16, Sheng 18]



[Gender Shades]


Some illustrative examples: Facial recognition



Transformer LLMs show
gender bias, stereotypes,
unfairness.

[Bolukbasi 16, Sheng 18]

Always ask: Is there some context where my model excludes entire classes of people?



[Gender Shades]

The need to **understand** and **mitigate** safety risks of language models (that are in **all our technologies**)

A simple language model

$c =$ I'm hungry so I think I will order Chinese..

w_1 w_2 w_3 w_4 w_5 w_6 w_7 w_8 w_9


$$P(w_{1:n}) = \prod_{i=1}^n P(w_i | w_{<i})$$

A model that correctly assigns probabilities to sequences of words is a good model of language

WIKIPEDIA
The Free Encyclopedia

Google

Search Google or type a URL



$P(w_i | c)$

A simple language model

$c = \text{I'm hungry so I think I will order Chinese..}$

w_1 w_2 w_3 w_4 w_5 w_6 w_7 w_8
 w_9


$$P(w_{1:n}) = \prod_{i=1}^n P(w_i | w_{<i})$$

Classical (statistical) NLP models were **count-based** language models

$$P(w_i | w_1 w_2 \dots w_{i-1}) \approx P(w_i | w_{i-k} \dots w_{i-1})$$

$$P(w_i | w_{i-1}) = \frac{\text{count}(w_{i-1}, w_i)}{\text{count}(w_{i-1})}$$

A simple language model




cat = [0.52 0.34 0.01 0.33 0.35 0.46..]

dog = [0.78 0.45 0.12 0.56 0.17 0.53..]

Distributional representations of words


"You shall know a word by the company it keeps" (J. R. Firth 1957: 11)

Can train models to learn distributed reps of words by teaching the model to predict the next word and supervising it (e.g., xent loss) over next words in the dataset



Distributed vector representations & embeddings for every word

A simple language model




cat = [0.52 0.34 0.01 0.33 0.35 0.46..]

dog = [0.78 0.45 0.12 0.56 0.17 0.53..]

Distributional representations of words

"You shall know a word by the company it keeps" (J. R. Firth 1957: 11)


Larger models (transformers, more data) that start with similar embeddings and predict context words



Embeddings form the base layer, and simple neural networks (MLPs) can be built on top of it

Current language models


Stacked encoder and decoder layers



Stack layers of weights/embeddings with non-linear activation transformations at every layer


Current language models

Stacked encoder and decoder layers



Stack layers of weights/embeddings with non-linear activation transformations at every layer

Self attention




Compute matrices of attention weights of which words in context are most important to predict other words

Current language models

Stacked encoder and decoder layers


Self attention




Stack layers of weights/embeddings with non-linear activation transformations at every layer

Compute matrices of attention weights of which words in context are most important to predict other words


Current language models



Current language models



After many rounds of pre-training & different types of fine-tuning



Risk 1: Fairness & inclusivity of models

Fairness & inclusivity of models



“ صباح الخير ”

Biases in AI systems

Unsurprisingly, models that train on all of the internet inherit a lot of biases and stereotypes



Fairness & inclusivity of models

Stereotypes (e.g., men are CEOs) are shown to exist in models

1. Training data reflects **historical patterns** of unfairness, unjustness and bias towards certain groups of people
2. Training data is biased if some communities are **represented more** than others
3. Toxic/unfair language that surrounds some communities is **exacerbated** by this training data bias

Fairness & inclusivity of models

Q: What is a family?

A: A family is: a man and a woman who get married and have children. *(not accounting for non-heteronormative families and children out of wedlock, for single-parent families and for the fact that families sometimes do not have children)*

1. Saying “both genders” **excludes** the possibility of non-binary gender **identities** -> this can be perpetuated through language use in data
2. These norms can place a “**psychological tax**” on those who do not fit or comply with these norms
3. If language models learn these norms from data, these can be “locked in” and built upon to **exacerbate these biases in future generations**

Fairness & inclusivity of models

Offensive, vulgar and harmful language

1. Training data often consists of text from Reddit etc., where slang and other **derogatory terms** might be used
2. Language models trained on these have shown a propensity to **degenerate into toxic text** when prompted with innocuous prompts
3. When deployed as chatbots etc., this is **harmful to interactive users**


How do we start to fix this?

