Using Speakers' Referential Intentions to Model Early Cross-Situational Word Learning

Frank, Goodman & Tenenbaum

presented by: Colin Kupitz

The Chicken or the Egg?

- When learning their first words, children face a joint-inference problem: trying to infer meaning and lexicon simultaneously
- Computationally hard without knowing one piece first

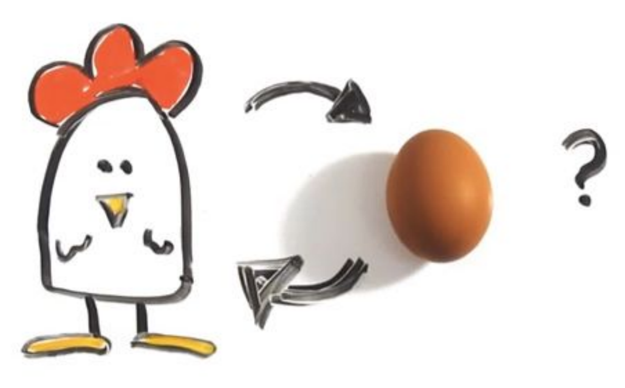

Accounts of Word Learning

- Social theories:
  - Depend on a rich understanding of the goals/intentions of speakers
- Cross-Situational:
  - Focus on the fact that words often refer to immediate environment of speaker

The Chicken *and* the Egg

- Current study presents a model that attempts to capture both aspects of the word-learning task - it simultaneously infers attempted communication and lexicon

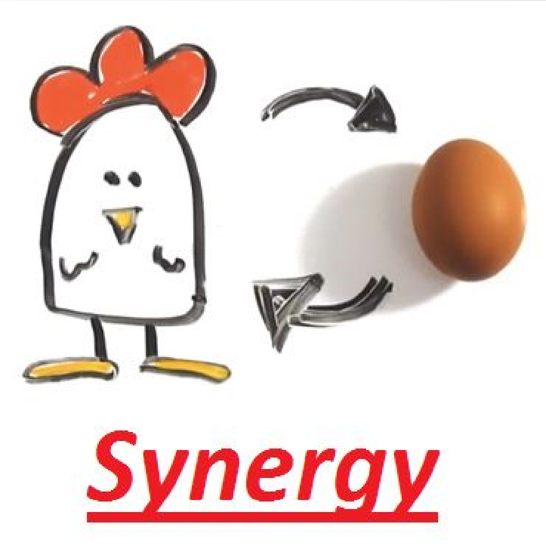

The Intentional Model

- What speakers intend to say is a function of the physical world around them
- Utterances are a function of what the speakers intend to say / how those intentions can be translated into language

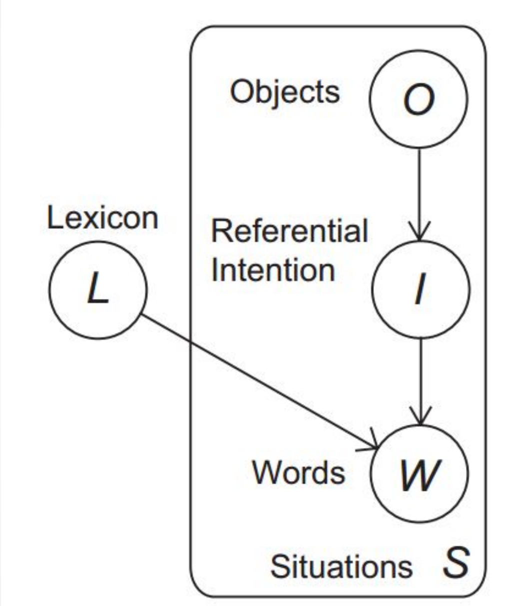

Assumptions

- Simple intentions - refer to objects at hand
- Basic-level objects / categories (no ‘aspect’ level references)
- L = Lexicons
- C = Corpus of Situations

