Inferring Mental States from Language Text

Lisa S. Pearl
lpearl@uci.edu

Computation of Language Laboratory
UC Irvine

Mark Steyvers, Brandon Turner, Joachim Vandekerckhove

AI/ML Seminar
Center for Machine Learning and Intelligent Systems
February 13, 2012

Mental states

People communicate these through the language they use:

- emotions like anger and embarrassment
- attitudes like confidence and disbelief
- intentions like persuasion or deception

Mental states through language text alone

Text message example:
You're in a rush, so when your friend texts you asking you to meet her later on, you text back a quick "Sure". She fires back a text asking if you're mad at her.

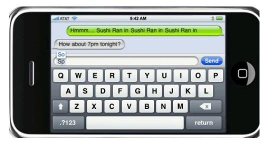

What happened?
Friend's inference: Terse message = angry

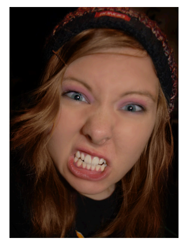

3

The process

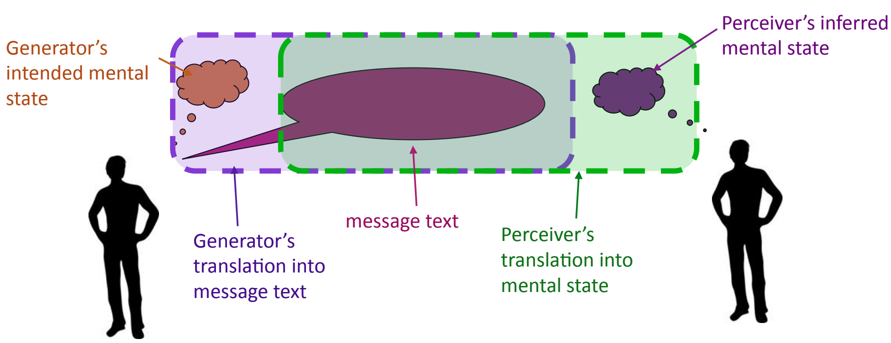

4

Research goals

Applied: Create tone-checker for email and text software

- automatic recognition

Psychological: Understanding the cognitive processes people use to transmit this information through language text

- linguistic cues
- processes underlying generation and perception: how related?

To address these, we need reliable data about the intended mental state of a message.

5

Getting reliable data

Use human-based computation (Kosurokoff 2001, von Ahn 2006) to construct a reliable database of messages expressing specific mental states. Specifically, use a **game with a purpose** (GWAP) (von Ahn and Dabbish 2004, von Ahn 2006, von Ahn, Kedia, and Blum 2006).

– "wisdom of the crowds" effect shown for many knowledge domains, including human memory, problem solving, and prediction (Steyvers et al. 2009, Turner & Steyvers 2011, Yi et al. (2012), Lee et al. (forthcoming))

– Snow et al. 2008: a relatively small number of non-expert annotations in natural language tasks can achieve the same results as expert annotation.

6

Using a game-with-a-purpose

Word Sleuth
Test your social language intelligence

http://gwap.ss.uci.edu

WordSleuth (Pearl & Steyvers 2010)
- encourages people to generate messages with a specific tone
- evaluates how these messages are perceived by others

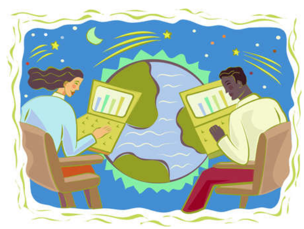

7

Word Sleuth
Test your social language intelligence

Asynchronous game play, with two roles: Word Sleuth & Expressor

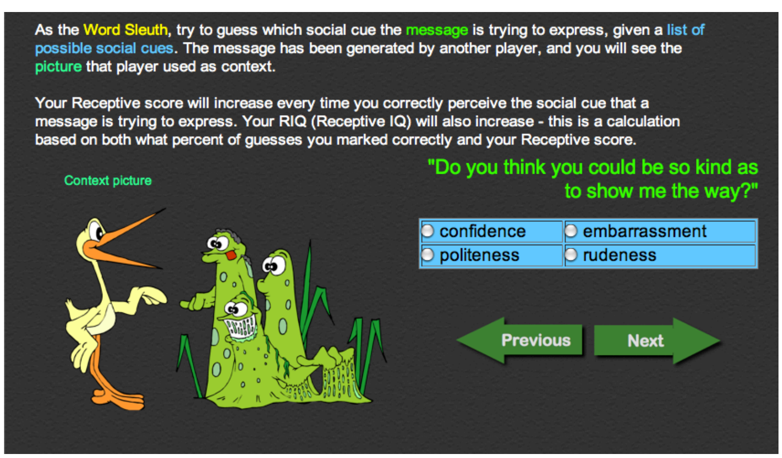

8

Word Sleuth

Test your social language intelligence

Asynchronous game play, with two roles: Word Sleuth & Expressor

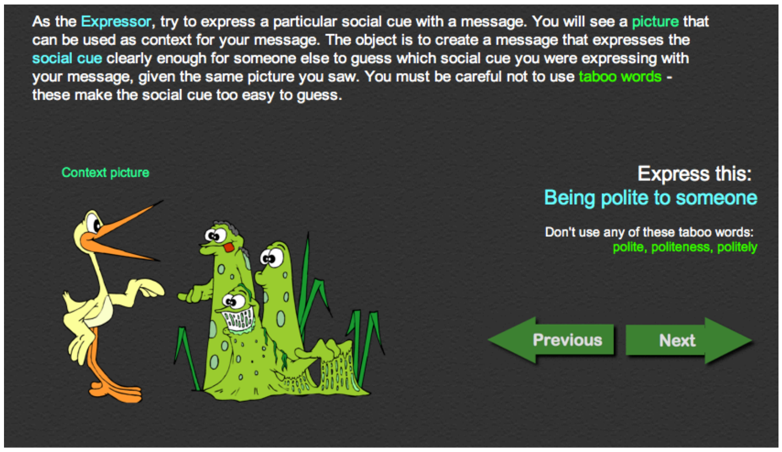

9

Word Sleuth

Test your social language intelligence

Asynchronous game play, with two roles: Word Sleuth & Expressor

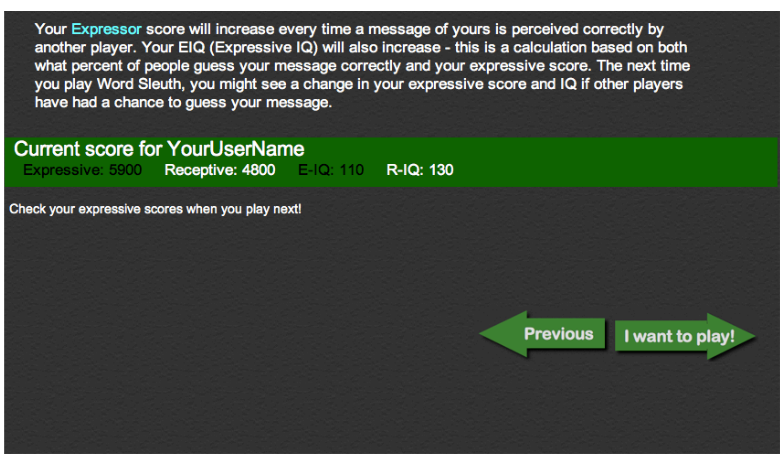

10

Word Sleuth
Test your social language intelligence

- Mental states currently explored (indicators for emotions, attitudes, and intentions):

politeness rudeness embarrassment formality
persuading deception confidence disbelief

- Messages are labeled by multiple participants.
- Participants never label their own messages.

