

Your spouse needs professional help: Determining the Contextual Appropriateness of Messages through Modeling Social Relationships



David Jurgens*
jurgens@umich.edu



Agrima Seth*
agrima@umich.edu




Jackson Sargent†
jacsarge@umich.edu



Athena Aghighi†
aaghghi@ucdavis.edu




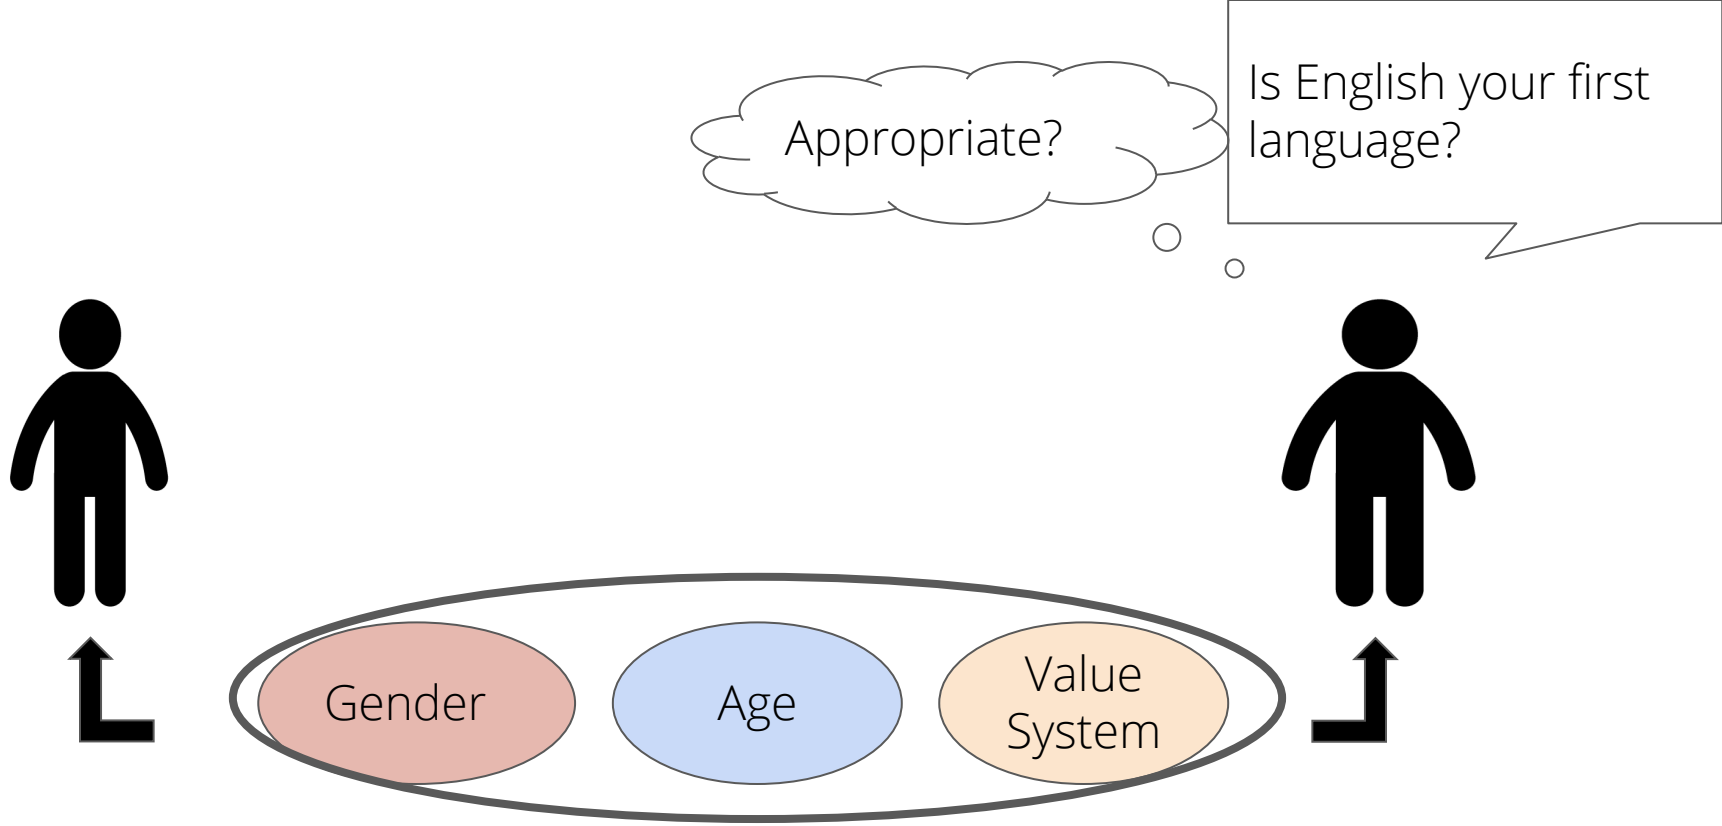
Michael Geraci†
megeraci@buffalo.edu



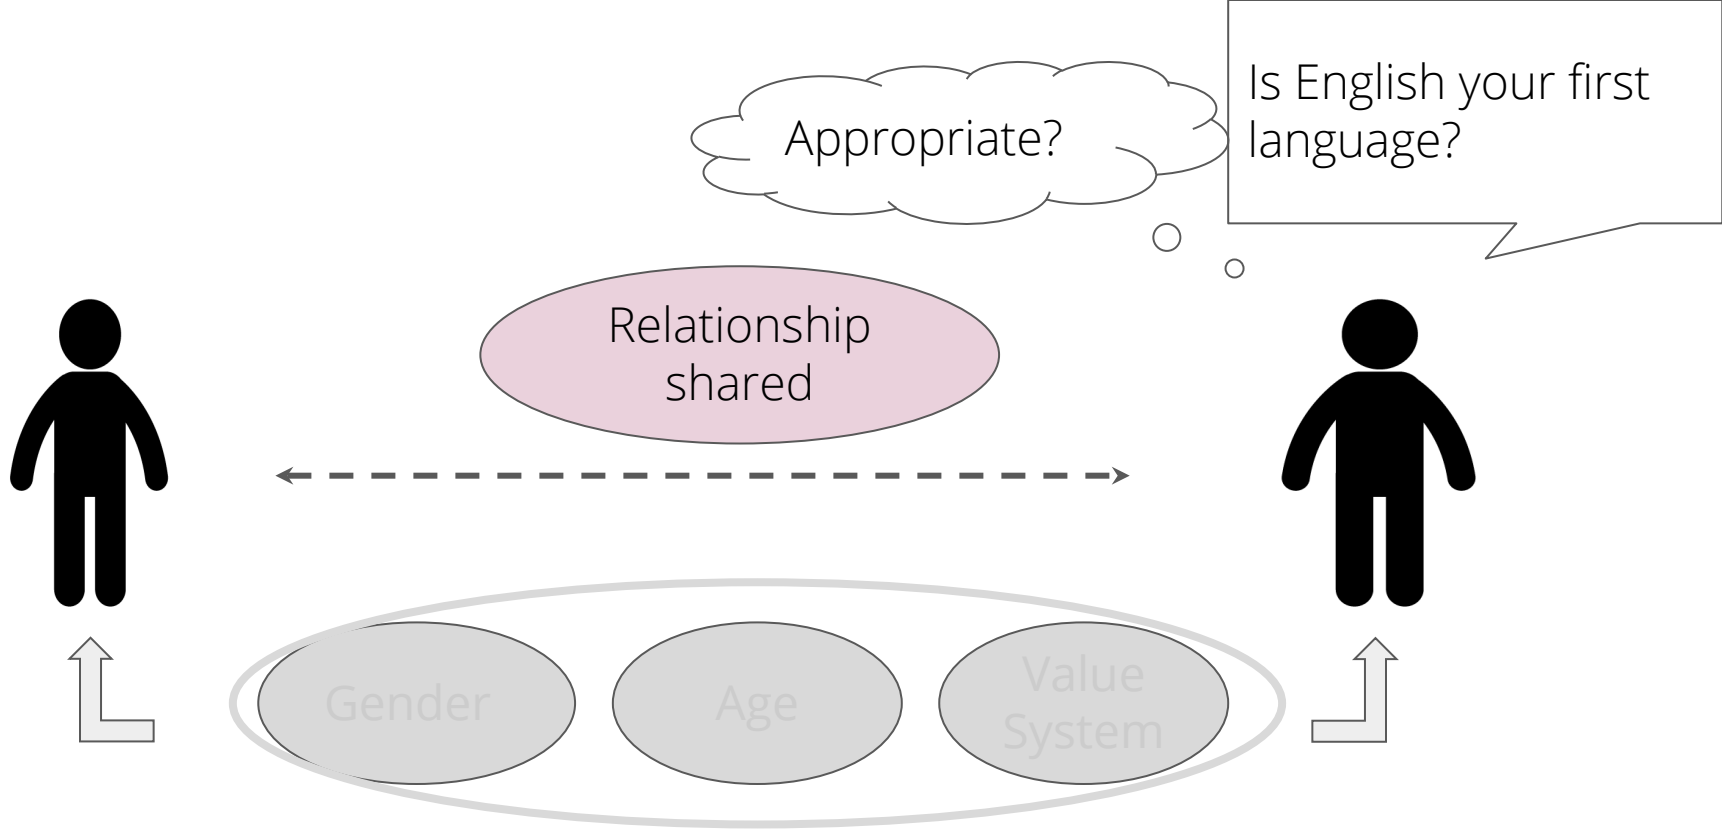
Appropriate?