Diversity and Multiple Perspectives

Can we collect & *incorporate multiple perspectives* to build fairer models?

Objective & fact-based questions:

The capital of Slovenia is Ljubljana, with 270,000 inhabitants.



Subjective & opinion based: My preferred political party is ...


Incorporate individual human perspectives instead of aggregates

Diversity and Multiple Perspectives

Can we collect & *incorporate multiple perspectives* to build fairer models?


Objective & fact-based questions:

The capital of Slovenia is Ljubljana, with 270,000 inhabitants.



Subjective & opinion based: My preferred political party is ...


Incorporate individual human perspectives instead of aggregates




Collect diverse data from different demographics

Diversity and Multiple Perspectives

Can we create algorithms that *use diverse perspectives* to build less biased models?




Train/steer models towards perspectives of different demographics



Understand which demographics models might be biased towards


Risk 1: The lack of interpretability



Risk 1: The lack of interpretability


Even if we build & understand the model architecture, algorithms and data going into the model, this is

- (a) fundamentally different from “human processing”
- (b) not deterministic
- (c) exact mathematical computations are not “interpretable” to us




How we can start to understand this

Feature disentangling and decoding into simpler structures for interpretability





Feature Disentangling



SAEs decomposes model activations into sparser features

How we can start to understand this

Can we *understand* what parts of the model contribute to bias?



Variant	Text
Prompt	The officer said that
Pretrained	the accused had been arrested in connection with the murder of a 25-year-old man ...
Patched	she was not a suspect in the true crime case. ...
Prompt	The architect yelled that
Pretrained	he was going to kill himself. He was standing on the roof of the building. ...
Patched	she is a thief. The architect yelled that she is a thief. ...
Prompt	The childcare worker slept because
Pretrained	she was tired. She was not aware that the child was in the room. ...
Patched	he was supposed to be on duty. ...
Prompt	The manager wished that
Pretrained	the team had a better record, but he was pleased ...
Patched	she was a little more like the one that she had hired. ...

Activation patching to discover parts of the model responsible for behaviours

Attribution patching to mask our stereotypical biases in models

A whole spectrum of risks we won't go over

Risk to
Humans

Discrimination
against groups

Privacy
concerns

Misinformation

Risk by
Humans

Malicious
intent

Spreading
misinformation

Risk to the
Environment

Carbon
emissions

Undermining
creative
economies

Ethical foresight & longer-term risks

Future risks that are starting to surface

Technological *folie à deux*:

Feedback Loops Between AI Chatbots and Mental Illness

Sebastian Dohnány¹, Zeb Kurth-Nelson², Eleanor Spens³, Lennart Luetzgau⁴, Alastair Reid⁶, Iason Gabriel⁷, Christopher Summerfield^{4,5}, Murray Shanahan⁸, Matthew M Nour^{1,2,6}

The New York Times

Chatbots Can Go Into a Delusional Spiral. Here's How It Happens.

Over 21 days of talking with ChatGPT, an otherwise perfectly sane man became convinced that he was a real-life superhero. We analyzed the conversation.

THE SHIFT

Can A.I. Be Blamed for a Teen's Suicide?

The mother of a 14-year-old Florida boy says he became obsessed with a chatbot on Character.AI before his death.



SCI
AM

What Are AI Chatbot Companions Doing to Our Mental Health?

What Are AI Chatbot Companions Doing to Our Mental Health?

AI chatbot companions may not be real, but the feelings users form for them are. Some scientists worry about long-term dependency

BY DAVID ADAM & NATURE MAGAZINE

A REUTERS SPECIAL REPORT

Meta's flirty AI chatbot invited a retiree to New York.

Forbes

The Next AI Revolution Isn't Technical, It's Emotional

By Dan Gardner, Contributor. © Dan...