11

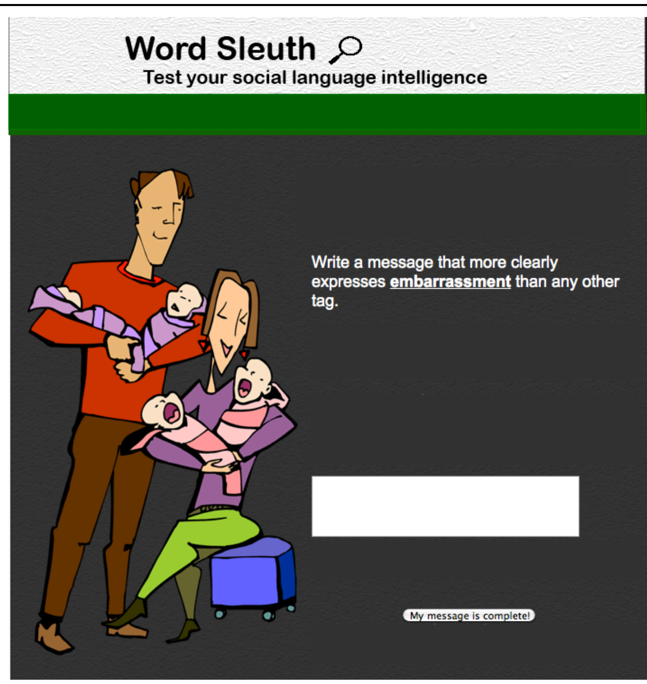

12

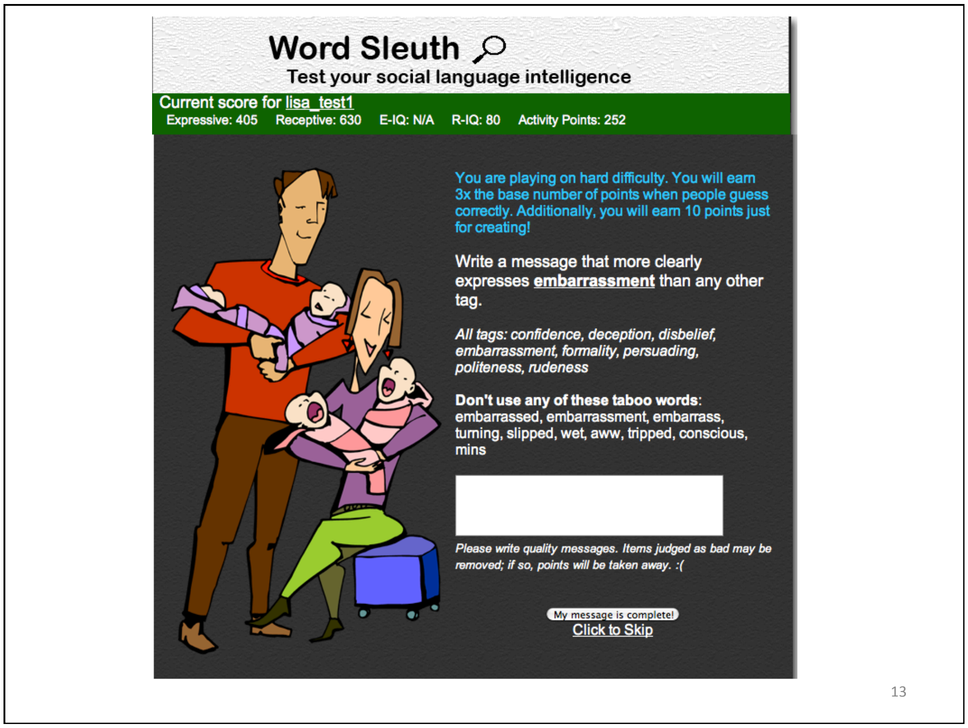
Word Sleuth
Test your social language intelligence
Current score for lisa_test1
Expressive: 405 Receptive: 630 E-IQ: N/A R-IQ: 80 Activity Points: 252
You are playing on hard difficulty. You will earn 3x the base number of points when people guess correctly. Additionally, you will earn 10 points just for creating!
Write a message that more clearly expresses embarrassment than any other tag.
All tags: confidence, deception, disbelief, embarrassment, formality, persuading, politeness, rudeness
Don't use any of these taboo words: embarrassed, embarrassment, embarrass, turning, slipped, wet, aww, tripped, conscious, mins
Please write quality messages. Items judged as bad may be removed; if so, points will be taken away. :(
My message is complete!
Click to Skip
13

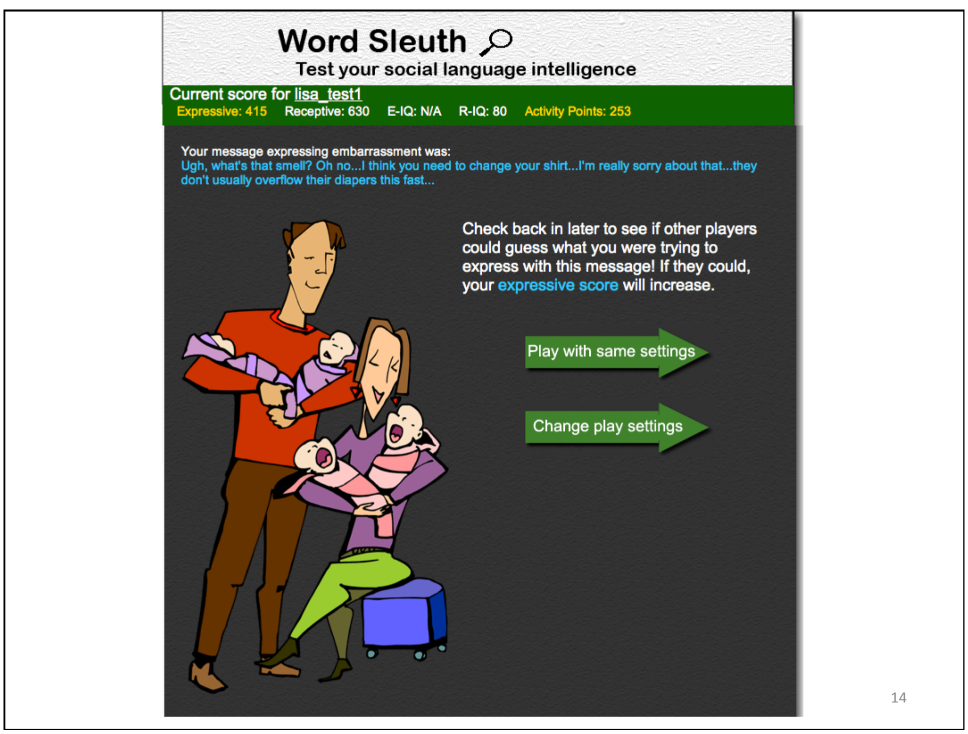
Word Sleuth
Test your social language intelligence
Current score for lisa_test1
Expressive: 415 Receptive: 630 E-IQ: N/A R-IQ: 80 Activity Points: 253
Your message expressing embarrassment was:
Ugh, what's that smell? Oh no...I think you need to change your shirt...I'm really sorry about that...they don't usually overflow their diapers this fast...
Check back in later to see if other players could guess what you were trying to express with this message! If they could, your expressive score will increase.
Play with same settings
Change play settings
14

