Pragmatic description generation with cooperative networks

Will Monroe

CS224U / LINGUIST 188/288

18 May 2016

Implicatures

Work is work.

Examples adapted from Grice (1970)

Implicatures

Will produced a series of sounds that corresponded closely to the tune of "Hey Jude."

Examples adapted from Grice (1970)

Implicatures: Grice's maxims

- "Make your contribution as informative as required [...]
- Do not make your contribution more informative than is required. [...]
- Do not say what you believe to be false. [...]
- Avoid obscurity of expression.
- Avoid ambiguity.
- Be brief (avoid unnecessary prolixity)."

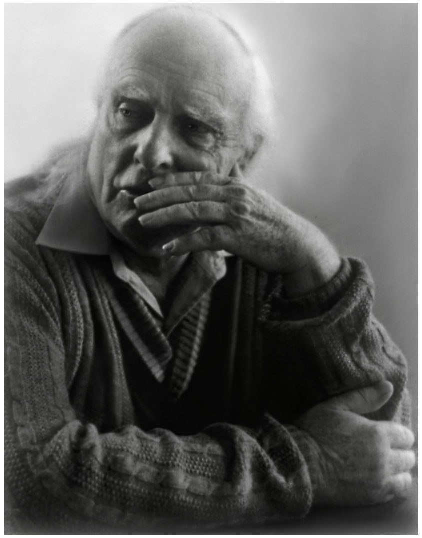

(Grice, 1970)

Image credit: Siobhan Chapman

Implicatures: Grice's maxims

Work is work.

“Make your contribution as informative as required”

Implicatures: Grice's maxims

Will produced a series of sounds that corresponded closely to the tune of "Hey Jude."

"Be brief (avoid unnecessary prolixity)."

Implicatures: Grice's maxims

How do you like my new haircut?

...It's shorter in the back!

"Be relevant."

Image credit: Gage Skidmore

Implicatures

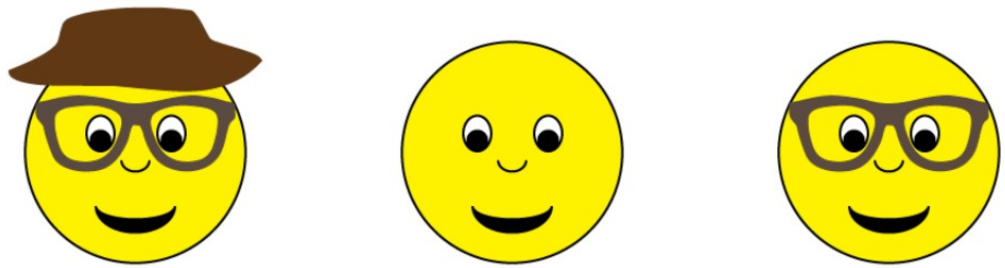

1 2 3

Image credit: Chris Potts

Implicatures

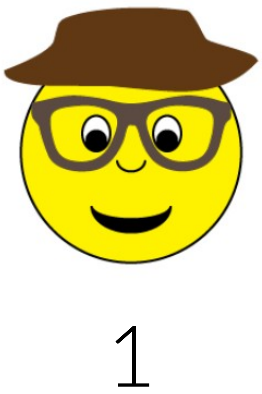

“glasses”

Image credit: Chris Potts

Implicatures

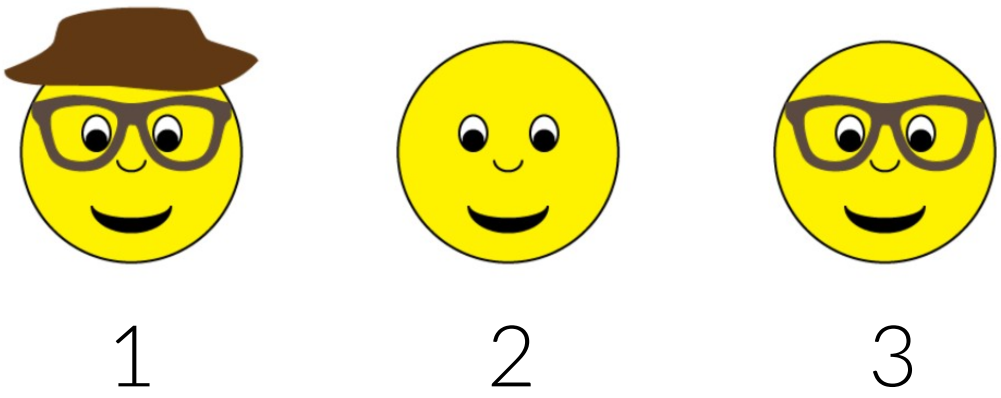

"person"

Image credit: Chris Potts

RSA: Bayesian pragmatic reasoning

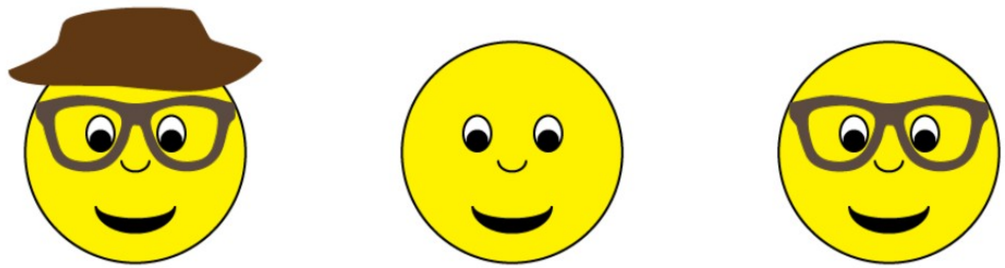

"hat"	✓	✗	✗
"glasses"	✓	✗	✓
"person"	✓	✓	✓

RSA: Bayesian pragmatic reasoning

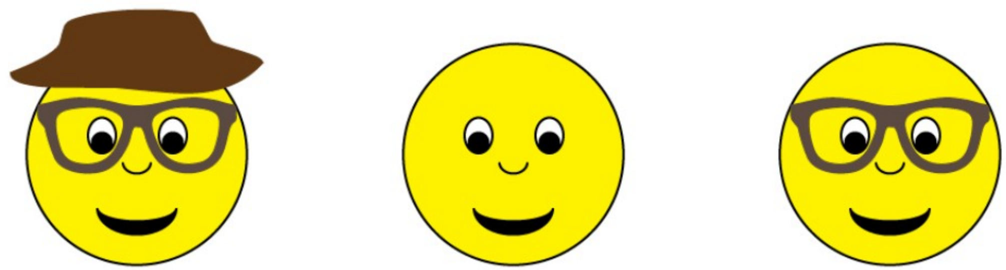

"hat"	1	0	0
"glasses"	0.5	0	0.5
"person"	0.33	0.33	0.33

literal (naive) listener

RSA: Bayesian pragmatic reasoning

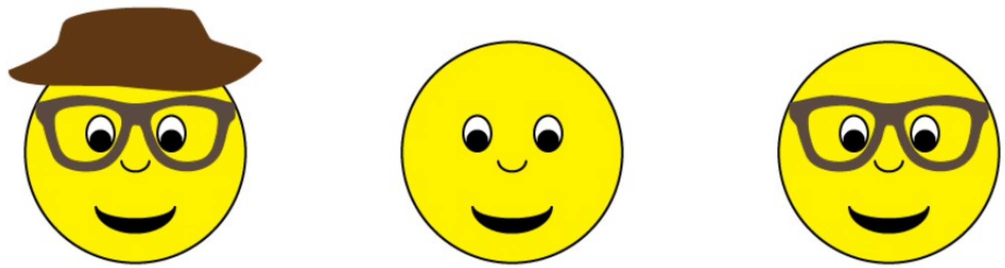

literal (naive) speaker

"hat"	0.33	0	0
"glasses"	0.33	0	0.5
"person"	0.33	1	0.5

RSA: Bayesian pragmatic reasoning

t

m

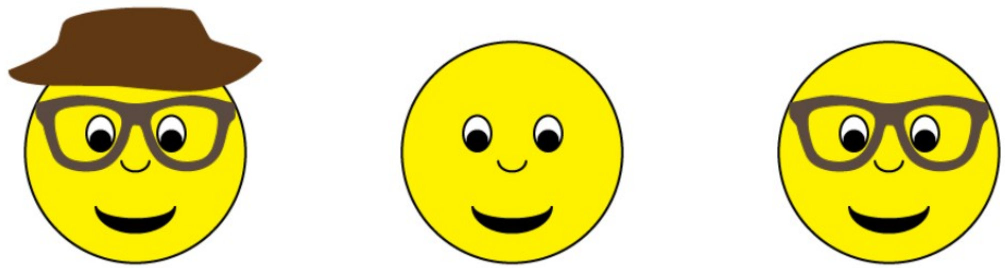

"hat"	1	0	0
"glasses"	0.4	0	0.6
"person"	0.18	0.55	0.27

pragmatic listener

$$L(t|m) = \frac{S(m|t)\,P(t)}{P(m)}$$

(Frank and Goodman, 2012)

Recent RSA conquests

- marked verbosity
  "Will produced a series of sounds..."
  "Clark got the car to stop."
- ignorance implicature
  A: "Does Barb live in Moscow?"
  B: "She lives in Russia..."
- metaphor
  "She's such a princess."
- hyperbole
  "a seven-million-dollar cup of coffee"

lexical uncertainty
(Smith et al. 2013;
Bergen et al. 2014)

question under
discussion
(Kao et al. 2014a/b)

Based on a slide by Will Hamilton

Two obstacles

1. Hand-written semantics

$$S \rightarrow C, \quad S \rightarrow \neg C$$
$$C \rightarrow m\,L\,s, \quad C \rightarrow m, \quad C \rightarrow s$$
$$L \rightarrow \vee, \quad L \rightarrow \wedge$$

$$[\![m]\!] = \{\{\text{Mary}\}, \{\text{Mary}, \text{Sue}\}\}$$
$$[\![s]\!] = \{\{\text{Sue}\}, \{\text{Mary}, \text{Sue}\}\}$$

$$\left(\begin{array}{ll} [\![\text{one}]\!] & = \{1,2,3\} \\ [\![\text{two}]\!] & = \{2,3\} \\ [\![\text{three}]\!] & = \{3\} \\ [\![\text{one or two}]\!] & = \{1,2,3\} \\ [\![\text{two or three}]\!] & = \{2,3\} \\ [\![\text{one or three}]\!] & = \{1,2,3\} \\ [\![\text{one or two or three}]\!] & = \{1,2,3\} \end{array} \right)$$

Can we learn from examples?