$$P(L|C) \propto P(C|L)P(L).$$

All about dat Bayes

Likelihood

- Product over situations of the probability components of the corpus (Words W, Objects O, Intentions I) - given the lexicon

$$P(C|L) = \prod_{s \in C} P(W_s, O_s, I_s|L).$$

Likelihood 2

- Words W and Objects O are independent (given Intentions I)
- Rewrite eqn 1 as the probability of the words given intentions and lexicon * probability of intentions given the physical context

$$P(C|L) = \prod_{s \in C} \sum_{I_s \subseteq O_s} P(W_s|I_s, L) \cdot P(I_s|O_s).$$

Likelihood 3

- Assume words W are independent (no syntax)
  - 2 causes for utterance: Referential or nonreferential
  - γ = probability a word is used referentially in given context
  - P_R = prob. of word utterance if it's used referentially
  - P_{NR} = prob. a word picked from lexicon at random

$$P(W_s|I_s, L) = \prod_{w \in W_s} \left[\gamma \cdot \sum_{o \in I_s} \frac{1}{|I_s|} P_{\mathrm{R}}(w|o, L) + (1-\gamma) \cdot P_{\mathrm{NR}}(w|L) \right].$$

JUDGING!

- Compare Intentional model to several other models
- Evaluate on the accuracy of lexicon and inferences regarding speakers' intent
- Each model produced a single summary statistic linking Words-> Objects

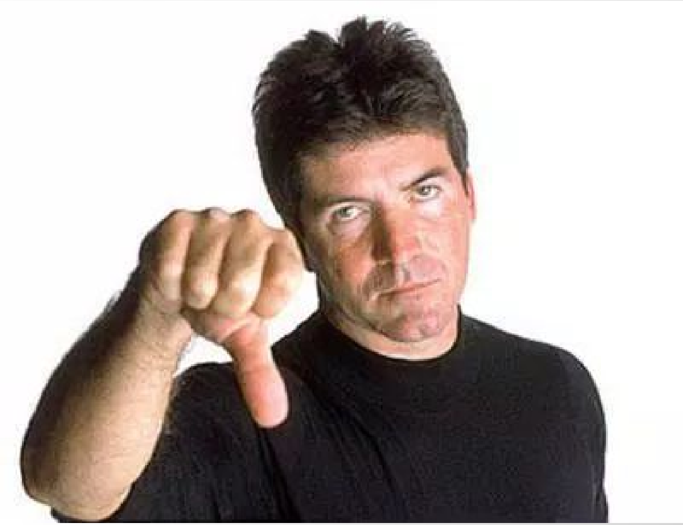

Still Judging!

- Chose the threshold for the summary stat. that maximized *F* score
  - The harmonic mean of precision (proportion of correct pairing) and recall (proportion of total correct pairings found)
- Compute scores relative to a gold-standard lexicon and gold-standard set of intentions from a human coder

Annnd Judged.

- Intentional Model beat the crap out of the others in ‘Best Lexicon’
- The more basic models had a large number of spurious lexical items.

A baby is required in every presentation right?

Well, here’s SuperBaby

Annnd Judged.

TABLE 1

Precision, Recall, and F Score of the Best Lexicon Found by Each Model When Run on the Annotated Data From the Child Language Data Exchange System

Model	Precision	Recall	*F* score
Association frequency	.06	.26	.10
Conditional probability (object|word)	.07	.21	.10
Conditional probability (word|object)	.07	.32	.11
Mutual information	.06	.47	.11
Translation model (object|word)	.07	.32	.12
Translation model (word|object)	.15	.38	.22
Intentional model	**.67**	**.47**	**.55**
Intentional model (one parameter)	.57	.38	.46

Words, words, words.