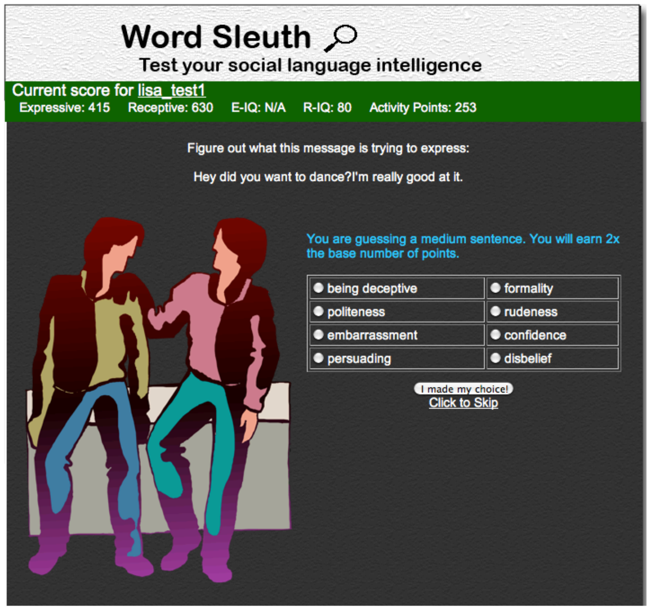
Word Sleuth
Test your social language intelligence
Current score for lisa_test1
Expressive: 415 Receptive: 630 E-IQ: N/A R-IQ: 80 Activity Points: 253
Figure out what this message is trying to express:
Hey did you want to dance?I'm really good at it.
You are guessing a medium sentence. You will earn 2x the base number of points.
being deceptive
formality
politeness
rudeness
embarrassment
confidence
persuading
disbelief
I made my choice!
Click to Skip

15

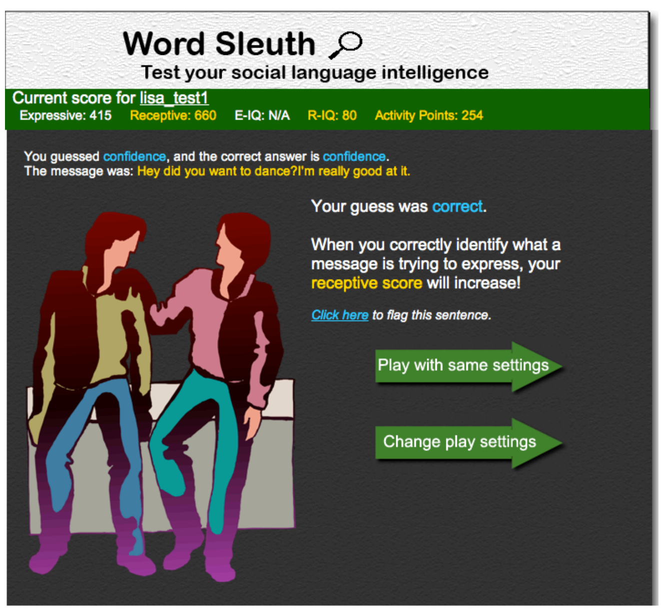
Word Sleuth
Test your social language intelligence
Current score for lisa_test1
Expressive: 415 Receptive: 660 E-IQ: N/A R-IQ: 80 Activity Points: 254
You guessed confidence, and the correct answer is confidence.
The message was: Hey did you want to dance?I'm really good at it.
Your guess was correct.
When you correctly identify what a message is trying to express, your receptive score will increase!
Click here to flag this sentence.
Play with same settings
Change play settings

16

Word Sleuth

Test your social language intelligence

(As of January 2012)

833 participants

3,832 messages created

22,689 labels inferred

- average of 5.92 inferred labels per message
- Over 900 messages with 10 or more labels

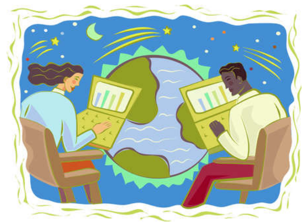

Humans are pretty good at this task:

- Overall accuracy of transmission: **0.74**

17

Word Sleuth

Test your social language intelligence

Not all mental states are created equal: some are more confusable than others

Inferred

Intended

	deception	politeness	rudeness	embarrassment	confidence	disbelief	formality	persuading	total labels
deception	**0.61**	0.05	0.05	0.02	0.07	0.04	0.02	0.13	2451
politeness	0.02	**0.72**	0.02	0.02	0.02	0.01	0.13	0.07	2666
rudeness	0.01	0.01	**0.85**	0.02	0.02	0.04	0.01	0.04	3252
embarrassment	0.03	0.05	0.02	**0.78**	0.01	0.07	0.02	0.02	3020
confidence	0.03	0.03	0.02	0.01	**0.81**	0.02	0.01	0.08	3204
disbelief	0.03	0.03	0.04	0.03	0.03	**0.80**	0.01	0.02	3369
formality	0.02	0.34	0.02	0.02	0.04	0.02	**0.46**	0.09	1810
persuading	0.05	0.04	0.02	0.00	0.08	0.01	0.02	**0.77**	2906
									22678

Deception & formality are harder:

- deception: involves a layer of semantic inversion
- formality: overlaps a lot with politeness

18

Word Sleuth
Test your social language intelligence

Sample messages

Intended Inferred	Message
confidence confidence	“here's the paper! i'm positive its really good this time”
rudeness rudeness	“You are the stinkiest person I've ever met.”
deception persuading	"I recommend that you take one step forward. Don't worry, it's not dangerous.”
formality politeness	“may i take the road on the left please”

19

Research Goal: Automatic Recognition

Prior research involving linguistic cues for identifying information in text has often used word-level cues (Anolli et al. 2002, Pang et al. 2002, Turney 2002, Zhou et al. 2004, Gupta & Skillicorn 2006).