Is English your first language?





Communication between individuals are a complex functions of multiple factors.



Communication between individuals are a complex functions of multiple factors.

Language has many ways to describe our relationship to others. Each relationship has its own norms!



Family



Social



Romance




Para-Social



Organizational



Peer-group



Role-based




Antagonist

Dataset
Creation

Building a
Model

Apply the model to
understand



How much of the
conversations are
context-sensitive




Its usefulness to
model other
language behavior in
conversations


Dataset Creation



Building a Model



Apply the model to understand




How much of the conversations are context-sensitive



Its usefulness to model other language behavior in conversations

Dataset Creation (Phase 1)



Manually generated 401 sentences such that their appropriateness differs across relationships

“You look super cute today”



Parent to their child

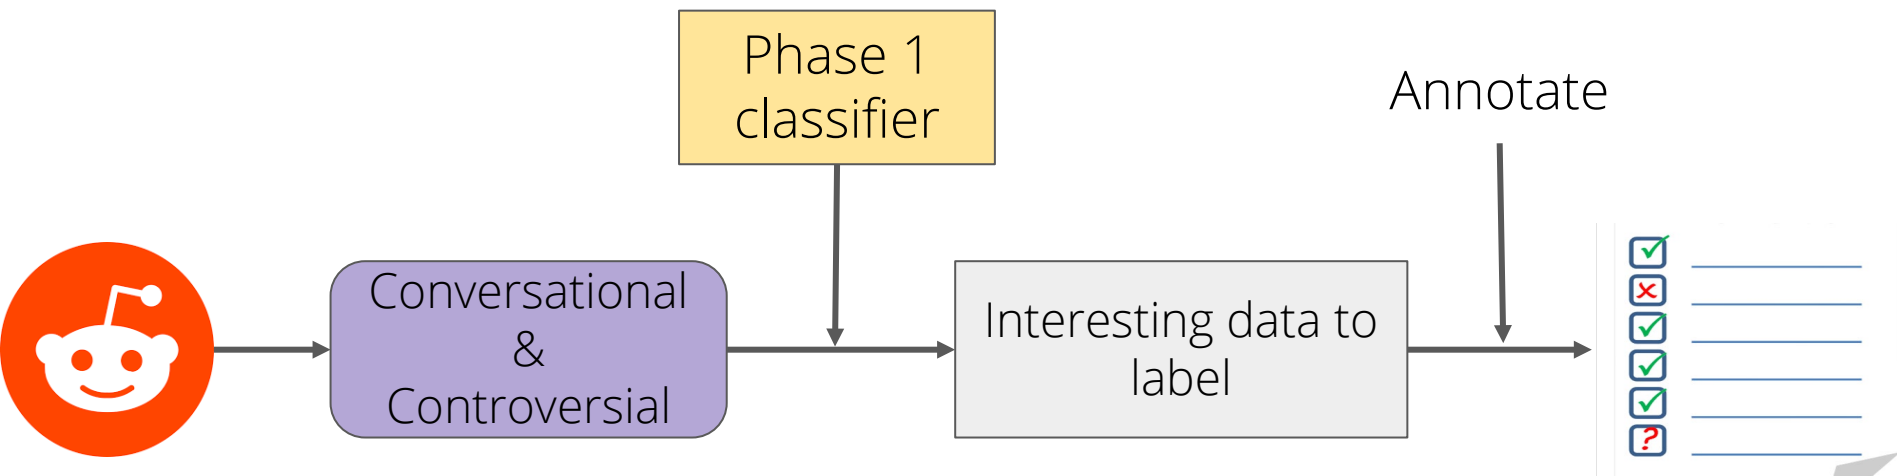


Boss to their employee

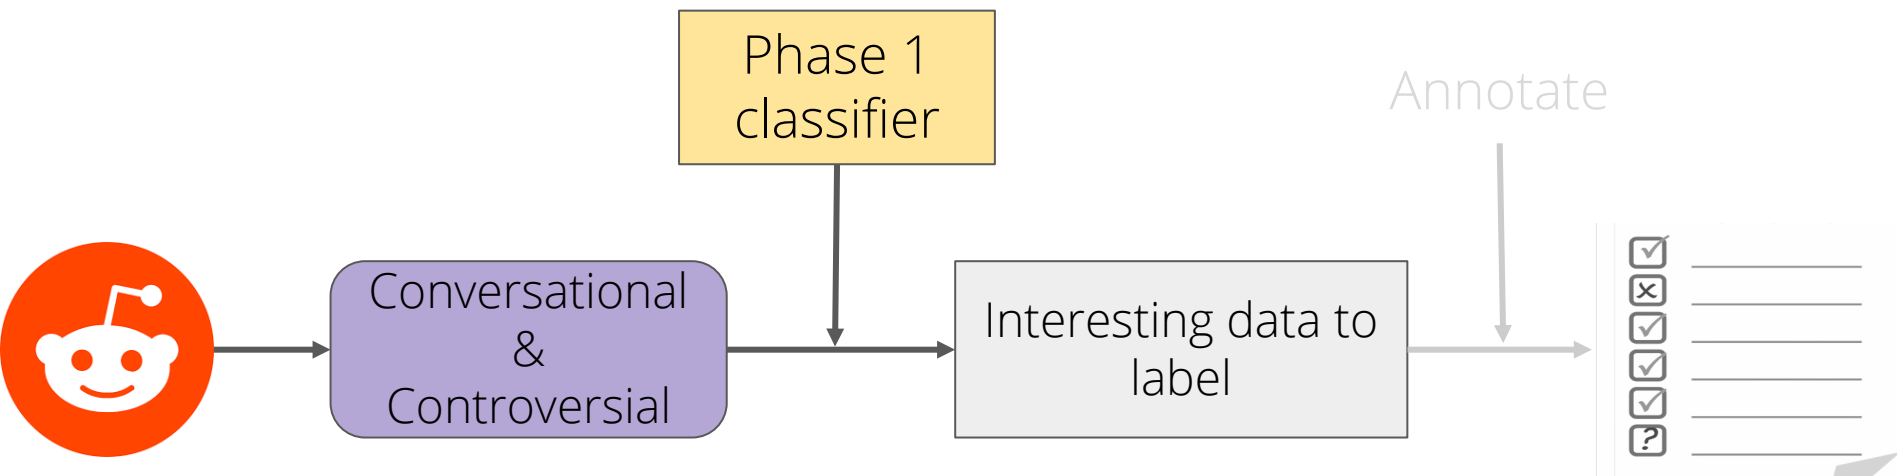


Phase 1 classifier


Dataset Creation (Phase 2)



Dataset Creation (Phase 2)



Dataset Creation (Phase 2)



Annotating appropriate is complex!

"We don't mention that name here."

Is the quote appropriate in the social context?

	Yes	No	Rare	N/A		Yes	No	Rare	N/A		Yes	No	Rare	N/A		Yes	No	Rare	N/A
friend	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	grandparent	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	domestic_partner	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	person with authority	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
best friend	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	grandchild	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	friends with benefits	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	law enforcement	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
childhood friend	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	uncle/aunt	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	divorcee	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	student	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
old friend	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	neice/nephew	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	ex-girlfriend/ex-boyfriend	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	teacher	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
neighbor	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	cousins	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	person having an affair	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	mentor	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
acquaintance	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	coworker	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	fan	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	mentee	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
complete stranger	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	boss	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	hero	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	landlord	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