(Bergen et al., 2015)

Attribute Selection for REG

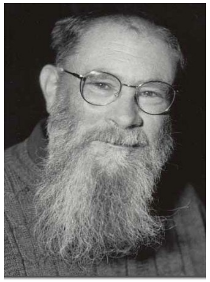  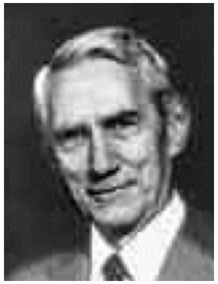

✓ *person*

✓ *beard*

glasses

¬tie

Attribute Selection for REG

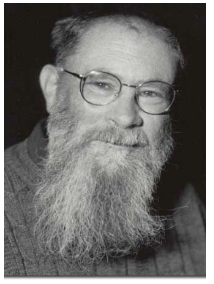

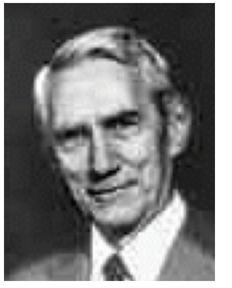

∅

person

beard

glasses

¬tie

person, beard

person, glasses

person, ¬tie

glasses, beard

...

The TUNA Corpus

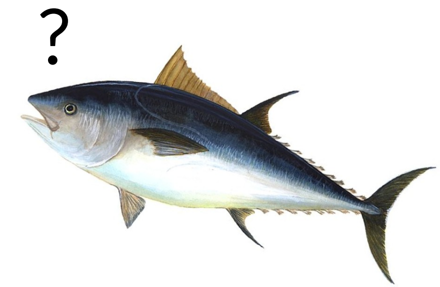

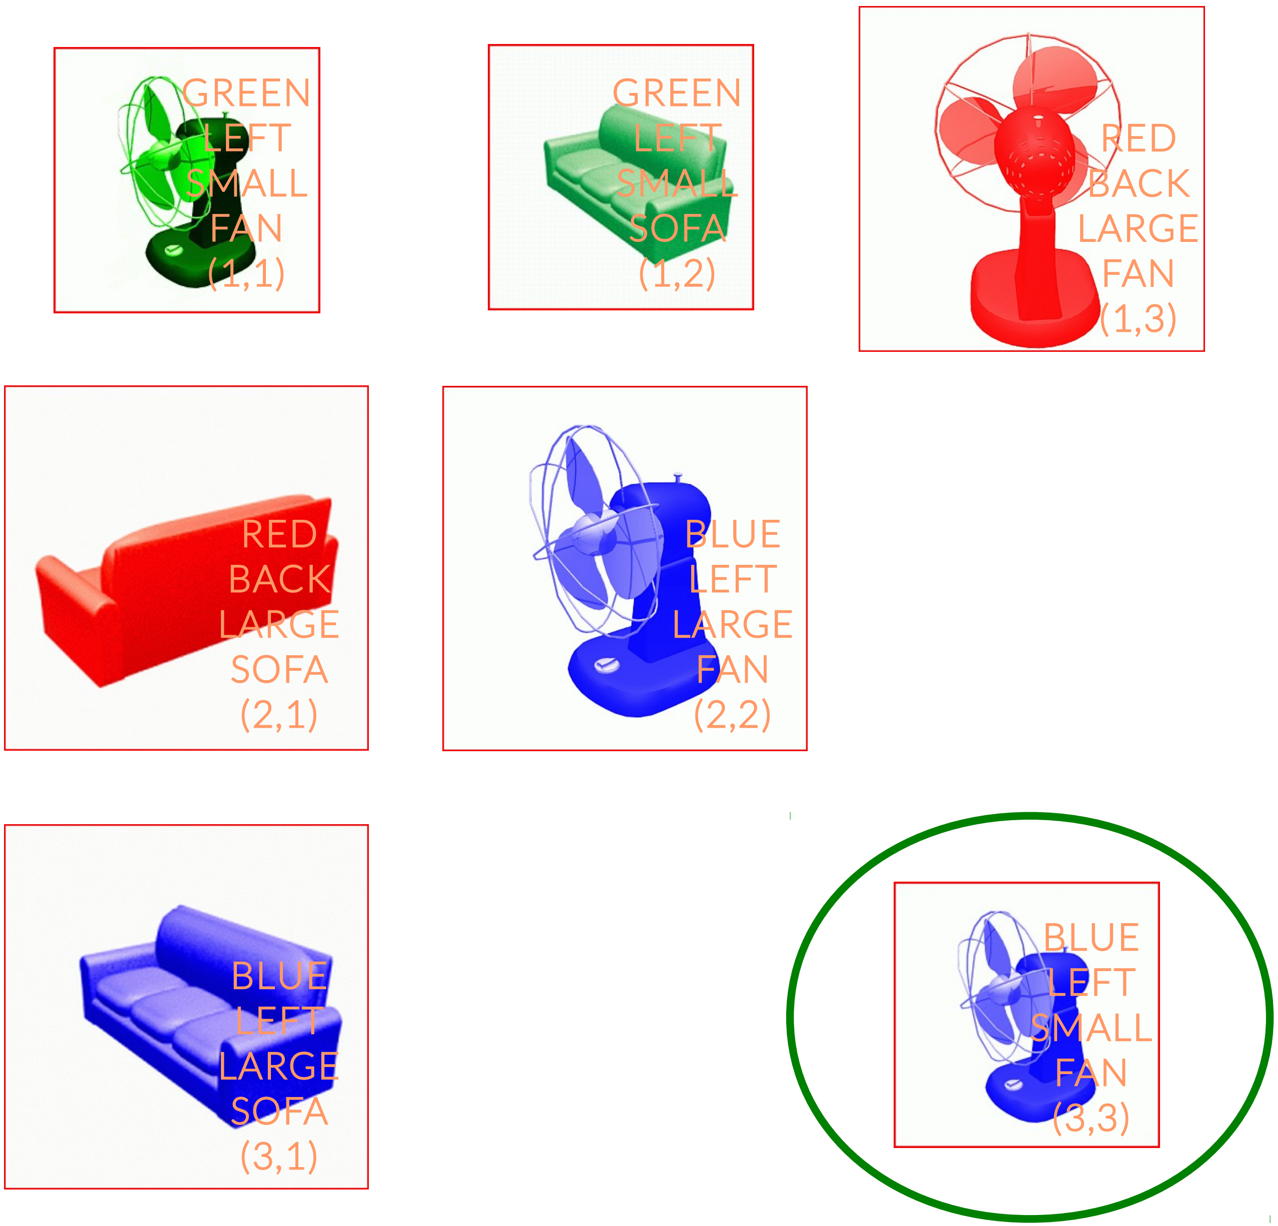

Human utterance:
“**blue fan small**”

Attributes:
blue, *fan*, *small*

(van Deemter et al., 2006; Gatt et al., 2007)

Learning through RSA

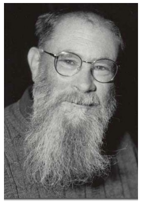

“beard”
“guy with the beard”
“guy with glasses”
...

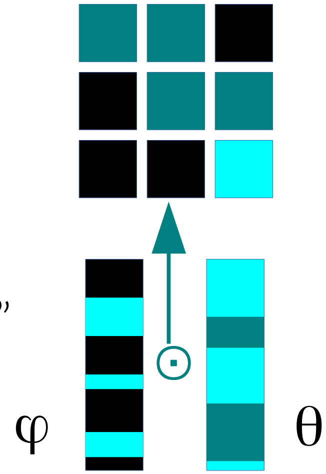

$$S_0(m|t,\theta) \propto \exp\left[\theta^T \varphi(t,m)\right]$$

(Monroe and Potts, 2015)

Learning through RSA

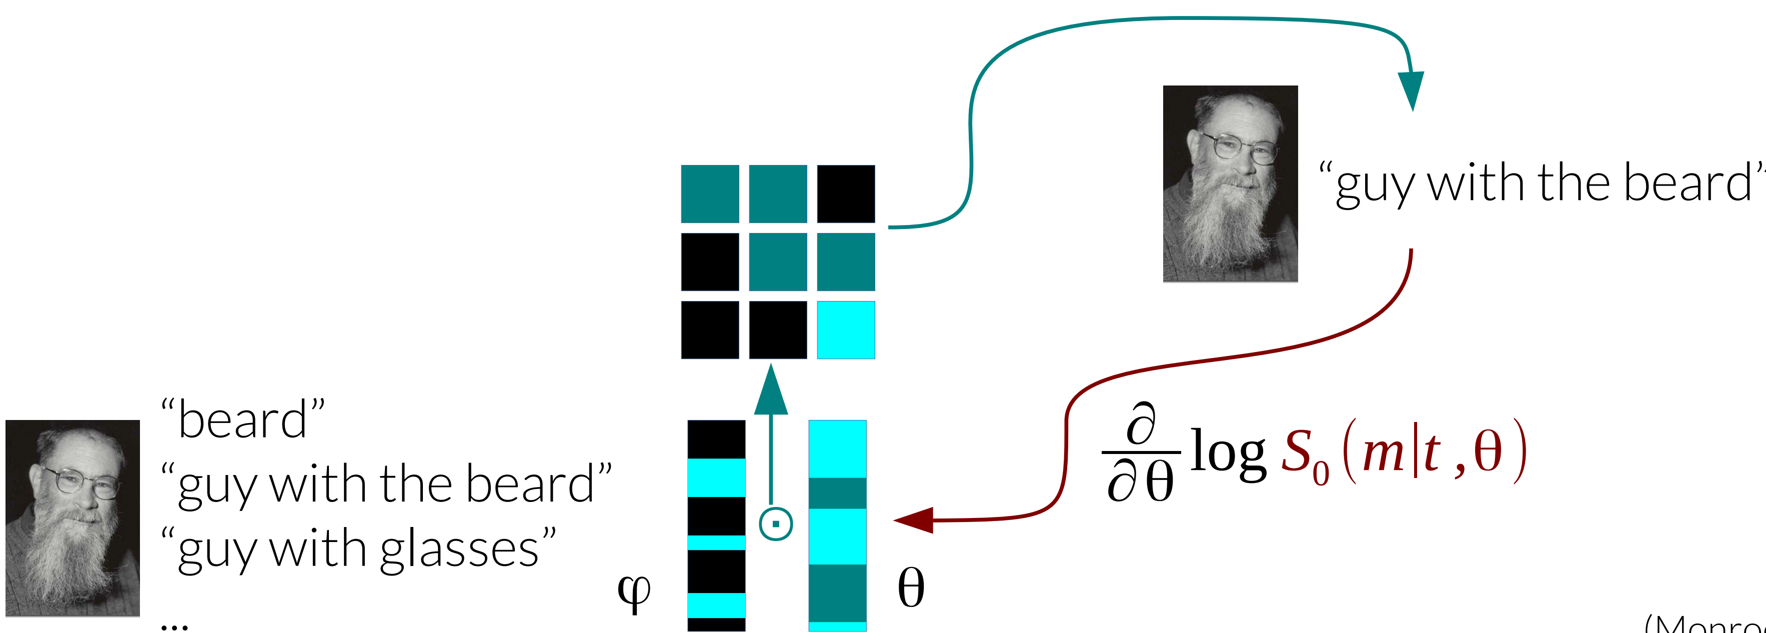

(Monroe and Potts, 2015)

Learning through RSA

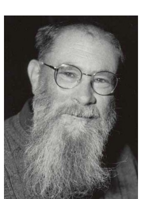

“beard”
“guy with the beard”
“guy with glasses”
...