Linguistic features: Basic shallow features (for now)

- unigrams, bigrams, and trigrams appearing more than once in the database
- number of word types, word tokens, and sentences per message
- average word length
- proportion of punctuation marks, 1st person pronouns, characters, digits
- average sentence length
- word type to word token ratio (more lexical diversity = higher score)
- average word log frequency of message for words appearing more than once in the database

~10,800 features

20

Research Goal: Automatic Recognition

Classifier goal: select the intended mental state from one of the eight options for a given message

- chance performance is 0.125
- 10-fold cross-validation (90% training and 10% test corpus)

Complete data set: 3832 messages

- baseline performance of choosing the most frequent label in the training set: (~0.137)

Filtered data set: 1303 messages where at least 2 people inferred the label and there was more than 50% agreement on the inferred label.

- hope: if there's more than 50% agreement on the label, it's likely reliable
- baseline performance of choosing the most frequent label in the training set: (~0.130)

21

Research Goal: Automatic Recognition

Naïve Bayes Classifier

- Complete data set: **0.578**
- Filtered data set: **0.619** (some improvement over complete dataset)

Classifier trained on either data set is between 4 and 5 times as good as chance (0.125) or the baseline strategy of choosing the most frequent label in the training set (0.130-0.137).

22

Research Goal: Automatic Recognition

Naïve Bayes Classifier: confusion matrix for filtered dataset

- All mental states are not created equal (total messages passing criteria varies significantly, some mental states are more confusable than others)

Inferred

Intended	deception	politeness	rudeness	embarrassment	confidence	disbelief	formality	persuading	total messages
deception	**0.19**	0.03	0.14	0.14	0.09	0.21	0.00	0.20	109
politeness	0.01	**0.49**	0.25	0.03	0.03	0.09	0.00	0.09	150
rudeness	0.00	0.04	**0.76**	0.03	0.03	0.09	0.00	0.05	217
embarrassment	0.01	0.01	0.11	**0.67**	0.07	0.11	0.00	0.02	179
confidence	0.01	0.02	0.07	0.05	**0.72**	0.09	0.00	0.05	184
disbelief	0.00	0.01	0.10	0.04	0.03	**0.79**	0.00	0.02	215
formality	0.01	0.31	0.31	0.04	0.04	0.08	**0.04**	0.17	72
persuading	0.01	0.04	0.16	0.02	0.04	0.05	0.00	**0.69**	177
									1303

Note: Precision is very good even though recall is very bad for deception and formality

Research Goal: Automatic Recognition

Sparse Multinomial Logistic Regression (SMLR) Classifier

(Krishnapuram et al. 2005) (λ = 0.05 [how strong a push for fewer non-zero weights])

- Complete data set: **0.584** [1118-1460 non-zero weighted features per state]
- Filtered data set: **0.665** [468-915 non-zero weighted features per state]

Classifier trained on either data set is nearly 5 times as good as chance (0.125), or the baseline strategy of choosing the most frequent label in the training set (0.130-0.137).

Research Goal: Automatic Recognition

Sparse Multinomial Logistic Regression (SMLR) Classifier

(Krishnapuram et al. 2005) (λ = 0.05 [how strong a push for fewer non-zero weights])

- Complete data set: **0.584** [1118-1460 non-zero weighted features per state]
- Filtered data set: **0.665** [468-915 non-zero weighted features per state]

Inferred

Intended	deception	politeness	rudeness	embarrassment	confidence	disbelief	formality	persuading	total messages
deception	**0.41**	0.01	0.16	0.08	0.15	0.07	0.04	0.08	109
politeness	0.06	**0.63**	0.10	0.04	0.05	0.07	0.02	0.03	150
rudeness	0.03	0.07	**0.64**	0.04	0.03	0.11	0.01	0.08	217
embarrassment	0.03	0.01	0.05	**0.75**	0.04	0.07	0.02	0.02	179
confidence	0.07	0.02	0.03	0.10	**0.70**	0.07	0.00	0.02	184
disbelief	0.03	0.01	0.05	0.06	0.04	**0.80**	0.01	0.00	215
formality	0.00	0.26	0.06	0.09	0.06	0.08	**0.38**	0.08	72
persuading	0.02	0.07	0.08	0.02	0.06	0.01	0.00	**0.74**	177
									1303

Note: Deception and formality are still harder, but precision is still pretty good

25

Research Goal: Automatic Recognition

Future directions:

- More sophisticated classifiers
- Better linguistic features
    - Psychological & machine learning goal: What kinds of linguistic cues are there for identifying these mental states in text?

26

Linguistic Features

Strongly weighted features from the SMLR classifier

[on the filtered dataset, with 1303 messages]

deception: - less use of ? and !
- more use of words/phrases like "promise", "I would never", and "I'm not"

politeness: - more use of words/phrases like "thanks" and "could you please"

rudeness: - less use of positive words like "nice" and beautiful", 1st pers pronouns
- more use of negative words like "annoying" and "idiot"

embarrassment: - more use of words like "forgot", "awkward", and "accidentally"
- less use of positive words like "great" and "good"

confidence: - more use of 1st pers pronouns ("me") and phrases like "I'm sure", "I know"
- less use of ? and negative adverbs like "never", "ca**n't**", "did**n't**"

disbelief: - more use of words/phrases like "surprised", "yeah right", and "no way", and ?

formality: - more use of titles like "sir", "miss", and "mrs" and modal verb "may"
- less use of contractions like "don't", and punctuation like ! and …

persuading: - more use of "you should", "guarantee", "trust me", and starting with "let's…"

27

Linguistic Features

More sophisticated linguistic features

- Syntactic: parts of speech (modal verbs: "should", "must", "may"; negation; negative adverbs), contractions, capitalization patterns, emoticons
- Lexical semantic classes: "promise" verbs, positive vs. negative valence, "apology" nouns and verbs, "certainty" adjectives, "surprise" nouns and verbs, titles of address ("sir"), "trust" verbs
  - Pull from WordNet synsets (Miller 1995, Fellbaum 1998) or from topic models whose topics use these keywords (Griffiths & Steyvers 2002)

Goal: improve machine learning & also understand what cues people are using, especially when they make mistakes

28

When people make mistakes

Agreement on "mistakes"?!

- "Oh dear, I'm so sorry…I had no idea this left a stain…maybe we should turn around before we cause any more damage…"
  Intended: politeness   Inferred (5 of 9): embarrassment

- "I can't believe John stood me up AGAIN, on our anniversary too."
  Intended: embarrassment   Inferred (6 of 8): disbelief

Question: Who's right? (Whose error is it?)

29

Research goal: Understand the cognitive processes

Why do people make mistakes?
What is the underlying process that produces the messages we observe?
How do people make inferences about the underlying mental state that produced those messages?