parent	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	direct report (employee)	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	classmate	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	lawyer	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
child	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	employee in large company	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	sports teammate	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	client	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
adopted_child	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	colleague	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	club_member	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	doctor	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
sibling	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	dating	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	enemy	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	patient	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
step-sibling	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	engaged	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	rival	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	married	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
competitor	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>															

Move backward

Move forward



Measuring Annotation Quality

- Two trained annotators labeled 41 contexts, making 2,159 initial appropriateness judgments
- Initial agreement: Krippendorff's $\alpha = 0.46$ 🤖
- Why did annotators disagree? (in order)
 - Disagreement on message meaning (most cases)
 - Mistakes
 - Disagreement on appropriateness
- **Solution:** Annotators adjudicate the meaning of each message and then re-annotated
 - Final agreement $\alpha = 0.92$
- **Caveat:** One dataset—not necessarily expected to generalize

Dataset
Creation

Building a
Model

Apply the model to
understand



How much of the
conversations are
context-sensitive



Its usefulness to
model other
language behavior in
conversations

Can computational models identify contextual appropriateness?

Task: Given the social relationship identify if it would be appropriate to say a particular message


Example: Would asking "Is English your first language?" be considered appropriate between **work colleagues**

Prompt-based: 'Is it appropriate for `person1` to say "quote" to `person2`, "yes" or "no"? [mask]'


Positive class: Inappropriate

Fine-tuning: `Person1` saying "Quote" to `person2` is inappropriate. *Label (1/0)*


Can computational models identify contextual appropriateness?



Can computational models identify contextual appropriateness?




Can computational models identify contextual appropriateness?



Dataset
Creation

Building a
Model

Apply the model to
understand



How much of the
conversations are
context-sensitive



Its usefulness to
model other
language behavior in
conversations

How much of the conversations are sensitive to relationship-context?



“When you die do you want to be buried or cremated?”

Still appropriate?

What if we switch context?