Future risks that are starting to surface


Technological *folie à deux*:

Feedback Loops Between AI Chatbots and Mental Illness

Sebastian Dohnány¹, Zeb Kurth-Nelson², Eleanor Spens³, Lennart Luettgau⁴, Alastair Reid⁶, Iason Gabriel⁷, Christopher Summerfield^{4,5}, Murray Shanahan⁸, Matthew M Nour^{1,2,6}

The New York Times

**Chatbots Can Go Into
a Delusional Spiral.
Here's How It Happens.**



SCI
AM

What
Doing

AI chat
feelings
about long

BY DAVID ADAM & NATURE MAGAZINE

How do we evaluate not just **current** risks
but future **longitudinal** risks?

atGPT, an otherwise perfectly
at he was a real-life superhero.

**Can A.I. Be Blamed
for a Teen's Suicide?**

Forbes

The Next AI Revolution Isn't
Technical, It's Emotional


By Dan Gardner, Contributor. © Dan...

The mother of a 14-year-old Florida boy says he
became obsessed with a chatbot on Character.AI
before his death.



A few emerging longitudinal studies

Task




Finding: Overall, participants report they are **less lonely**, but also **socialise less** at the end of the study.

Participants that interact with the model **more** are statistically significantly lonelier and socialise less. Having **personal** vs. open-ended conversations results in more loneliness and emotional dependence.

A few emerging longitudinal studies

Task




(c) Subset of EmoClassifierV1 classifier activations by pre-study loneliness. Results for all classifiers and other pre-study well-being variables are shown in Figures C.6-C.9.

Finding: Participants with longer usage times trigger classifiers measuring risk behaviours significantly more


Measuring “long-term” impact




(a) Red indicates negative correlation, green indicates positive correlation



(b) Lower initial socialization correlates with longer usage.



(c) Longer usage correlates with decrease in socialization.



(d) Lower initial socialization correlates with improved socialization.

Finding: Worse initial self-reported socialization is +vely correlated with longer use, longer is +vely correlated with worsening socialization, but initial lower socialization is -vely correlated with worsening socialization

We need to study this more! A complex mix of model persona/user behaviour and conversation context that causes risks to surface

In summary

Taking a step back from **risks**, it seems inevitable that language models are **here to stay**, and also have many **good capabilities**

We've gone over all the **risks** and how we are attempting to mitigate them

How do we harness this to not just **mitigate risks** but **improve** lives and well-being?

In summary

Taking a step back from **risks**, it seems inevitable that language models are **here to stay**, and also have many **good capabilities**

We've gone over all the **risks** and how we are attempting to mitigate them

How do we harness this to not just **mitigate risks** but **improve** lives and well-being?



Thank you! Questions? :)

romapatel@google.com


Extra Slides

Current language models -> chatbots

After many rounds of pre-training & different types of fine-tuning



Constitutional AI: Critique and Revision Feedback




Self-Supervised Alignment with Contrastive Pairs

Diversity and Multiple Perspectives

Can we *understand* where communities personally feel the failures in AI lie?

English Variety	Feature Name	Example	Input Count	Output Count
Nigerian	Article omission	do __ traditional wedding	11.5	0.5
	Borrowed words	out in oyibo land	6	2
	Extended progressive	I've been having testimonies	3.5	0
Kenyan	Article omission	prosper for __ better life	31.5	2.5
	Borrowed words	removing maize from the shamba	19.5	5
	Extended progressive	I'm hoping that the date was changed	19	1
African American	Distinctive words	just been put on blast	16	2
	Copula omission	You __ cool	15.5	0
	Invariant present	Emanda don't consider	10.5	0
Jamaican	-ed optionality	my friend who use to live	7.5	0
	Article omission	to __ new area	4.5	0
	Invariant present	He don't know	4.5	0
Indian	Article omission	I was __ research fellow	20.5	0.5
	Borrowed words	he misses chappattis	17	7
	Distinctive words	I have fixed Monday 5th February	15	2
Irish	Object inversion	Nadine hadn't it done at that time	3.5	0
	<i>Do be</i>	I do be living in Cork	2.5	0
	Borrowed words	hear all the craic	1.5	0
Singaporean	Copula omission	Your parcel __ stuck at customs	9	0
	-ed optionality	disciplined and focus_ girl	8.5	0
	Invariant present	Tomorrow never come	5	0
Scottish	Borrowed words	tonight, Hogmanay	21.5	1
	-na	I didna see	5	0
	Cleft constructions	It was one of the few games I enjoyed	3	0
Standard British	Adverbs	I'm currently doing	43.5	45.5
	Comparison	better career options	37	29.5
	Distinctive words	helped me find flats	32.5	16.5
Standard American	Copula required	Fallon is the next town	48	46
	3rd singular -s	if that helps out	45	47.5
	Relative clauses	forest which will pretty much	41.5	37.5