TABLE 3

The Best Lexicon Found by the Intentional Model

Word	Object	Word	Object
bear	**bear**	**lamb**	**lamb**
bigbird	**bird**	laugh	cow
bird	**duck**	meow	baby
birdie	**duck**	mhmm	hand
book	**book**	**mirror**	**mirror**
bottle	bear	**moocow**	**cow**
bunnies	**bunny**	oink	pig
bunnyrabbit	**bunny**	on	ring
hand	**hand**	**pig**	**pig**
hat	**hat**	put	ring
hiphop	mirror	**ring**	**ring**
kittycat	**kitty**	**sheep**	**sheep**

One parameter Intentional Model?

- κ, how likely in-lexicon words were to be used NR, and γ, the prob. of using words R, were set to their maximum *a posteriori* values.
- Uhh, is this cheating?

Intentions?

- Intentional Model similarly the best overall for Intentions
- Not the best at Recall
  - Recall = proportion of total correct pairing that were found
  - Likely due to more ‘shotgun’ approach that arises with a gigantor lexicon as in the Associate Frequency model (highest Recall)

Intentions

TABLE 2

Precision, Recall, and* F *Score for the Referential Intentions Found by Each Model, Using the Lexicons Scored in Table 1

Model	Precision	Recall	*F* score
Association frequency	.27	**.81**	.40
Conditional probability (object|word)	.59	.36	.45
Conditional probability (word|object)	.32	.79	.46
Mutual information	.36	.37	.37
Translation model (object|word)	.57	.41	.48
Translation model (word|object)	.40	.57	.47
Intentional model	**.83**	.45	**.58**
Intentional model (one parameter)	.77	.36	.50

Intentional Model Advantages

- Distinguish between referential / non-referential words.
- Allow for ‘empty’ intentions.
- Model prefers sparse, one-to-one lexicons.