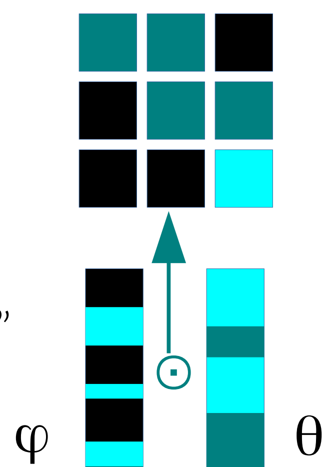

$$S_0(m|t,\theta) \propto \exp[\theta^T \varphi(t,m)]$$

(Monroe and Potts, 2015)

Learning through RSA

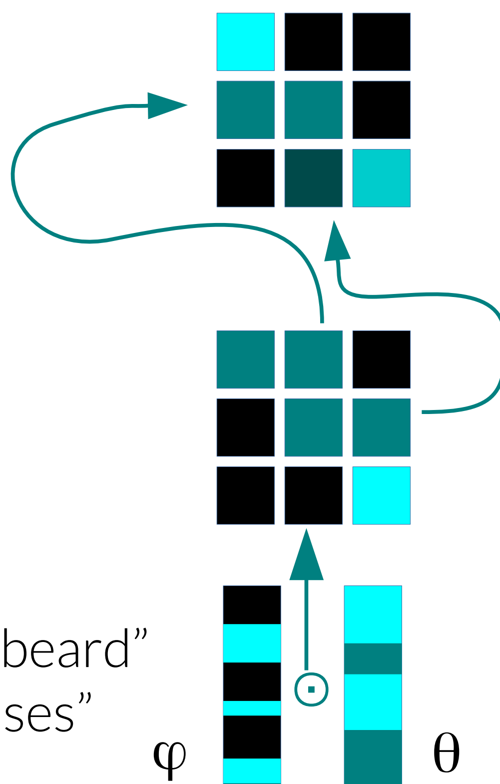

$$L_1(t|m,\theta) \propto S_0(m|t,\theta)$$

$$S_0(m|t,\theta) \propto \exp\left[\theta^T \varphi(t,m)\right]$$

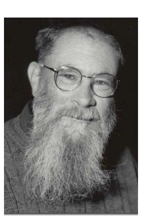

"beard"
"guy with the beard"
"guy with glasses"
...

(Monroe and Potts, 2015)

Learning through RSA

$$S_1(m|t,\theta) \propto L_1(t|m,\theta)$$

$$L_1(t|m,\theta) \propto S_0(m|t,\theta)$$

$$S_0(m|t,\theta) \propto \exp\left[\theta^T \varphi(t,m)\right]$$

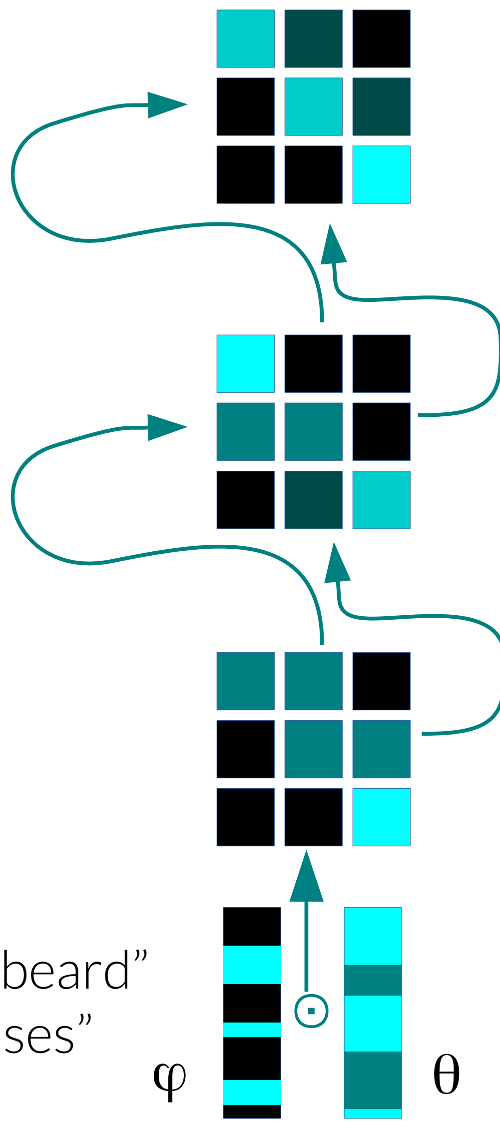

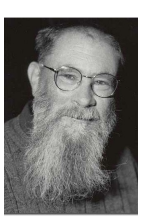

"beard"
"guy with the beard"
"guy with glasses"
...

(Monroe and Potts, 2015)