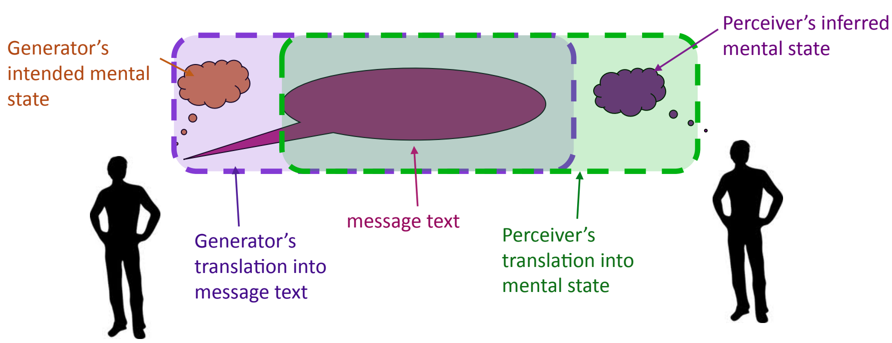

30

Using the data

- The intended label tells us about generative processes
  - How is a message created on the basis of some latent mental state?
- The inferred label tell us about perceptive processes
  - How is a message perceived and interpreted?
- Can we simultaneously learn from the intended label and the inferred label?
  - Can the triplet (intended label, message text, inferred label) inform us about the transmission process, and perhaps individual expertise?

31

One idea for the process: Analysis by synthesis

- In a generative model framework, the act of understanding a stimulus is based on inverting the generative process ("analysis by synthesis" [AxS])
  - A likely interpretation for a message would be based on simulating a forward process of generating that message, and choosing the mental state that most likely generated that message.

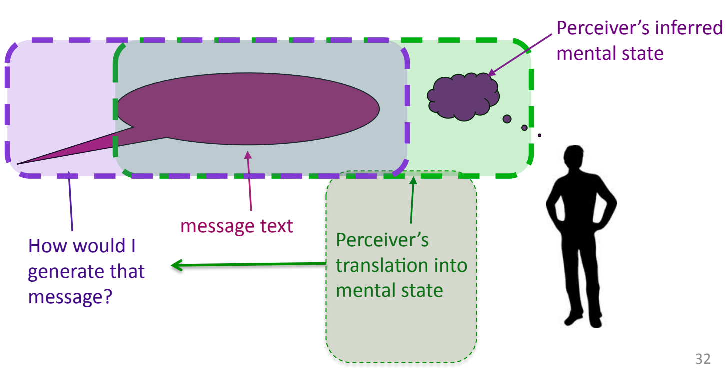

32

One idea for the process: Analysis by synthesis

- Similarly, the act of generating a stimulus can be based on inverting the inferential process ("synthesis by analysis" [SxA])
  - To generate a message, you would simulate a backward process of inferring a mental state from that message, and choose the message that expressed the intended mental state.

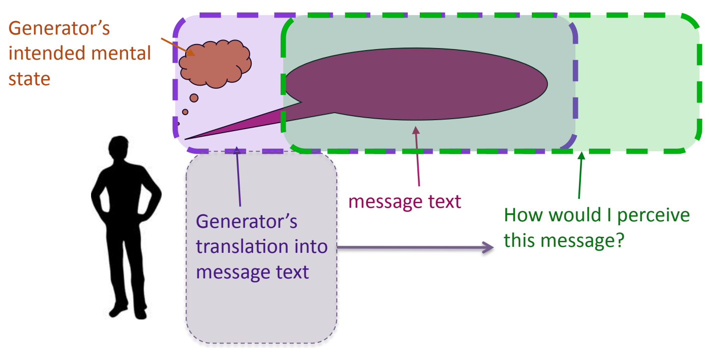

33

One idea for the process: Analysis by synthesis

- This idea is fundamental to many different cognitive process and machine learning models
  - Deep belief nets for handwriting analysis (e.g. Hinton & Salakhutdinov 2006, Hinton 2007)
  - Bayesian models of perception (Kersten & Yuille 2003, Cremers & Yuille 2003, Kertsen, Mamassian, & Yuille 2004, Lu & Yuille 2005, Chen, Zhu, Yuille, & Zheng 2009)
  - Linguistic theories on language comprehension & production (Bever & Poeppel 2010 for a review)
  - Models of causal reasoning (Baker, Saxe, & Tenenbaum 2009, Griffiths & Tenenbaum 2009, Kemp, Goodman, & Tenenbaum 2011, Goodman, Ullman, & Tenenbaum 2011)
  - Models of theory of mind & social goals (Baker, Goodman, & Tenenbaum 2008, Ullman, Baker, Macindoe, Evans, Goodman, & Tenenbaum 2010)

34

Another idea for the process: No analysis by synthesis

- Idea: The process of perception is really not at all related to the process of generation. These are completely separate abilities, which are not correlated at all. (This would clearly violate the assumptions of many successful models in perception, machine learning, and linguistics.)

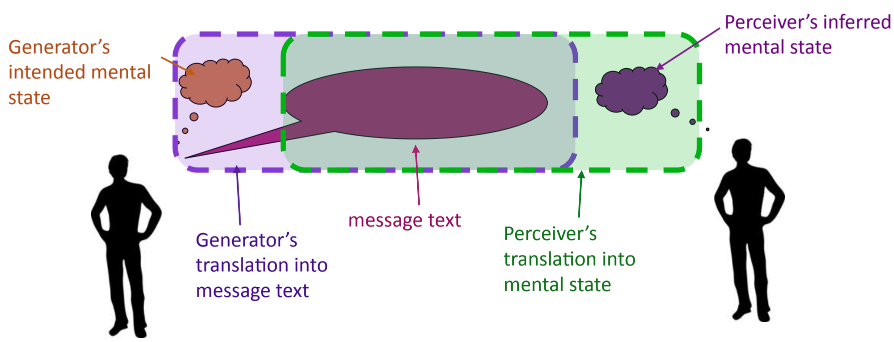

35

Cognitive Process Models: Overview

Analysis by Synthesis [AxS] (and Synthesis by Analysis [SxA])

- Simple: Compete Overlap
  - inference and generation use the same ability
- More Nuanced: Partial Overlap
  - inference and generation are different abilities, but both are used whether you're generating or inferring

No Analysis by Synthesis

- No Overlap
  - inference and generation are different abilities and inference is only used when inferring while generation is only used when generating

36

Analysis by Synthesis [AxS, SxA]: Simple formalization = Complete Overlap

Generator g: Use tr_g to translate from m_i to m_t. [SxA]
Perceiver p: Use tr_p to translate from m_t to m_{inf}. [AxS]

Transmission errors could be due to the poor tr_g or tr_p.
From any given triplet (m_i, m_t, m_{inf}), we can infer tr_g and tr_p.

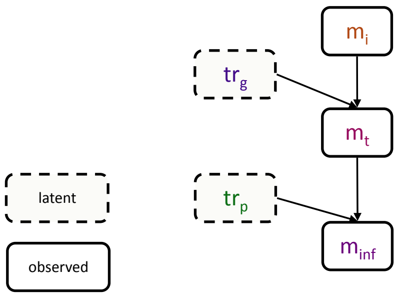

Each subject has one ability: translational ability

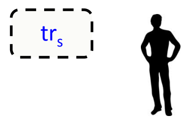

37

Analysis by Synthesis [AxS, SxA]: Simple formalization = Complete Overlap

Expectations:

- This ability is involved for each person for each message.
- A person's generative ability is completely identical to their perceptual ability.