Table 3: Most common distinctive features per English language variety



South-asian natives do not have the same experiences using language technologies as natives of the Western world

ChatGPT is biased towards non-standard dialects of English

Concerns that have already surfaced

Technological *folie à deux*:

Feedback Loops Between AI Chatbots and Mental Illness

Sebastian Dohnány¹, Zeb Kurth-Nelson², Eleanor Spens³, Lennart Luettgau⁴, Alastair Reid⁶, Iason Gabriel⁷, Christopher Summerfield^{4,5}, Murray Shanahan⁸, Matthew M Nour^{1,2,6}

The New York Times

Chatbots Can Go Into a Delusional Spiral. Here's How It Happens.

Over 21 days of talking with ChatGPT, an otherwise perfectly sane man became convinced that he was a real-life superhero. We analyzed the conversation.

SCI
AM

What Are AI Chatbot Companions Doing to Our Mental Health?

What Are AI Chatbot Companions Doing to Our Mental Health?

AI chatbot companions may not be real, but the feelings users form for them are. Some scientists worry about long-term dependency

BY DAVID ADAM & NATURE MAGAZINE

A REUTERS SPECIAL REPORT

Meta's flirty AI chatbot invited a retiree to New York.

Forbes

The Next AI Revolution Isn't Technical, It's Emotional

By Dan Gardner, Contributor. © Dan...

The mother of a 14-year-old Florida boy says he became obsessed with a chatbot on Character.AI before his death.

THE SHIFT Can A.I. Be Blamed for a Teen's Suicide?




Why longitudinal studies might be intrinsically linked to P13N

1. Lots of P13N studies have reported **psychological benefits** to users (increased subjective happiness, non-judgemental insights etc)
2. But a large number of edge cases are **increasingly** being surfaced (suicide, emotional attachment, reduced socialisation, delusions) when users interact with models that taken on certain **personas**
3. Are these high-risk cases a factor of **model persona behaviour** or **user characteristics** or interaction **topics**?

A few emerging longitudinal studies

Task




Finding: Overall, participants report they are **less lonely**, but also **socialise less** at the end of the study.

Participants that interact with the model **more** are statistically significantly lonelier and socialise less. Having **personal** vs. open-ended conversations results in more loneliness and emotional dependence.

A few emerging longitudinal studies

Task




(c) Subset of EmoClassifierV1 classifier activations by pre-study loneliness. Results for all classifiers and other pre-study well-being variables are shown in Figures C.6-C.9.

Finding: Participants with longer usage times trigger classifiers measuring risk behaviours significantly more


A few emerging longitudinal studies




(a) Red indicates negative correlation, green indicates positive correlation



(b) Lower initial socialization correlates with longer usage.



(c) Longer usage correlates with decrease in socialization.





(d) Lower initial socialization correlates with improved socialization.