Learning through RSA

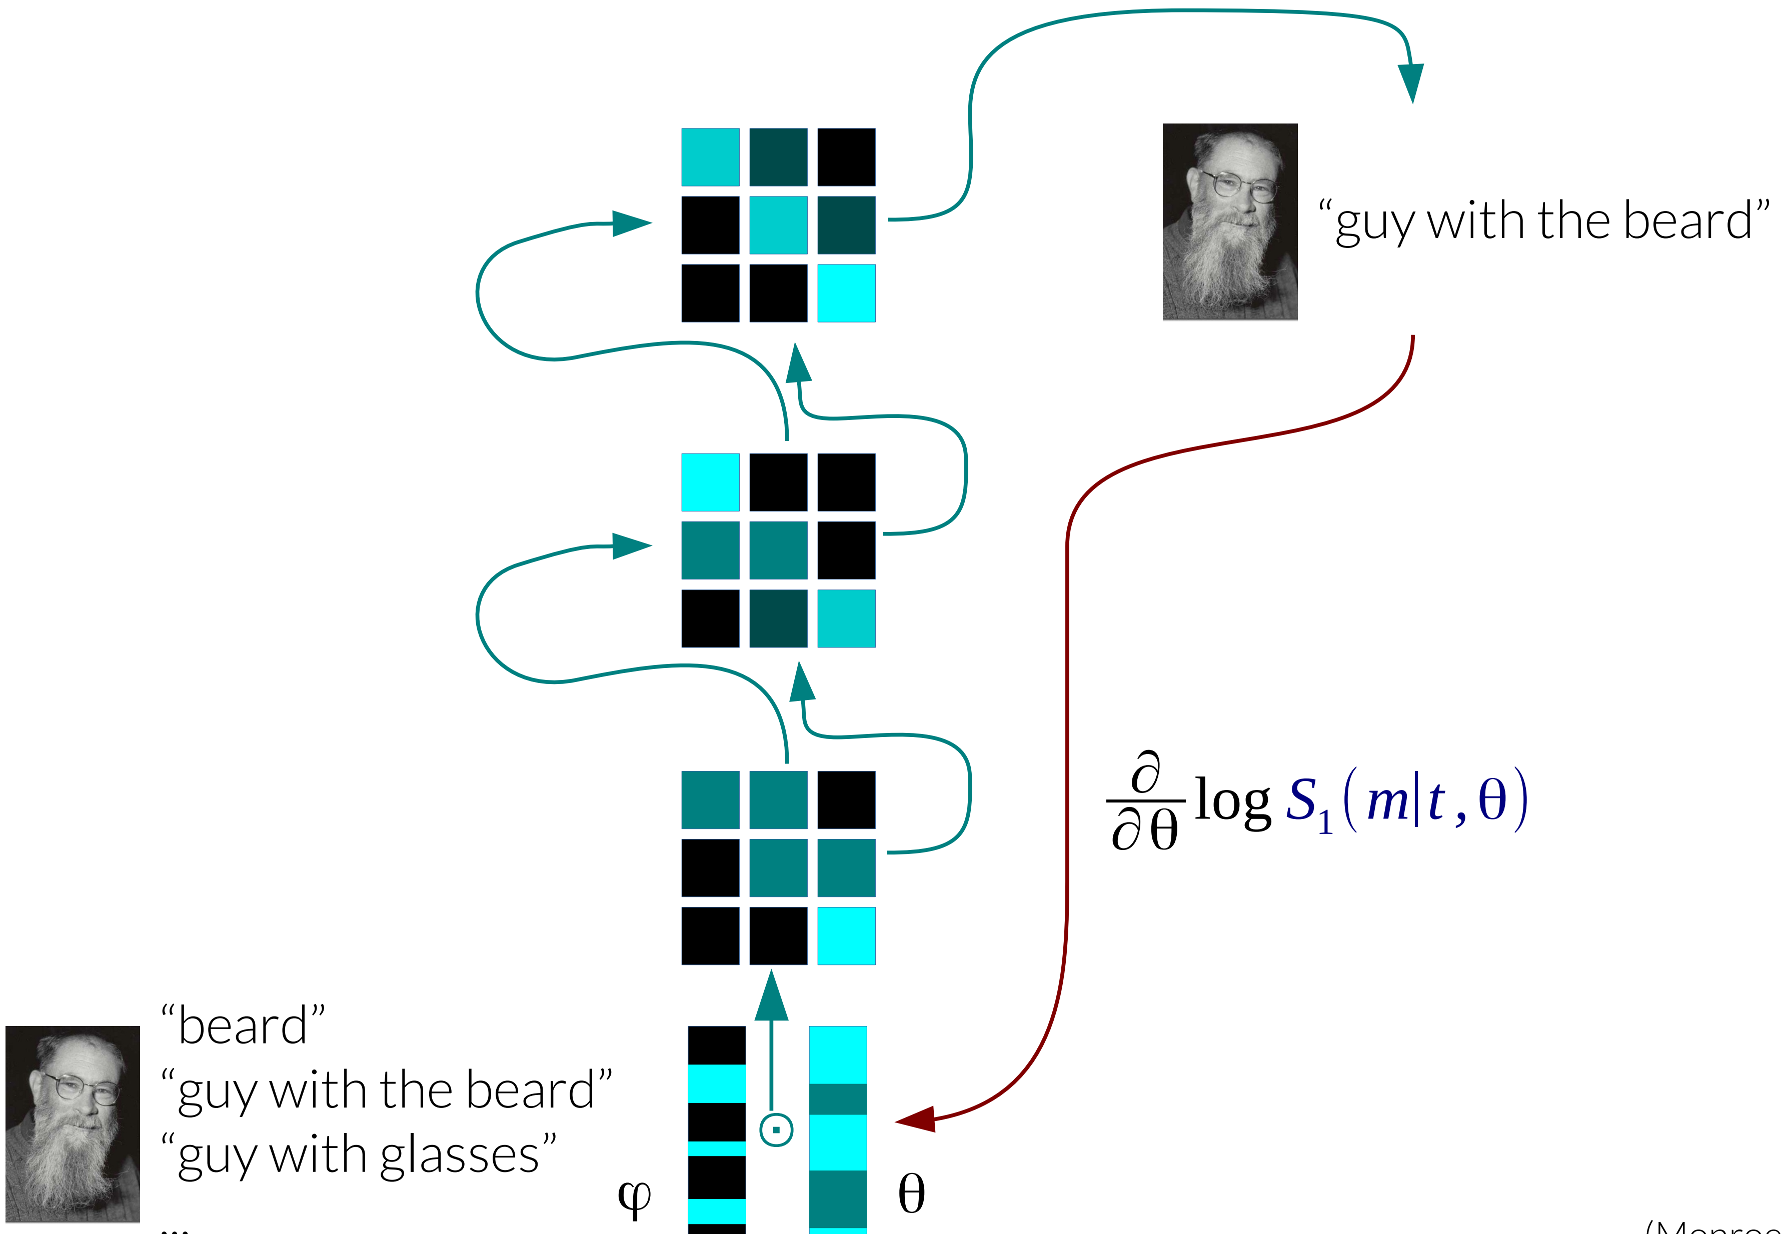

(Monroe and Potts, 2015)

Goals

Features

Avoid hand-built lexicon

Learn human quirks

Goals	Features
Avoid hand-built lexicon	➢ cross-product
Learn human quirks	➢ generation features

Goals

Avoid hand-built lexicon

Learn human quirks

Features

- cross-product

  BLUE → *blue*, BLUE → *fan*, ...

- generation features

Goals

Avoid hand-built lexicon

Learn human quirks

- people overproduce colors

Features

- cross-product

  BLUE → *blue*, BLUE → *fan*, ...

- generation features

  attribute type

  {*color*}

Goals	Features
Avoid hand-built lexicon	➢ cross-product BLUE → *blue*, BLUE → *fan*, ...
Learn human quirks • people overproduce colors • attributes fit into a hierarchy	➢ generation features attribute type {*color*} attribute pairs (pos/neg) {*type*}+{*color*}, {*color*}+¬{*size*}

Goals

Avoid hand-built lexicon

Learn human quirks

- people overproduce colors
- attributes fit into a hierarchy
- certain utterance lengths are preferred

Features

- cross-product

  BLUE → *blue*, BLUE → *fan*, ...

- generation features

  attribute type

  {*color*}

  attribute pairs (pos/neg)

  {*type*}+{*color*}, {*color*}+¬{*size*}

  message size

  {*2 attrs*}, {*3 attrs*}, ...

Example: dataset

train

person, glasses

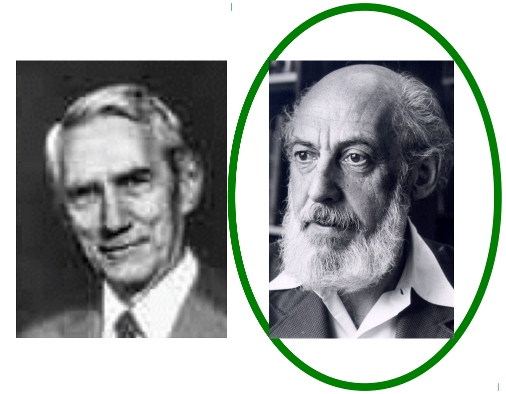

person, beard

test

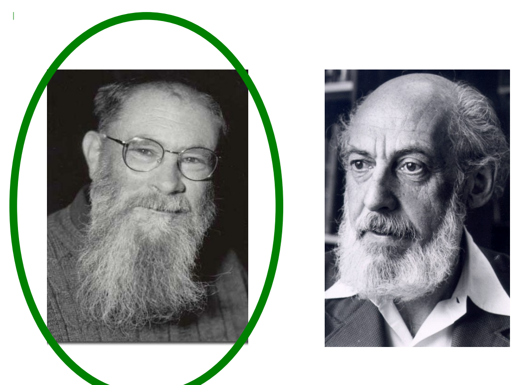

?

Example: distributions

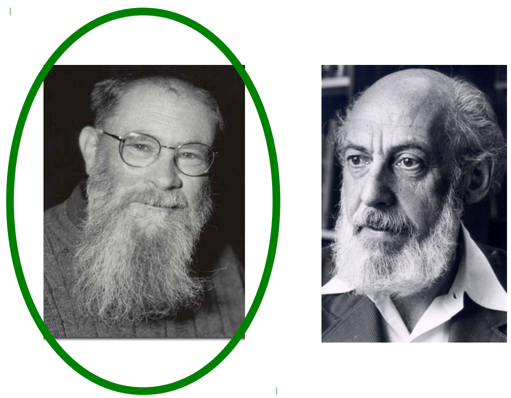

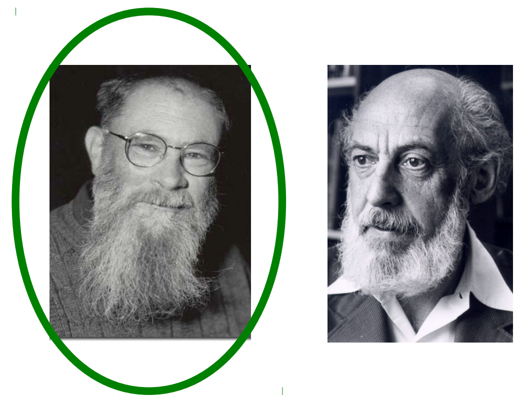

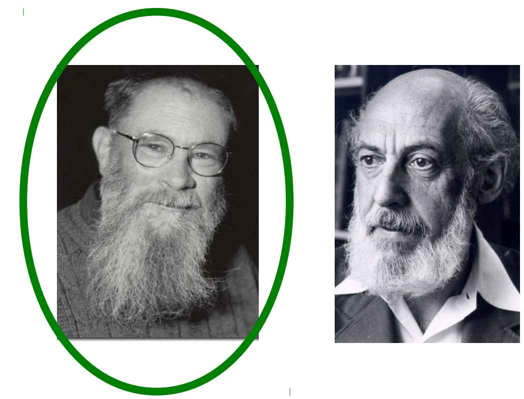

Ø
person
glasses
beard
person, glasses
person, beard
glasses, beard
all

RSA $\quad S_0 \quad S_1$

Example: distributions

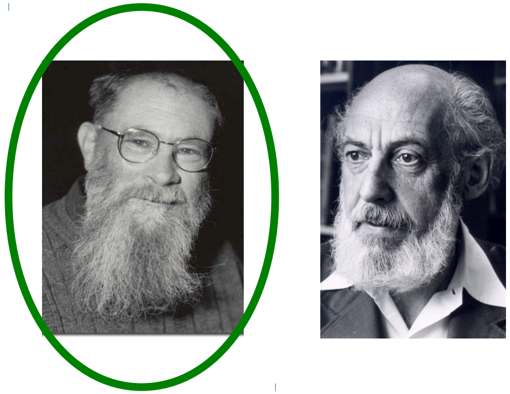

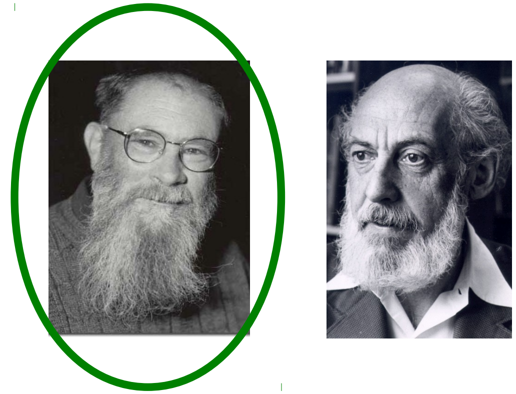

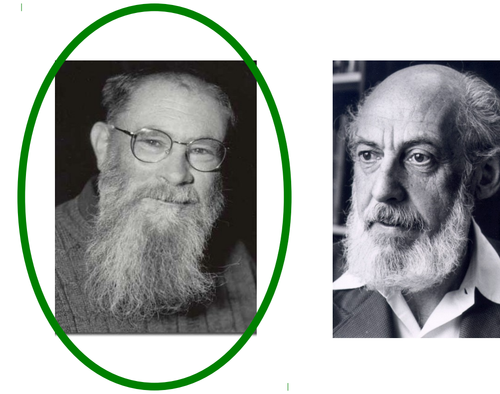

Ø	.08	.25
person	.08	.25
glasses	$\boxed{.17}$	0
beard	.08	.25
person, glasses	$\boxed{.17}$	0
person, beard	.08	.25
glasses, beard	.17	0
all	.17	0