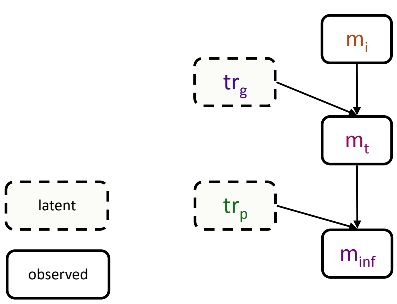

Each subject has one ability: translational ability

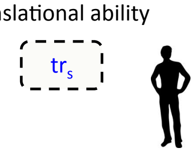

38

Model Specification

Model the decision process as a multinomial decision tree.

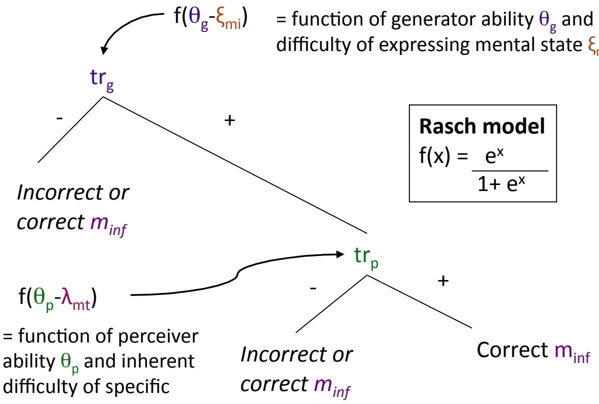

= function of generator ability θ_g and inherent difficulty of expressing mental state ξ_{mi}

= function of perceiver ability θ_p and inherent difficulty of specific message expressing particular mental state λ_{mt}

Rasch model

$$f(x) = \frac{e^x}{1+e^x}$$

Complete overlap

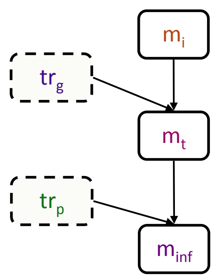

39

Model Specification

Model the decision process as a multinomial decision tree.

Each subject has one ability: translational ability

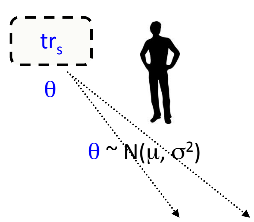

$$p(m_{inf} = m_i \mid m_i) = f(\theta_g - \xi_{mi}) f(\theta_p - \lambda_{mt}) + \phi$$

Probability of accidentally choosing correct m_{inf}

$\phi = \alpha(1 - f(\theta_g - \xi_{mi}) f(\theta_p - \lambda_{mt}))$

α estimated from empirical confusion matrix data or idealized to 1/ # states

Rasch model

$$f(x) = \frac{e^x}{1+e^x}$$

Complete overlap

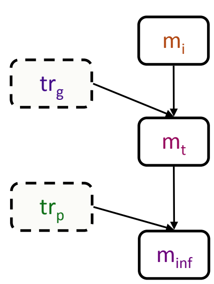

40

Model Specification

Represented in graphical notation.

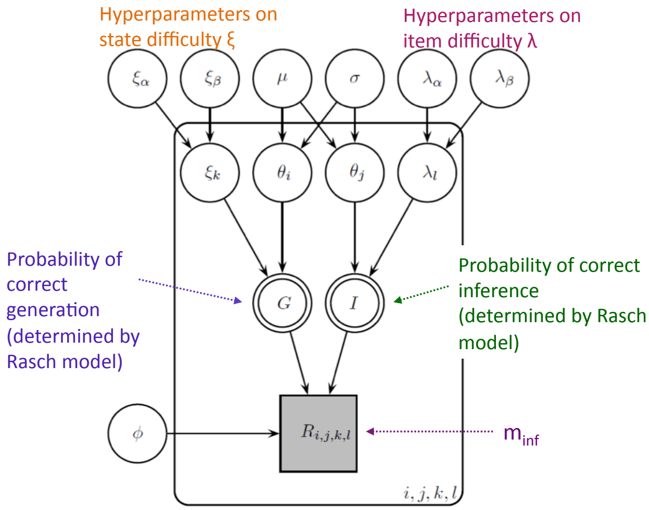

Complete overlap

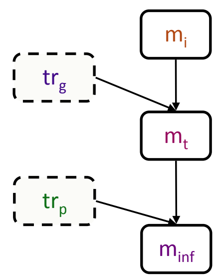

41

Analysis by Synthesis [AxS, SxA]: More nuanced formalization = Partial Overlap

Generator g: Use inf_g and gen_g to translate from m_i to m_t. [SxA]
Perceiver p: Use gen_p and inf_p to translate from m_t to m_{inf}. [AxS]

Transmission errors could be due to the poor gen_g, inf_g, gen_p, or inf_p.
From any given triplet (m_i, m_t, m_{inf}), we can infer gen_g, inf_g, gen_p, or inf_p.

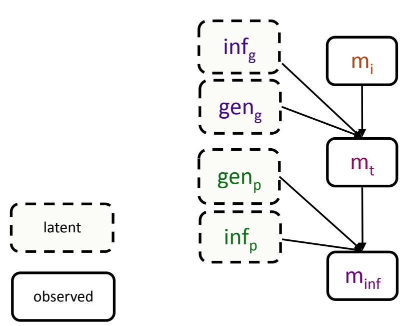

Each subject has two abilities:
generative ability & inferential ability

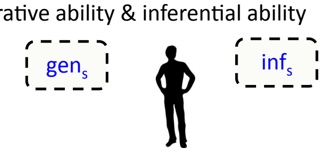

42

Analysis by Synthesis [AxS, SxA]: More nuanced formalization = Partial Overlap

Expectations:

- Both abilities are involved for each person for each message.
- Abilities do not need to be equal – this allows for outside factors to affect one process but not another (e.g. more likely to think people intend to persuade rather than deceive, so this affects interpretation more than generation).

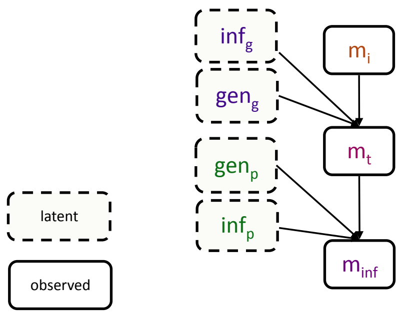

Each subject has two abilities:
generative ability & inferential ability

43

Model Specification

Model the decision process as a multinomial decision tree.