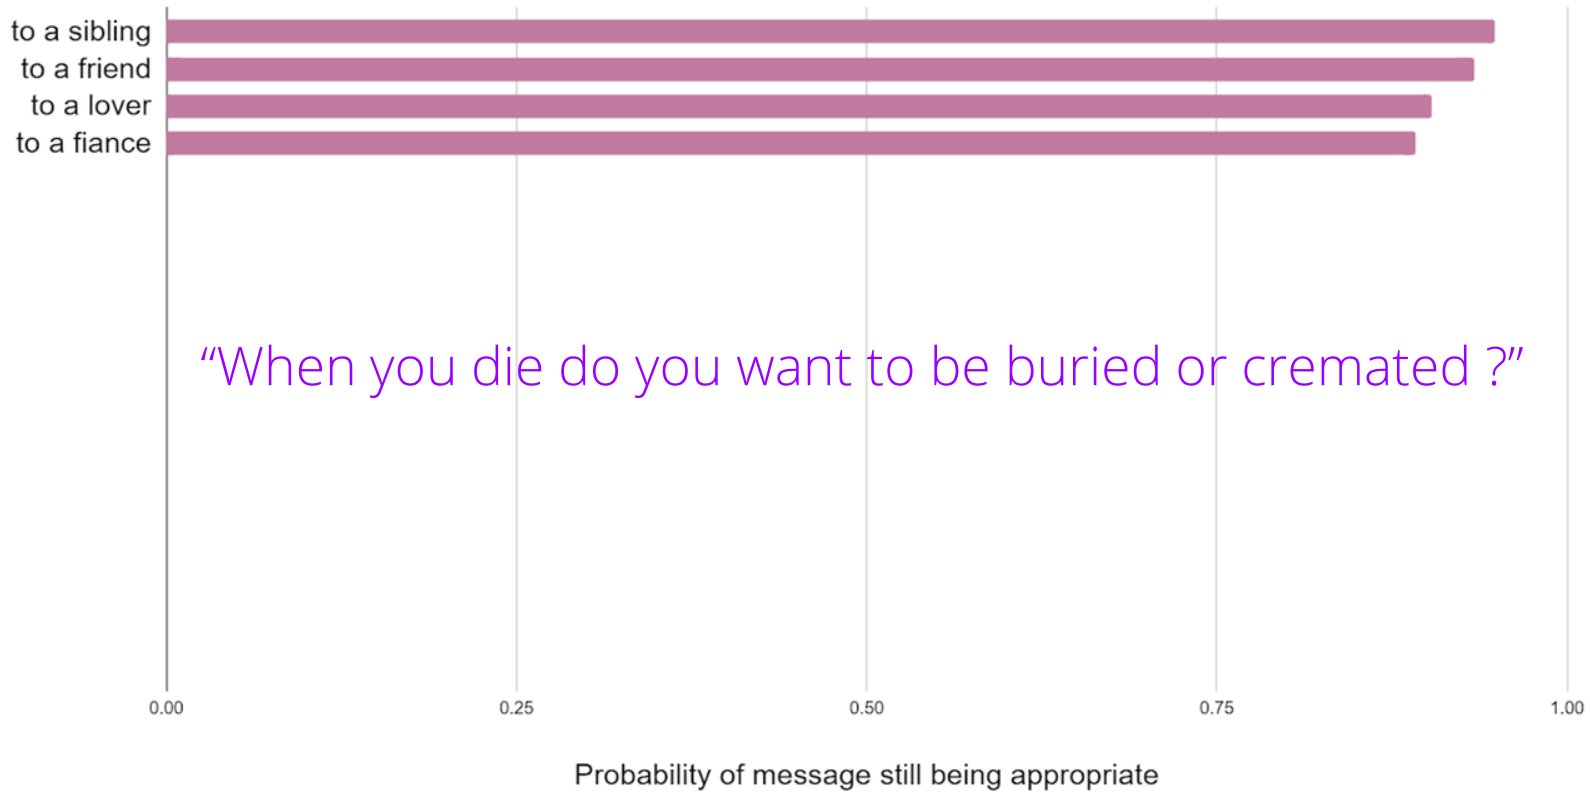


“When you die do you want to be buried or cremated?”