RSA S_0 S_1

Example: distributions

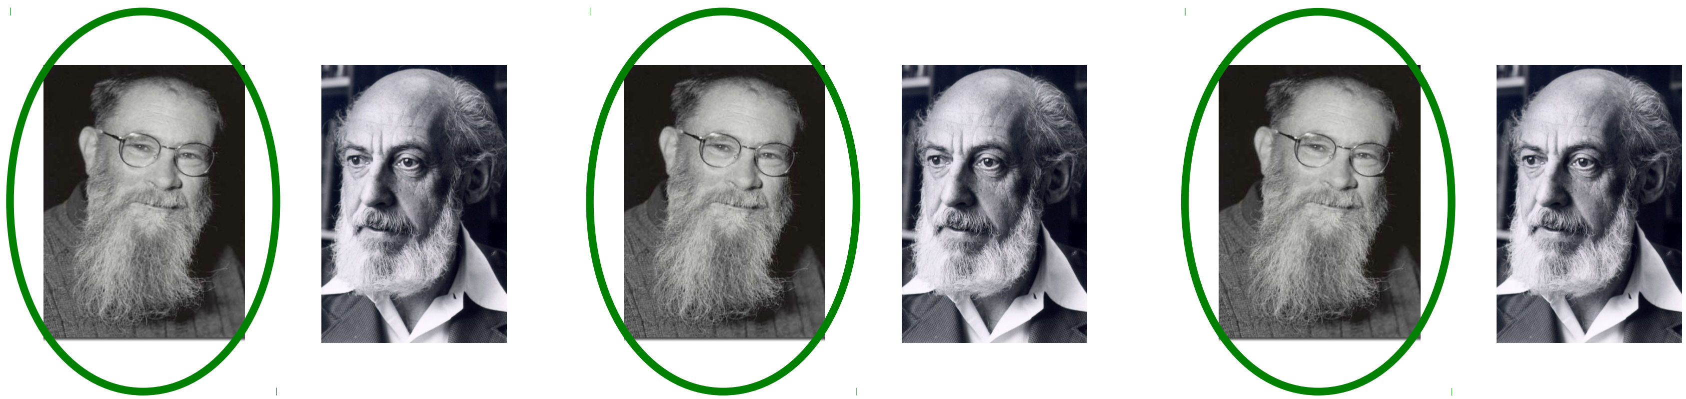

	RSA		S_0		S_1	
Ø	.08	.25	.03	.00		
person	.08	.25	.22	.10		
glasses	$\boxed{.17}$	0	.03	.00		
beard	.08	.25	.03	.04		
person, glasses	$\boxed{.17}$	0	$\boxed{.22}$	.01		
person, beard	.08	.25	$\boxed{.22}$	.74		
glasses, beard	.17	0	.03	.00		
all	.17	0	.22	.10		

Example: distributions

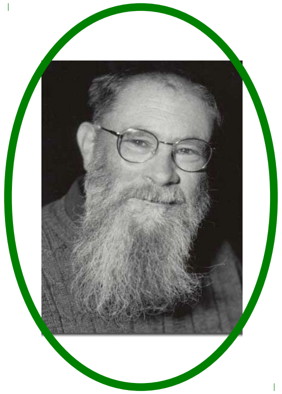 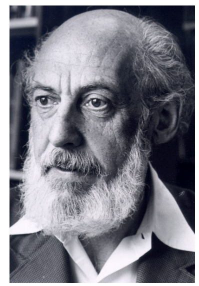 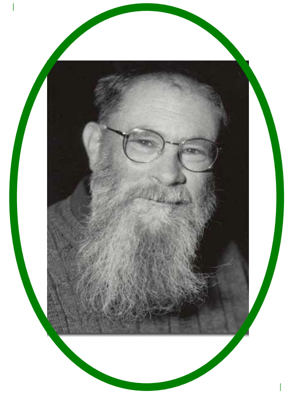 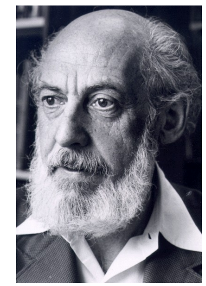 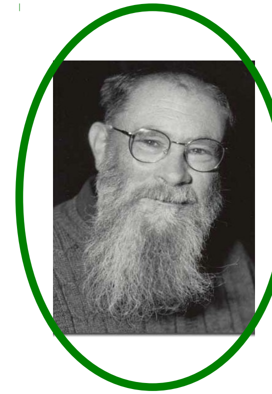 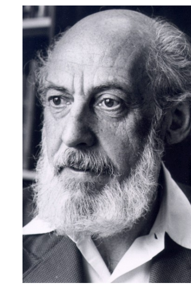

Ø	.08	.25	.03	.00	.10	.11
person	.08	.25	.22	.10	.16	.13
glasses	.17	0	.03	.00	.11	.07
beard	.08	.25	.03	.04	.08	.17
person, glasses	.17	0	.22	.01	.18	.08
person, beard	.08	.25	.22	.74	.12	.19
glasses, beard	.17	0	.03	.00	.10	.11
all	.17	0	.22	.10	.16	.11
	RSA		S_0		S_1	

Example: distributions

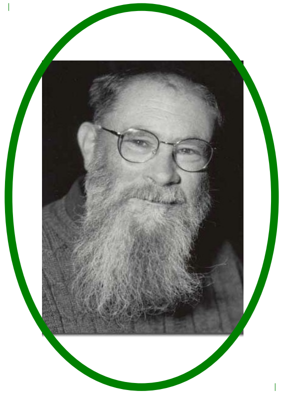 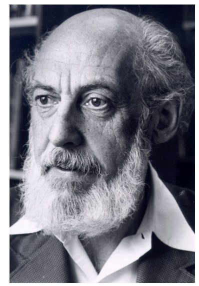 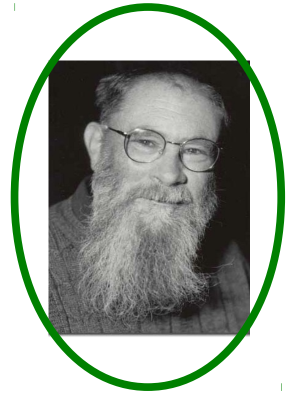 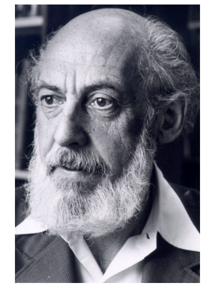 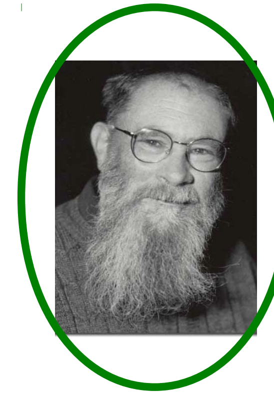 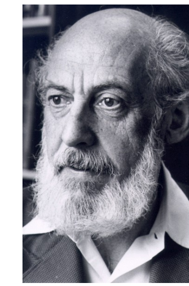

	RSA		S_0		S_1	
Ø	.08	.25	.03	.00	.10	.11
person	.08	.25	.22	.10	.16	.13
glasses	.17	0	.03	.00	.11	.07
beard	.08	.25	.03	.04	.08	.17
person, glasses	.17	0	.22	.01	.18	.08
person, beard	.08	.25	.22	.74	.12	.19
glasses, beard	.17	0	.03	.00	.10	.11
all	.17	0	.22	.10	.16	.11