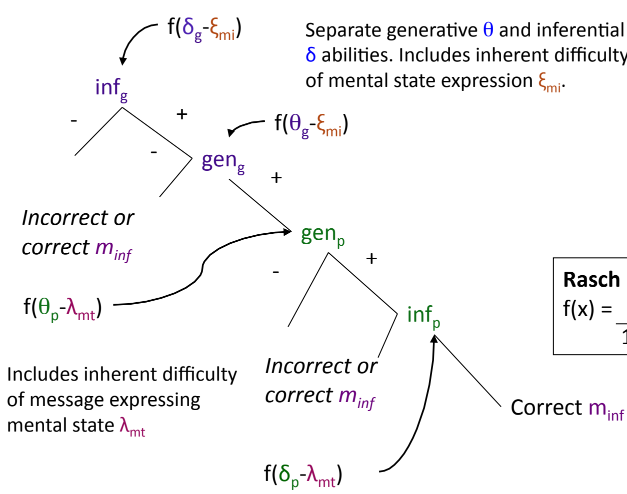

Partial overlap

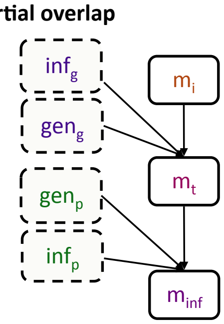

Rasch model

$$f(x) = \frac{e^x}{1+e^x}$$

44

Model Specification

Model the decision process as a multinomial decision tree for each model type

Each subject has two abilities:
generative ability & inferential ability

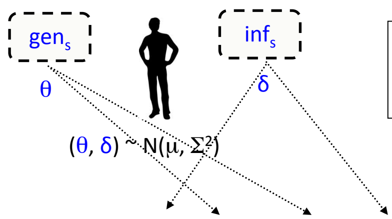

Rasch model
$f(x) = \frac{e^x}{1+e^x}$

Partial overlap

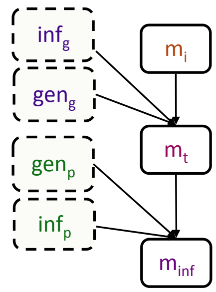

$p(m_{inf} = m_i \mid m_i) = f(\delta_g - \xi_{mi})\, f(\theta_g - \xi_{mi})\, f(\theta_p - \lambda_{mt})\, f(\delta_p - \lambda_{mt}) + \phi$

Probability of accidentally choosing correct m_{inf}

$\phi = \alpha(1 - f(\delta_g - \xi_{mi})\, f(\theta_g - \xi_{mi})\, f(\theta_p - \lambda_{mt})\, f(\delta_p - \lambda_{mt}))$

α estimated from empirical confusion matrix data or idealized to 1/ # states

45

Model Specification

Represented in graphical notation.

Hyperparameters on state difficulty ξ

Hyperparameters on item difficulty λ

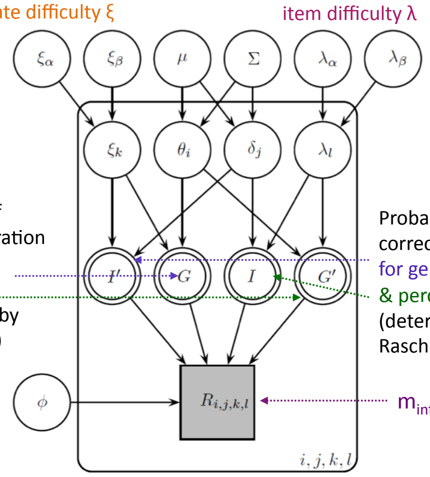

Probability of correct generation for generator & perceiver (determined by Rasch model)

Probability of correct inference for generator & perceiver (determined by Rasch model)

m_{inf}

Partial overlap

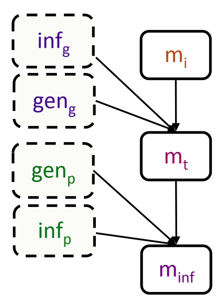

46

No analysis by synthesis = No Overlap

Generator g: Use gen_g to translate from m_i to m_t.
Perceiver p: Use inf_p to translate from m_t to m_{inf}.

Transmission errors could be due to the poor gen_g or inf_p.
From any given triplet (m_i, m_t, m_{inf}), we can infer gen_g and inf_p.

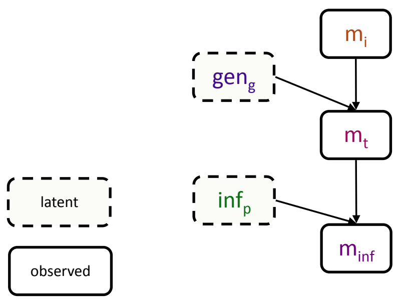

Each subject has two abilities:
generative ability & inferential ability

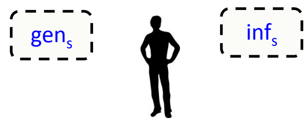

47

Model Specification

Model the decision process as a multinomial decision tree.

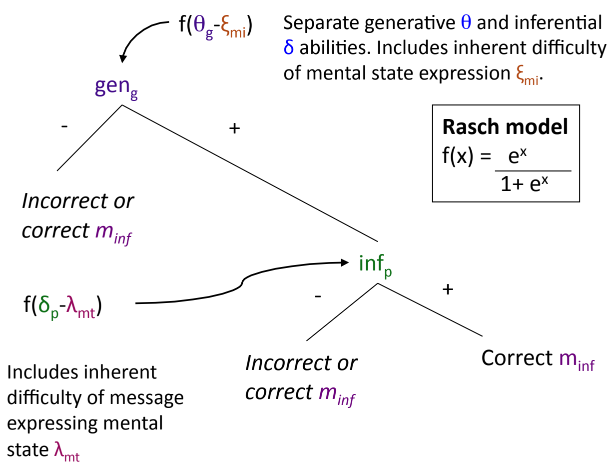

Rasch model
$$f(x) = \frac{e^x}{1+e^x}$$

No overlap

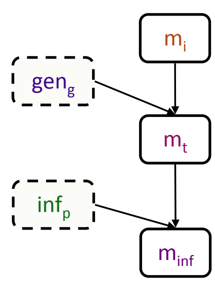

48

Model Specification

Model the decision process as a multinomial decision tree.

Each subject has two abilities:
generative ability & inferential ability

gen_s θ inf_s δ

$(\theta, \delta) \sim N(\mu, \Sigma^2)$

Rasch model

$f(x) = \frac{e^x}{1+e^x}$

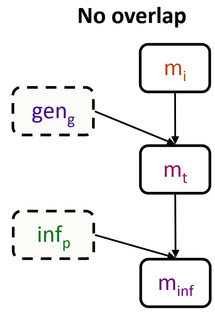

$p(m_{inf} = m_i \mid m_i) = f(\theta_g - \xi_{mi}) f(\delta_p - \xi_{mt}) + \phi$

Probability of accidentally choosing correct m_{inf}

$\phi = \alpha(1 - f(\theta_g - \xi_{mi}) f(\delta_p - \lambda_{mt}))$

α estimated from empirical confusion matrix data or idealized to 1/ # states

Model Specification

Represented in graphical notation.