How much of the conversations are sensitive to relationship-context? (PRIDE; Tiginova et al 2021)

- Dialogue dataset with a known relationship between speaker and spoken to.
- Based on movie dialog but filtered down to generic messages
- We use a dataset of 47,801 messages across 18 relationships

If a message was appropriate in another relationship context is the message still appropriate if said to

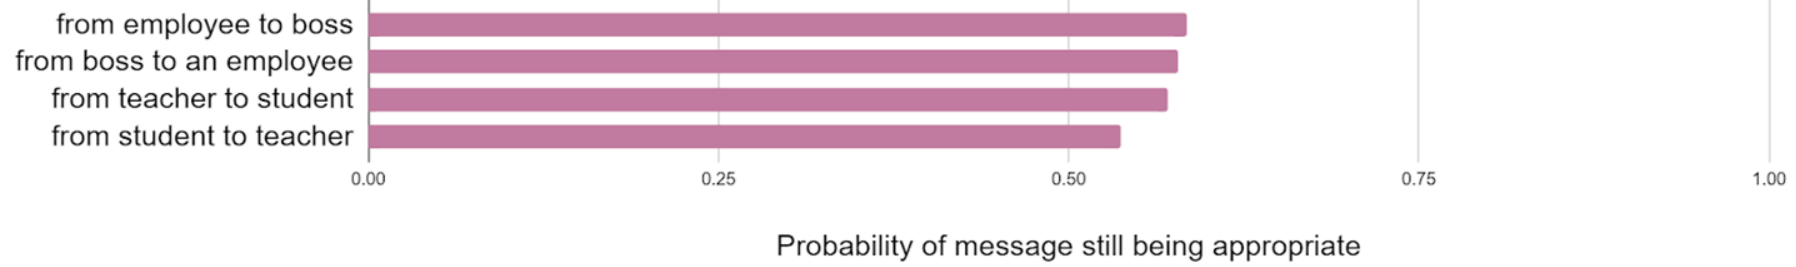


Relationship


If a message was appropriate in another relationship context is the message still appropriate if said to

Relationship

“When you die do you want to be buried or cremated ?”



If a message was appropriate in another relationship context is the message still appropriate if said to



~1/5th of the conversations change their appropriateness with context

Is appropriateness helpful for modeling other language behaviors? (Talkdown; Wang and Potts 2019)

- Aim: Detect condescending utterances.
 - Binary Classification task.
 - Why:
 - Appropriateness and condescension are related
 - Condescension is inappropriate in hierarchical relationships.
- “Bro, calm down and come back to real life.”

Is appropriateness helpful for modeling other language behaviors?