Experimental results

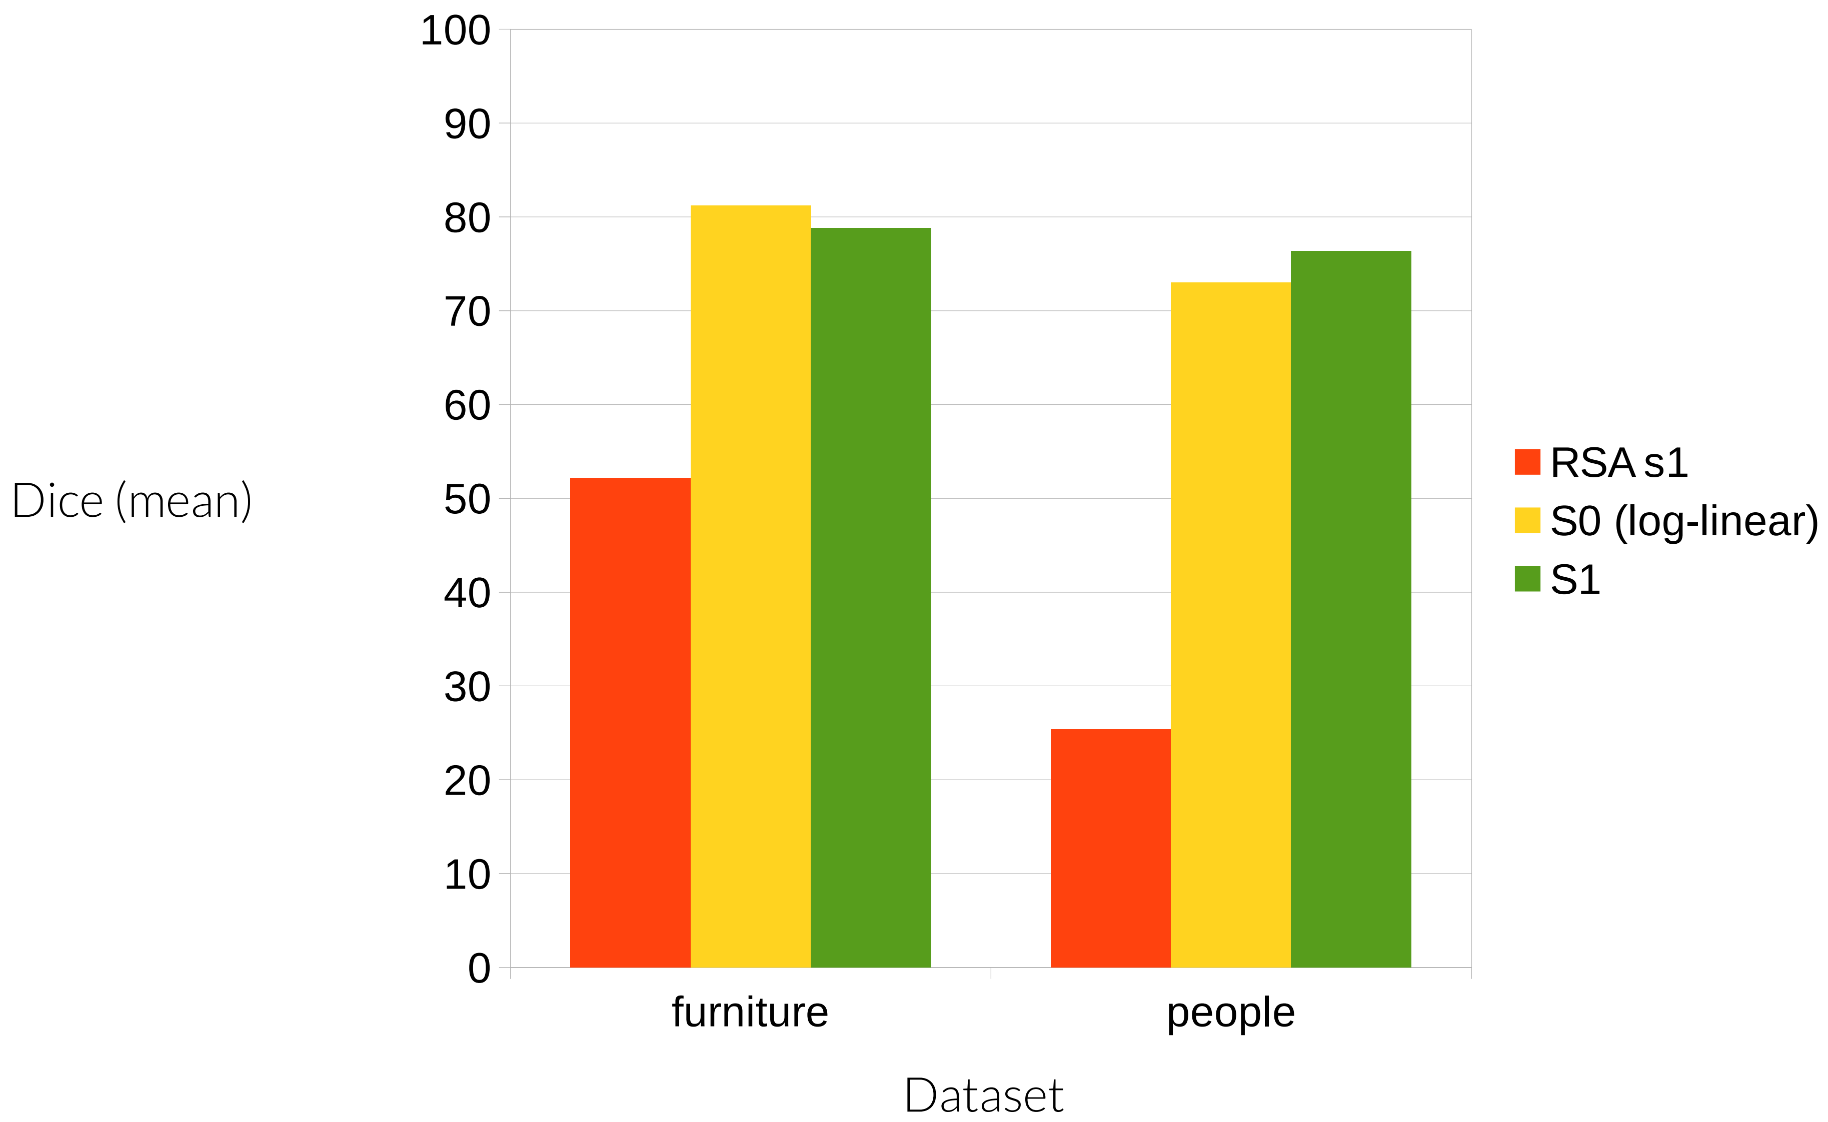

Analysis

TUNA *people* dataset

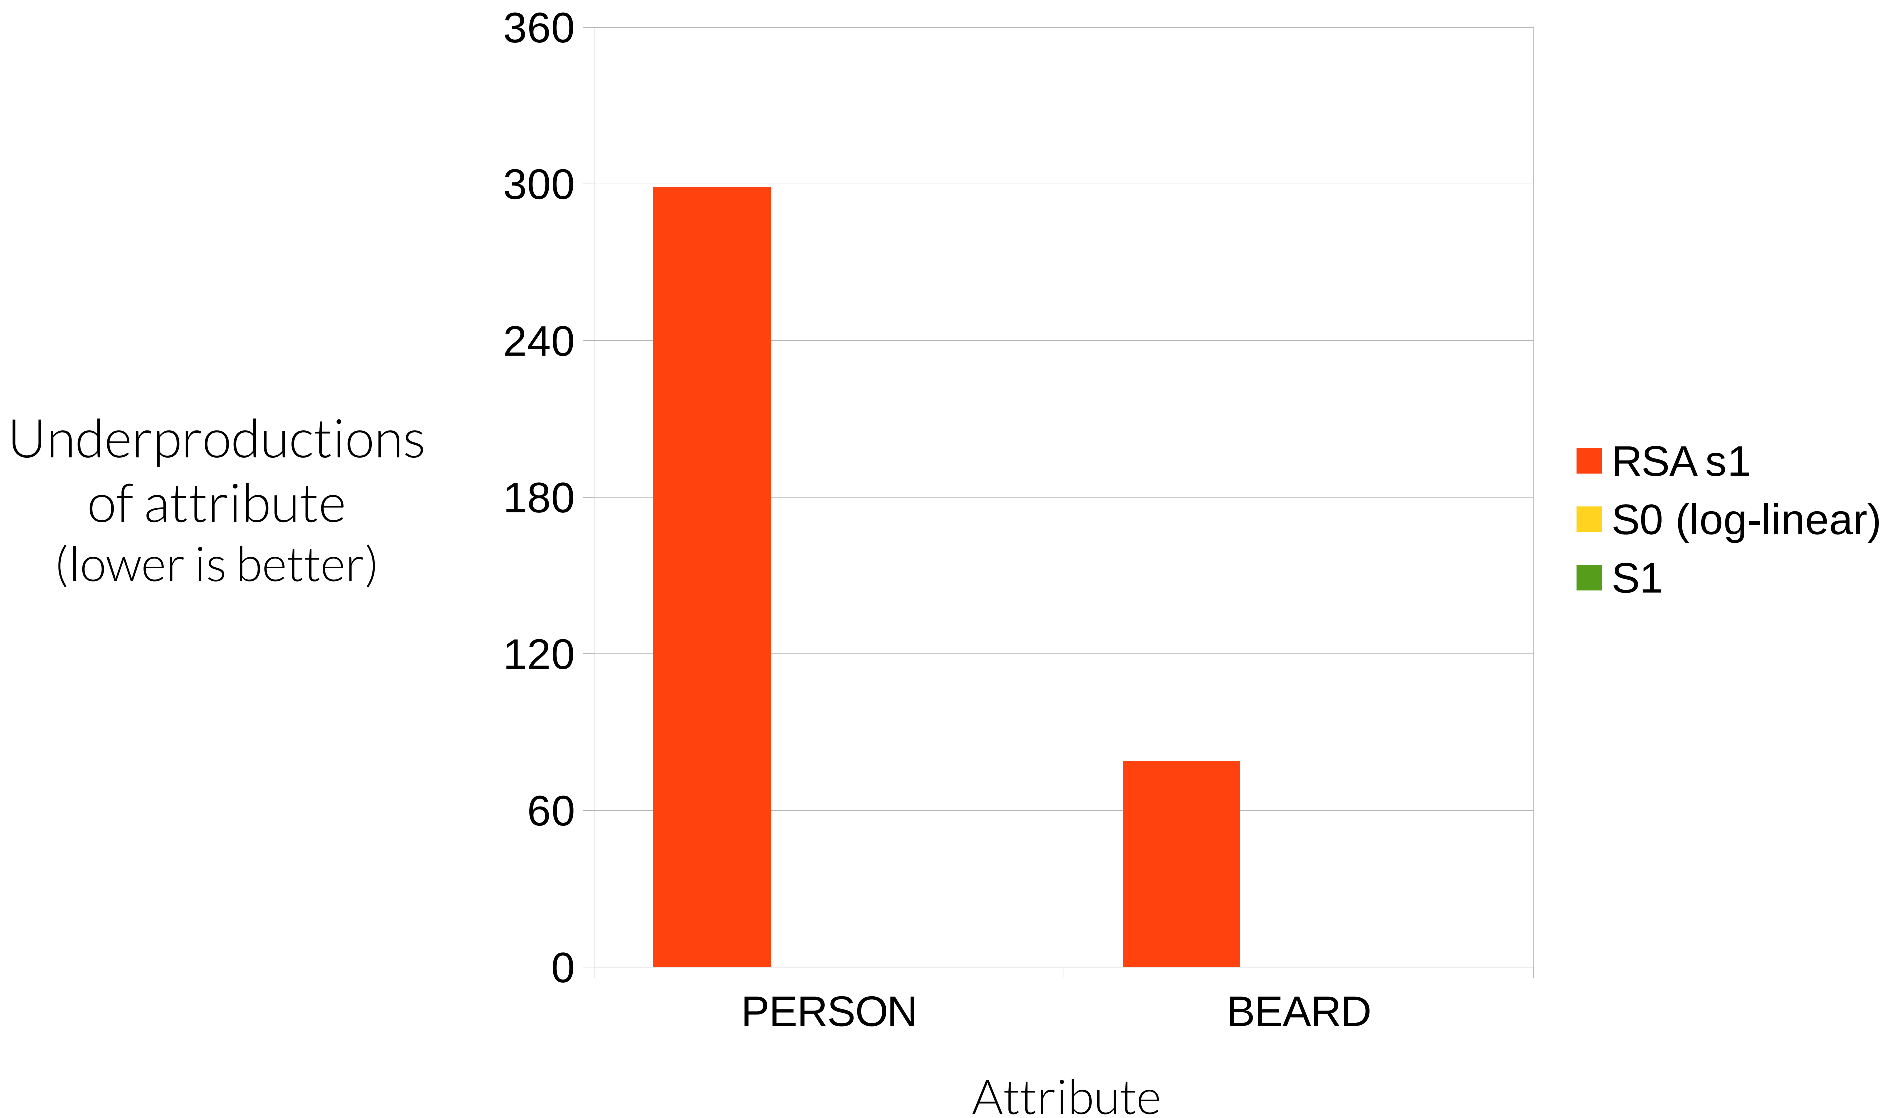

Analysis

TUNA *people* dataset

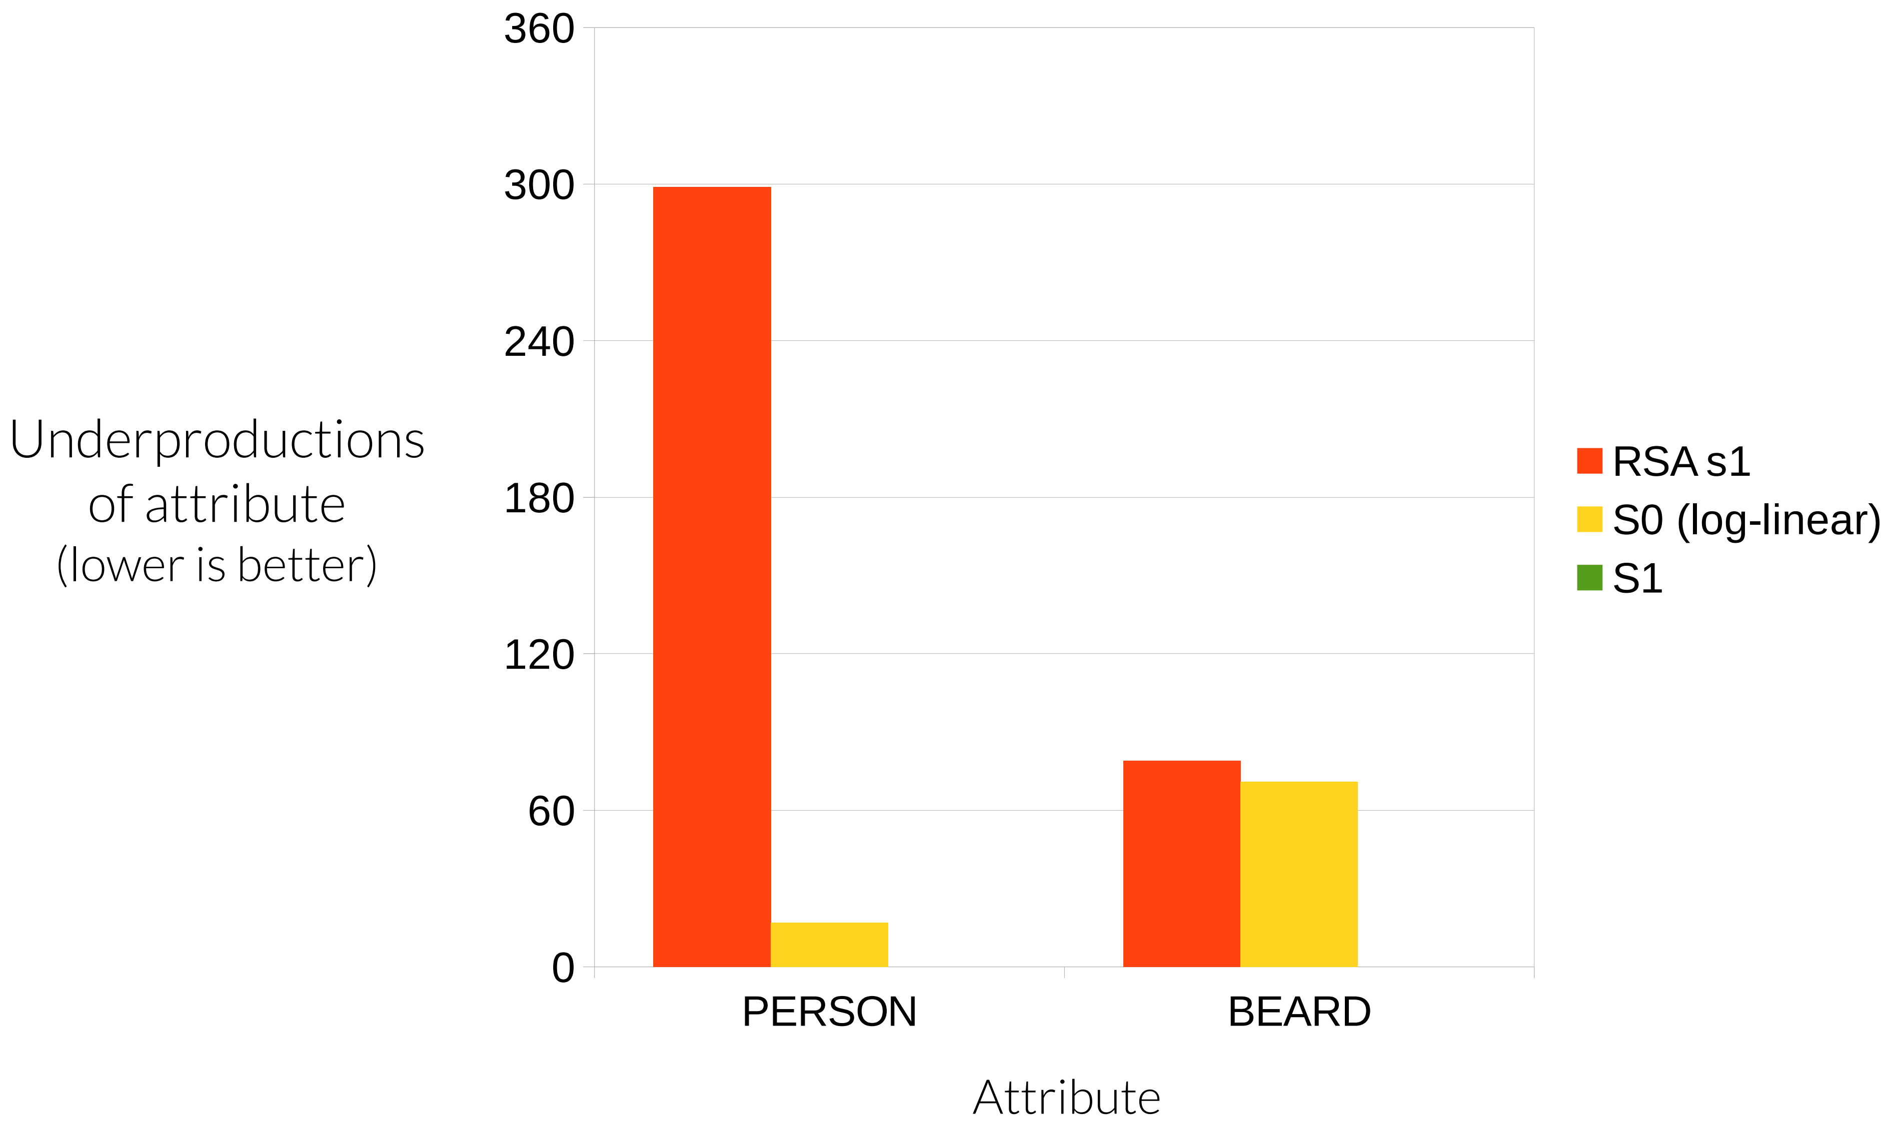

Analysis

TUNA *people* dataset

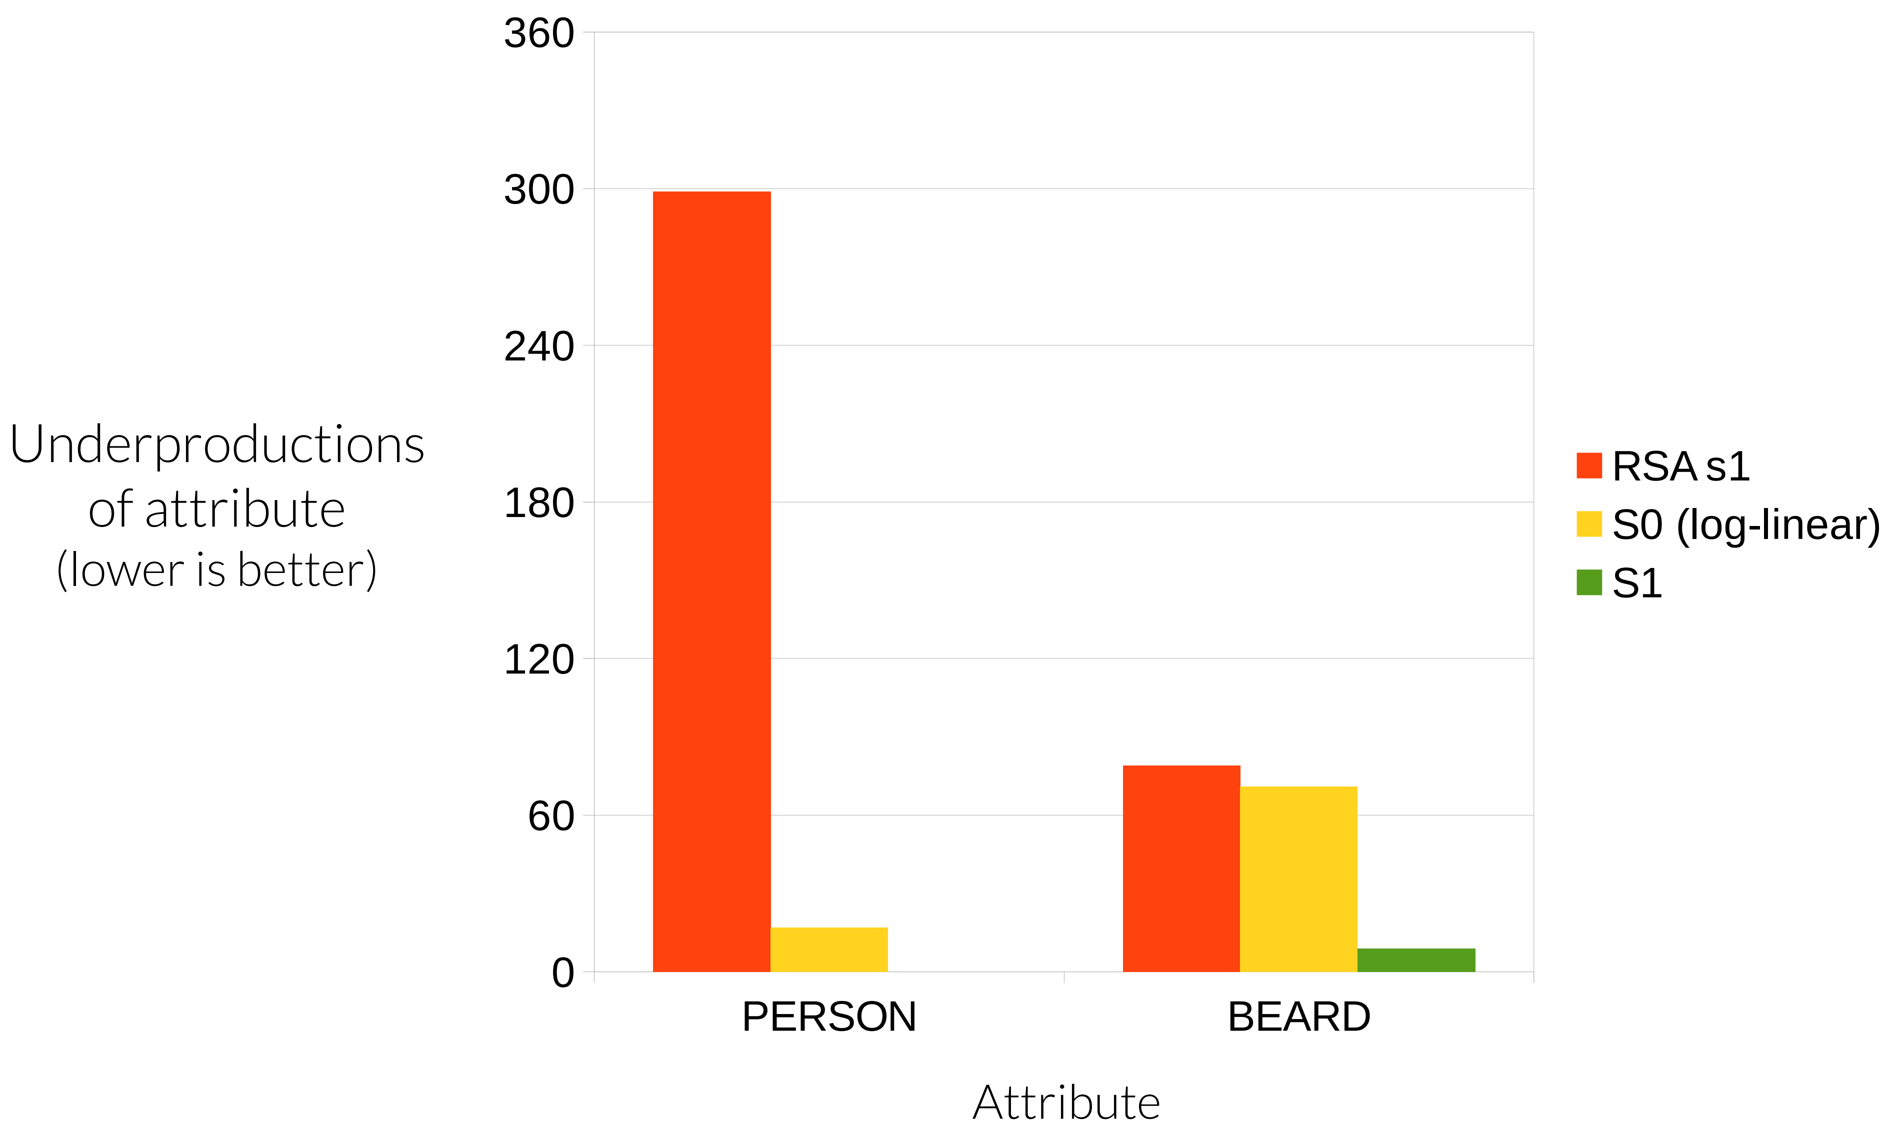

Two obstacles

1. Hand-written semantics

$$S \rightarrow C, \quad S \rightarrow \neg C$$
$$C \rightarrow m L s, \quad C \rightarrow m, \quad C \rightarrow s$$
$$L \rightarrow \vee, \quad L \rightarrow \wedge$$

$$[\![m]\!] = \{\{\text{Mary}\}, \{\text{Mary}, \text{Sue}\}\}$$
$$[\![s]\!] = \{\{\text{Sue}\}, \{\text{Mary}, \text{Sue}\}\}$$