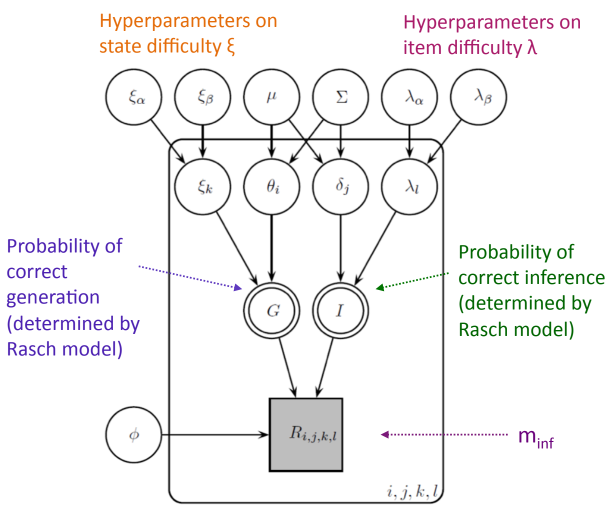

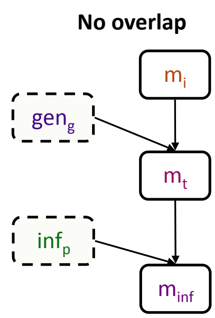

Testing the models

We train them on the data we have from humans, which consists of an intended mental state (m_i), message text (m_t), and an inferred mental state (m_{inf}). Each model will set the parameters as best it can to cover these training data.

Individual expertise parameters

Analysis by Synthesis

- Complete overlap: Each subject s has tr_s (or $gen_s = inf_s$)
- Partial overlap: Each subject s has gen_s and inf_s (gen_s not correlated with inf_s)

No Analysis by Synthesis

- No overlap: Each subject s has gen_s and inf_s (gen_s not correlated with inf_s)

51

Ability correlation

Do we see correlation in generation θ and inference δ abilities if we don't explicitly build it in? (Does complete overlap [θ = δ] happen naturally when trying to explain the observable data?)

No Overlap Model

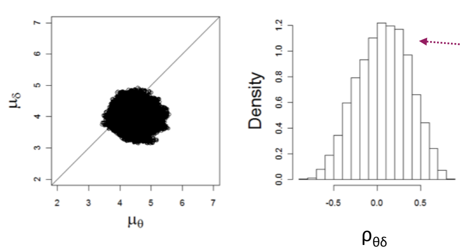

Not really – though there is some correlation in abilities (slight positive shift for the normal curve).

People vary more on their generation abilities, as compared with their inference abilities.

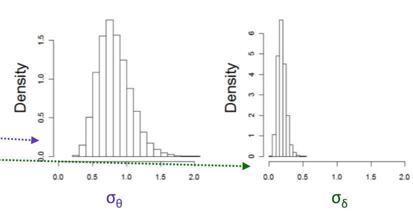

52

Ability correlation

Do we see correlation in generation θ and inference δ abilities if we don't explicitly build it in? (Does complete overlap [$\theta = \delta$] happen naturally when trying to explain the observable data?)

Partial Overlap Model

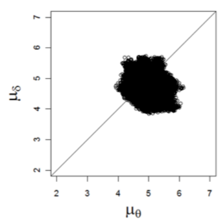

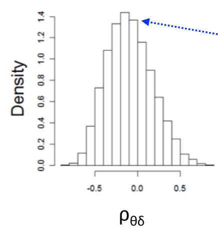

Really not – if anything, there's a slight negative shift suggesting these abilities are negatively correlated.

People still vary more on their generation abilities, as compared with their inference abilities.

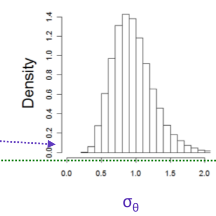

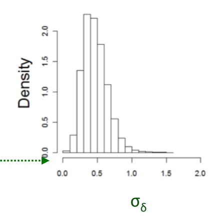

53

Model success at prediction

We might think that the No Overlap and Partial Overlap models are winning out – but it's worth asking how well these models predict the data.

- Natural expectation: Models with more ability parameters involved do better (No Overlap and Partial Overlap have separate gen_s and inf_s ability parameters [θ and δ]). Do these models do significantly better than the Complete Overlap model at predicting the observable data?

Use the Deviance Information Criterion **[DIC]** (Spiegelhalter et al. 2002) to compare models (smaller DIC values indicate better models).

- DIC includes both model fit and number of parameters involved
- Particularly useful for assessing models with posterior distributions attained by Markov chain Monte Carlo simulation

We'd also like to use confusion matrix data to see if and how these model predictions differ from human transmission behavior.

54

Test 1: Predict the inferred mental state

Given the intended mental state and a message, how well can the model predict the inferred mental state? (Given the generator's and perceiver's abilities)

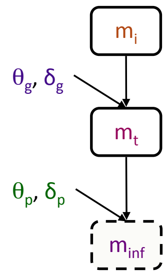

"What is subject s_2 likely to think this message is, if it was produced from the following mental state by subject s_1?"

55

Test 2: Predict the intended mental state

Given a message and the inferred mental state, how well can the model predict the intended mental state (given the generator's abilities)?

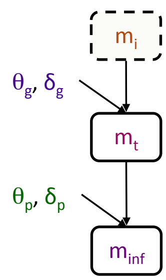

"What is subject s_1 likely to have intended this message to express, if it was interpreted by s_2 as this mental state?"

56

Integrating linguistic cues
Test 3: Predict both the intended and inferred mental state

Given a message and its linguistic cues, how well can the model predict the intended mental state and the inferred mental state? (Given the generator's and perceiver's abilities)

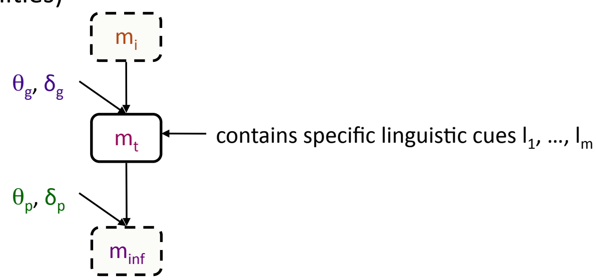

"Given this message generated by subject s_1 and interpreted by subject s_2, what is subject s_1 likely to have intended and what is subject s_2 likely to think this message expresses?

Requires us to know individual expertise with respect to linguistic cues, as opposed to just mental state type

57

Testing cognitive process modeling assumptions

- Our data allows us to test a major assumption of the generative modeling framework
- Perhaps there are processes in understanding that are unrelated to generation
  - This suggests limits on a simplistic "analysis by synthesis" framework (especially for this task)
  - Maybe better machine learning results for other tasks can be obtained with models that allow for more complex versions of analysis by synthesis (specifically: allow for some differences in the analysis component as compared to the synthesis component).

58

Big Picture

Machine Learning Goal: Create tone-checker for email and text software

- leads to identification of linguistic cues that are useful for machine learning

Psychological Goal: Understanding the cognitive processes people use to transmit this information through language text

- linguistic cues that humans use
- how do humans create these messages?
- leads to testing a modeling assumption used in both cognitive process and machine learning models

Thank You!

Galina Tucker **Shannon Stanton** **Joseph Nunn**
Uma Patel **Lawrence Phillips** **Sue Braunwald**
Hilary Cunningham **Sarah Pieper**

Audience at NAACL 2010 Emotions Workshop

The members of the UCI CoLaLab

Computation of
Language
Laboratory

UC Irvine