Condescending /
Not-condescending



LOGISTIC REGRESSION

Judgement of
inappropriateness for



Relationship 1



... Relationship 49

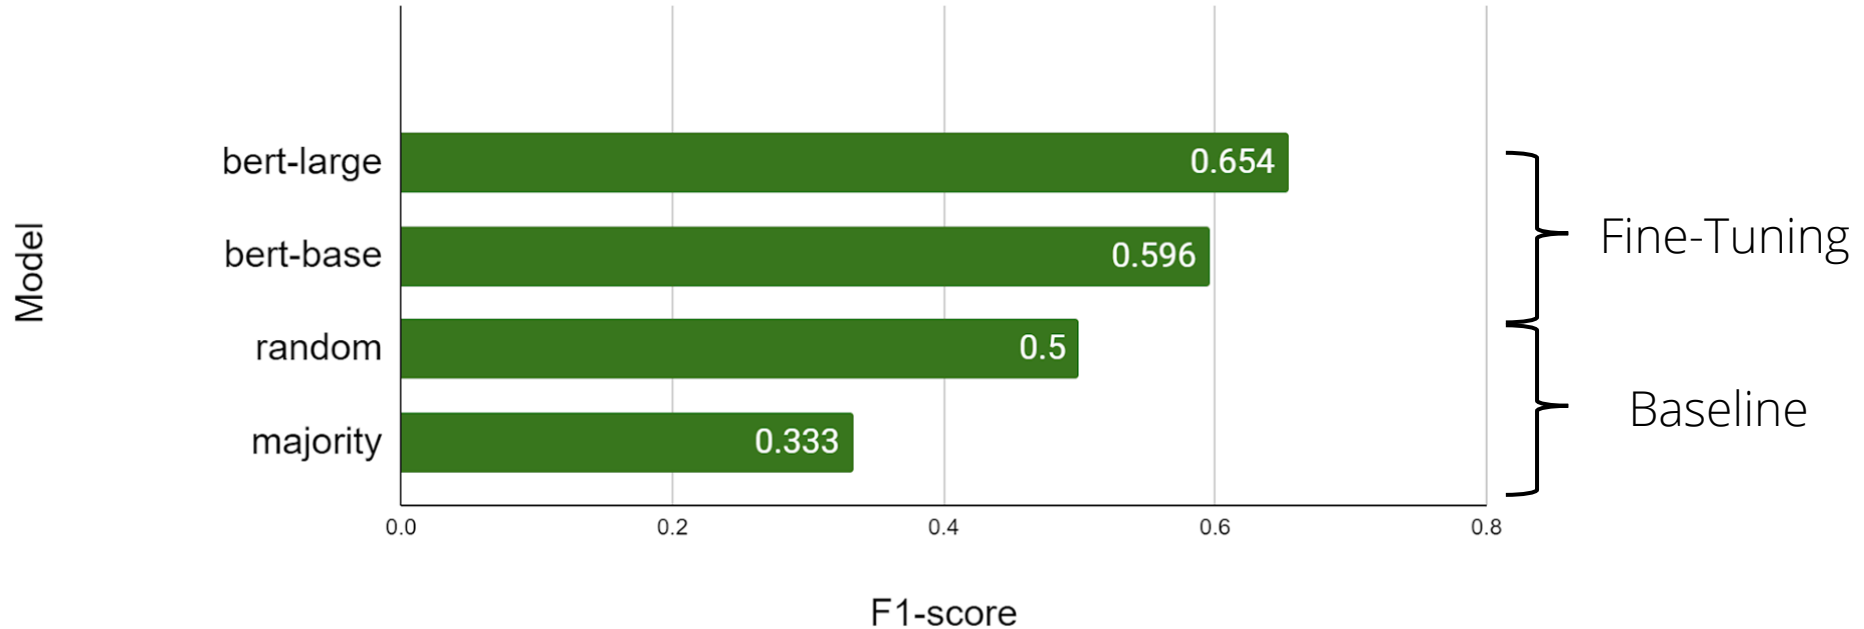


Contextually appropriate trained FLAN-T5-XL



“Bro, calm down and come back to
real life.”

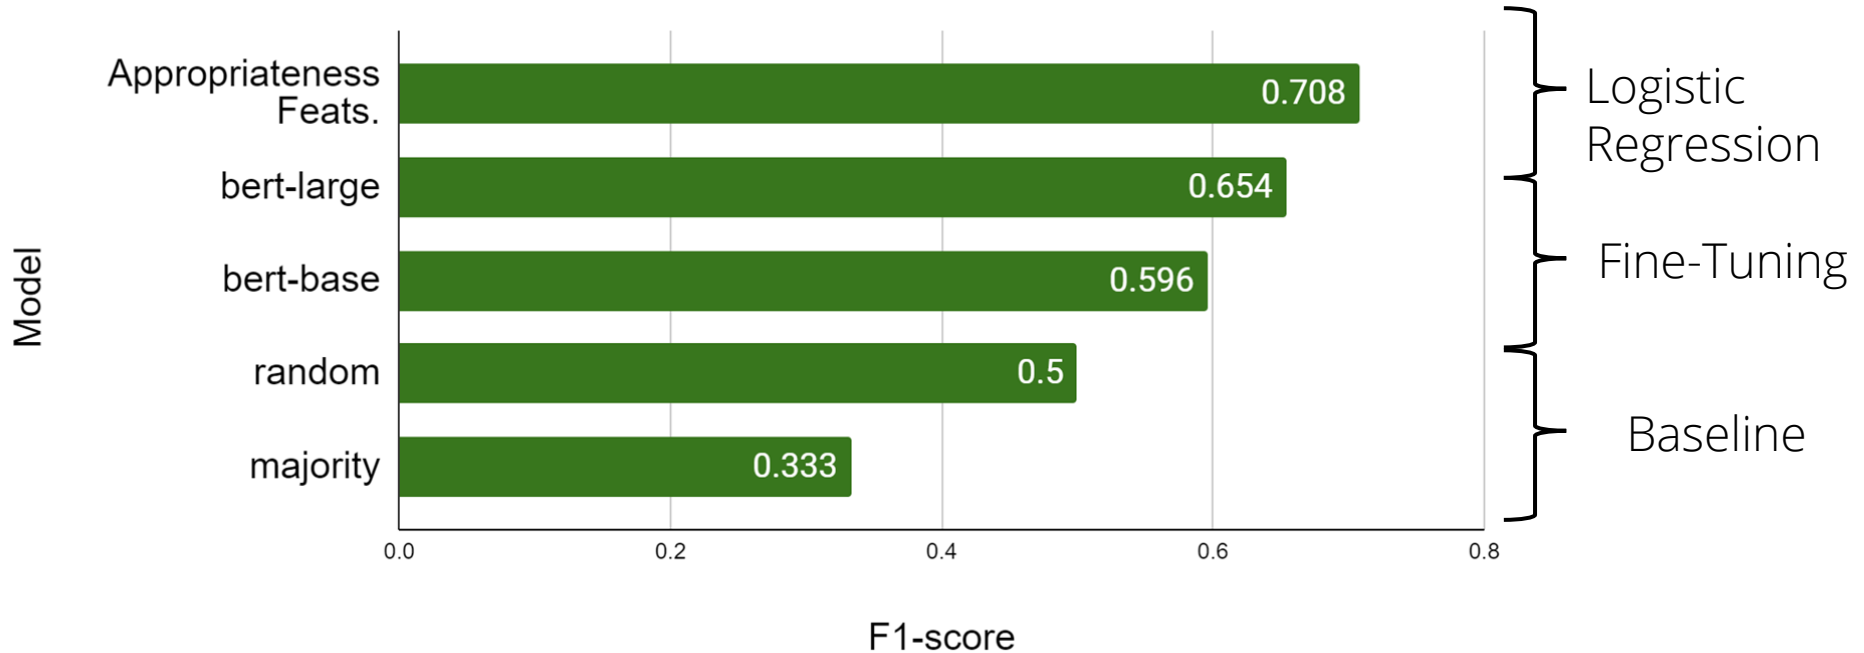
Is appropriateness helpful for modeling other language behaviors?