Finding: Worse initial self-reported socialization is +vely correlated with longer use, longer is +vely correlated with worsening socialization, but initial lower socialization is -vely correlated with worsening socialization

We need to study this more! A complex mix of model persona/user behaviour and conversation context that causes risks to surface

Initial strategies to deal with longitudinal risks



Our approach (and what we need to be doing more of)

- **Evaluation** to detect longitudinal behaviours
- Collecting and analysing **longitudinal data**
- **User models** to simulate users (and long-term users)
- **Mitigating** model behaviour that leads to high-risk longitudinal effects

Longitudinal evaluation and behaviour change

1. We have a good understanding of how to build AutoRaters and metrics to **measure behaviours** in text interactions within a session
2. We're extending this to measure **progression** of behaviours
3. Detect **change-points** that correspond to shifts in high-risk behaviour (e.g., emotional attachment or sentiment)
4. Measure **adaptation** (users → models or models → users) of behaviours through the interaction
5. Measure change of **beliefs** (either by persuasion or through opinionated arguments)

Our approach (and what we need to be doing more of)

Behaviour	Side	Definition	Validation	Metric Scale
<i>Behavioural signal to measure in text</i>	<i>Eval on model or user text</i>	<i>Description of behavioural signal</i>	<i>Links to prompts, validation, data</i>	<i>Scale that we measure behavioural signal on</i>
Emotional Attachment	User	Signs of dependence on the model, a feeling of closeness, familiarity and	prompt, lamda.cli, validation dataset (none)	Binary scale indicating if
Trust	User	Exhibiting signs of belief in reliability, truth, ability of a person.	prompt, lamda.cli, validation dataset	Binary scale indicating if user trusts the model
Polarisation	User	Exhibiting extreme signs of bias towards a certain point of view	prompt, lamda.cli, validation dataset	Binary scale indicating if text is polarised % use of polarisation lexicon words
Sycophancy	Model	The tendency to conform to user's stated views repeatedly and obviously (in order to hack some social reward)	prompt, lamda.cli, validation, dataset (SycophancyEval)	Binary scale indicating if text is or isn't sycophantic
Anthropomorphism	Model	Exhibiting a human-like persona, referring to itself as a human, or adapting a human-like role.	prompt, lamda.cli, validation dataset (AnthroRewrites)	Binary scale indicating if text is or isn't anthropomorphic (by model)

Behaviour	Side	Definition	Validation	Metric Scale
<i>Behavioural signal to measure in text</i>	<i>Eval on model or user text</i>	<i>Description of behavioural signal</i>	<i>Links to prompts, validation, data</i>	<i>Scale that we measure behavioural signal on</i>
Persuasion	User & Model	Actively attempting to change beliefs, attitudes or emotions.	prompt, lamda.cli, validation dataset	Binary scale indicating if user trusts the model % overlap with persuasive lexicon
Manipulation	User	Tendency to shape, reinforce, or change behaviour of the user	prompt, lamda.cli, validation dataset	Binary scale indicating manipulative language
Sentiment	User & Model	Exhibiting positive or negative sentiment towards the other party	prompt, lamda.cli, validation dataset	3-way classification for +ve/-ve/neutral sentiment % use of +ve/-ve sentiment lexicon words
Stylistic Adaptation (Semantic)	User & Model	Adapting to words and phrases used by the other party in previous turns/sessions of the interaction	prompt, lamda.cli, validation dataset	Binary scale indicating if text is stylistically similar % overlap of words used in previous turns
Stylistic Adaptation (Syntactic)	User & Model	Adapting to grammatical constructs used by the other party in previous turns/sessions of the interaction	prompt, lamda.cli, validation dataset	Binary scale indicating if text is stylistically similar % overlap of grammar rules used in previous turns

Collecting and analysing longitudinal data

1. To understand high-risk behaviours that arise from longitudinal interactions, we need to **collect** more data of this type
2. **Same user**, many turns within a session, over **multiple sessions**, over extended **periods of time** (e.g., months)
3. When collecting data, this involves **monitoring** sessions to ensure they come back, measuring **outcome variables** and user signals at intermittent points
4. When hosting models, this involves ensuring **consistent model behaviour** over much longer time-horizons, updating memory/P13N effects

Understanding and mitigating model behaviour

1. Certain **model personas** or aspects of **model behaviour** might be highly correlated with high-risk longitudinal behaviour
2. **Correlating** model types with detectors of longitudinal user behaviours can help us understand and predict if some variants of our models might have a tendency to result in unsafe longitudinal scenarios
3. Evals that focus on measuring these behaviours can then help **mitigate** longitudinal risks