Mutual Exclusivity

- 1-to-1 = Soft support for the concept of *mutual exclusivity*

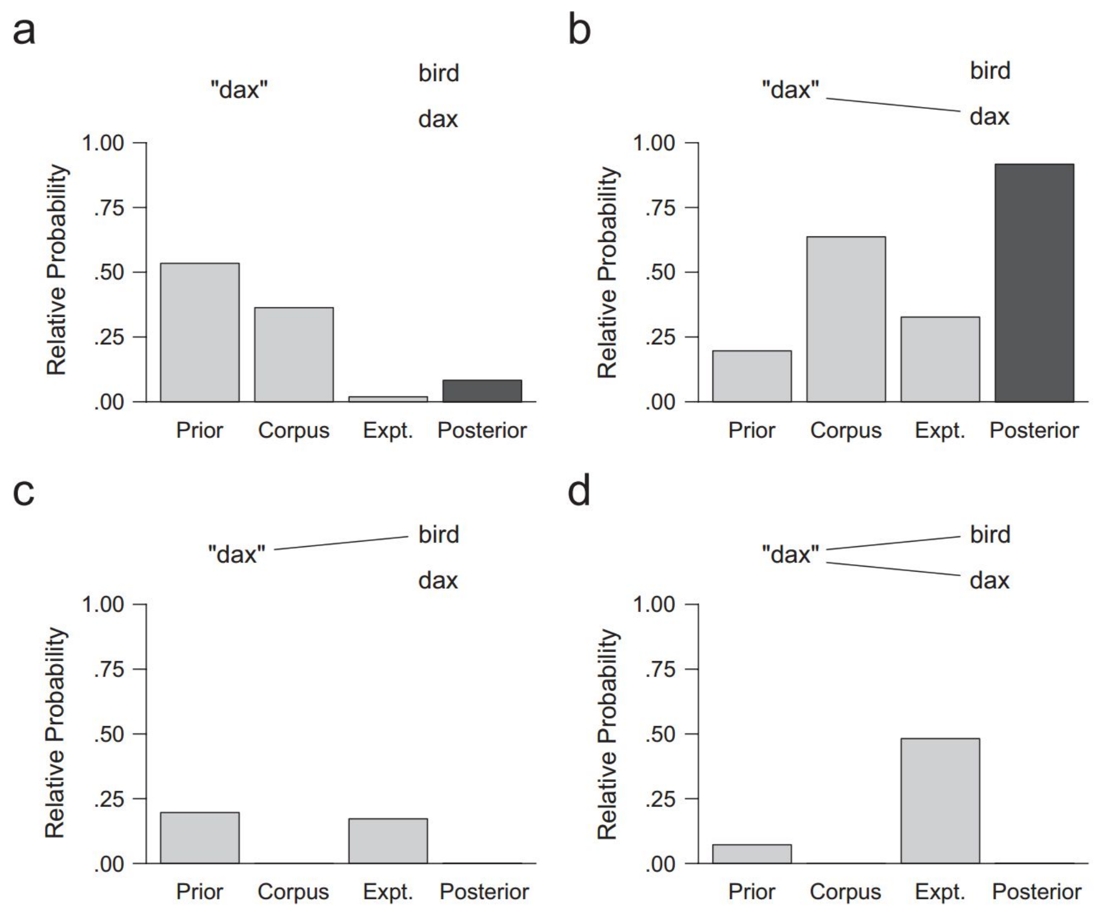

One-Trial Learning

- The model captures one-trial learning well; most of the baseline models do as well, so they’re not too excited about it.

Object Individuation

- Similarly, the model predicts object individuation.
  - model simulation on right

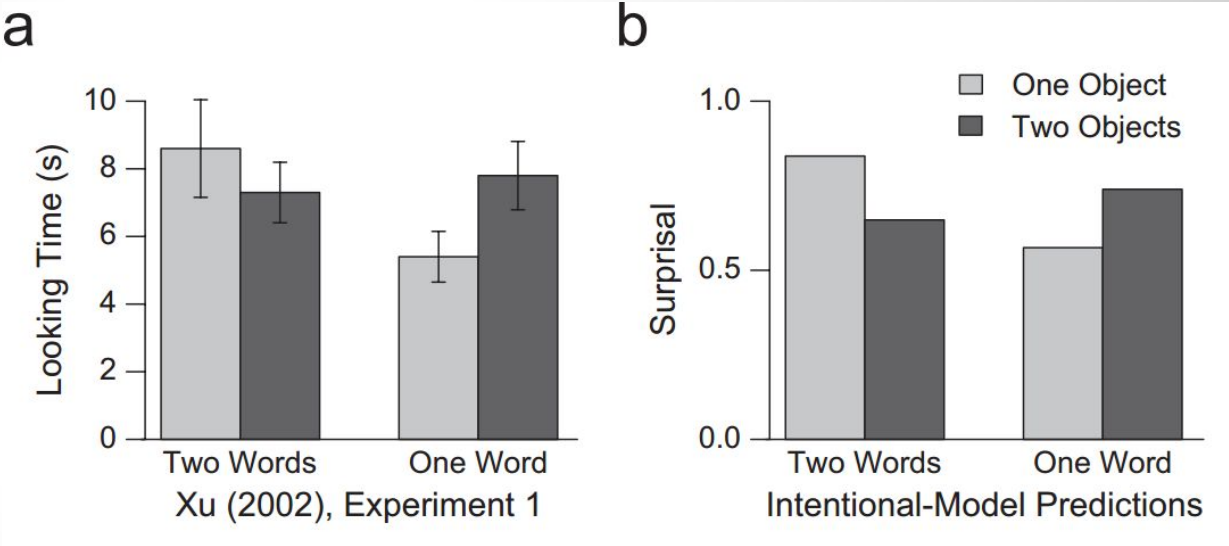

End Main Article

Cognitive Development in LangAc

- LangAc - one equation, 2 unknowns?
  - Maybe one equation, 2 *partially* unknown variables
- Children are strange, and we don’t account for that often enough
  - We may be focusing on too many explanations that speak in ‘adult’ centric terms - children’s observations and understandings may be more unintuitive than we think