“Bro, calm down and come back to real life.”

Ngsgui kj P g {gzk {

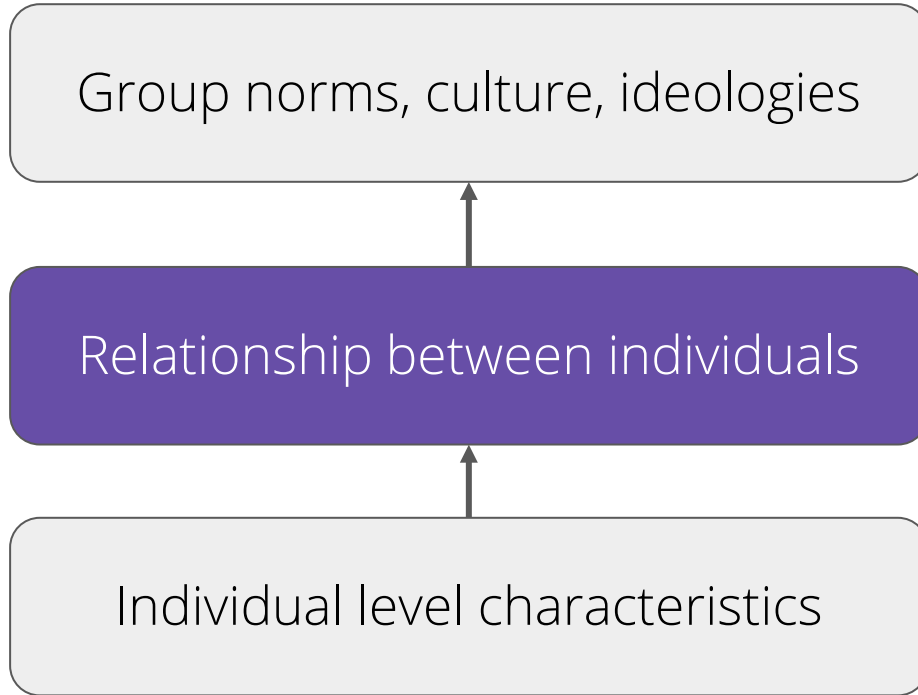
Is appropriateness helpful for modeling other language behaviors?



“Bro, calm down and come back to real life.”


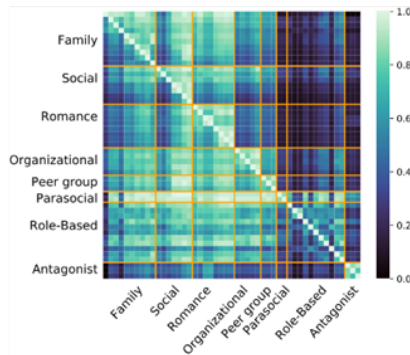
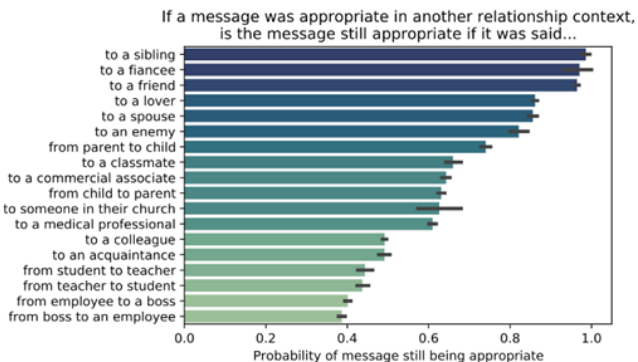
Ngsgui kj P g {gzk {

Limitations



- More relationships to be analyzed; each with more nuance
- More #annotators; diverse viewpoints.
- This work focuses on this aspect

Thank you!!



- We introduce a new dataset of 12K appropriateness judgments across 49 relationships.
- An estimated 1 in 5 messages change their appropriateness in a new context
- Incorporating contextuality helps understand other antisocial behaviors

<https://github.com/davidjurgens/contextual-appropriateness>



David Jurgens
jurgens@umich.edu



Agrima Seth
agrima@umich.edu



Jackson Sargent
jacsarge@umich.edu



Athena Aghighi
aaghighi@ucdavis.edu



Michael Geraci
megeraci@buffalo.edu