$$\left(\begin{array}{ll} [\![\text{one}]\!] & = \{1,2,3\} \\ [\![\text{two}]\!] & = \{2,3\} \\ [\![\text{three}]\!] & = \{3\} \\ [\![\text{one or two}]\!] & = \{1,2,3\} \\ [\![\text{two or three}]\!] & = \{2,3\} \\ [\![\text{one or three}]\!] & = \{1,2,3\} \\ [\![\text{one or two or three}]\!] & = \{1,2,3\} \end{array}\right)$$

Can we learn from examples?

(Bergen et al., 2015)

Two obstacles

1. Hand-written semantics

$$S \rightarrow C, \quad S \rightarrow \neg C$$
$$C \rightarrow m\,L\,s, \quad C \rightarrow m, \quad C \rightarrow s$$
$$L \rightarrow \vee, \quad L \rightarrow \wedge$$

$$[\![m]\!] = \{\{\text{Mary}\}, \{\text{Mary}, \text{Sue}\}\}$$
$$[\![s]\!] = \{\{\text{Sue}\}, \{\text{Mary}, \text{Sue}\}\}$$

$$\left(\begin{array}{ll} [\![\text{one}]\!] & = \{1,2,3\} \\ [\![\text{two}]\!] & = \{2,3\} \\ [\![\text{three}]\!] & = \{3\} \\ [\![\text{one or two}]\!] & = \{1,2,3\} \\ [\![\text{two or three}]\!] & = \{2,3\} \\ [\![\text{one or three}]\!] & = \{1,2,3\} \\ [\![\text{one or two or three}]\!] & = \{1,2,3\} \end{array}\right)$$

Can we learn from examples? ✔

(Bergen et al., 2015)

Two obstacles

2. Exhaustive enumeration of utterances and worlds

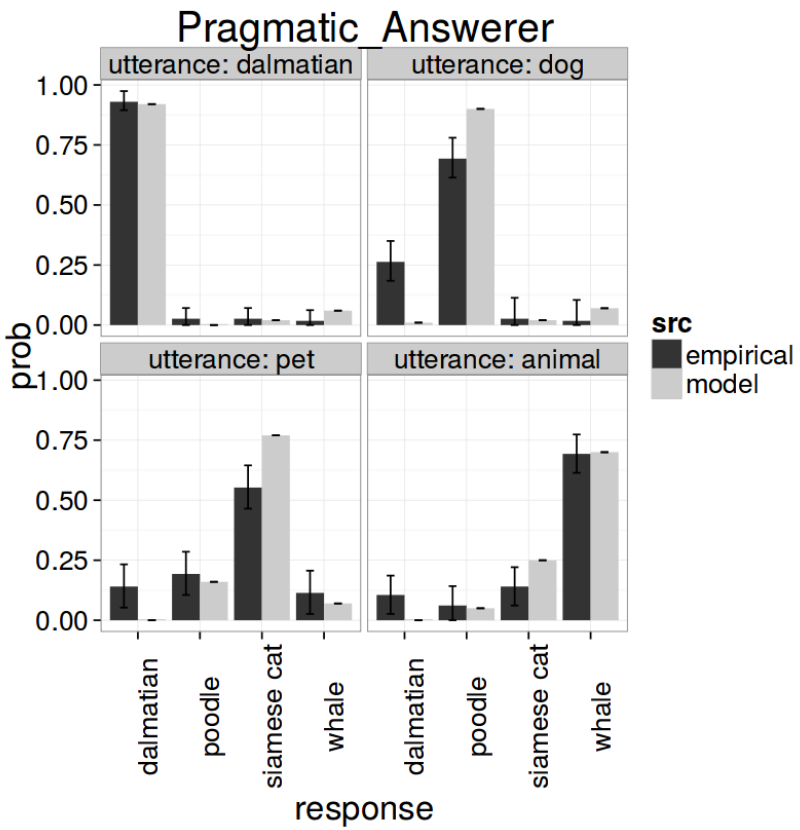

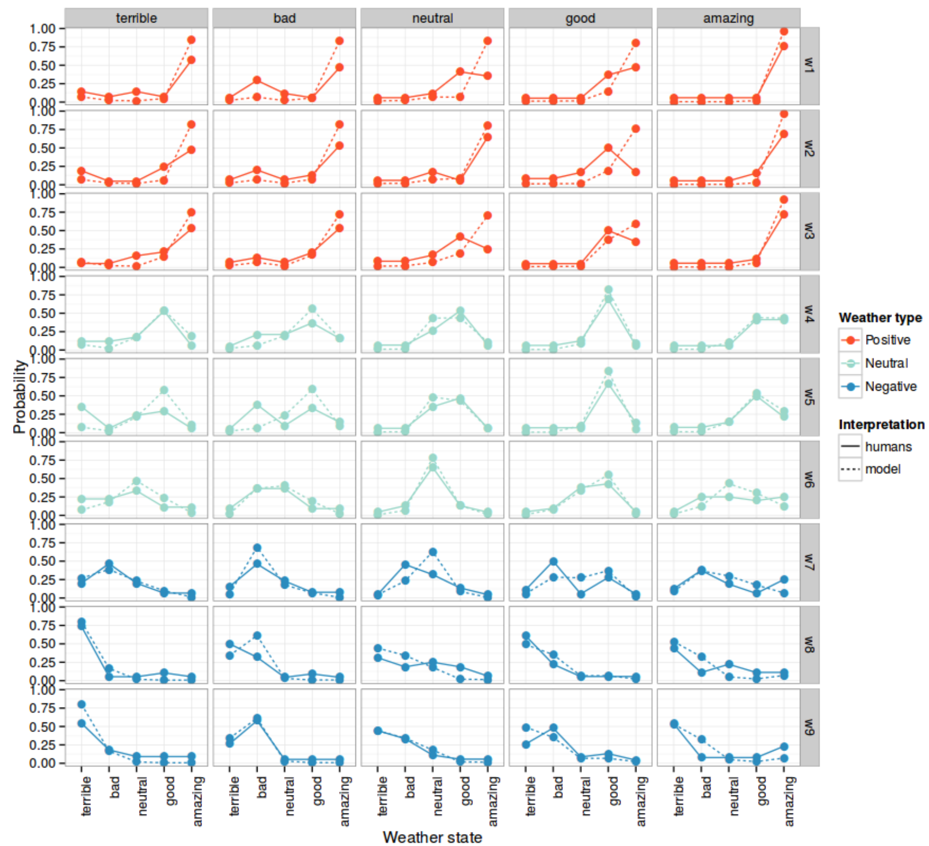

Can we generalize efficiently?

(Hawkins et al., 2015; Kao and Goodman, 2015)

Task: modeling color descriptions

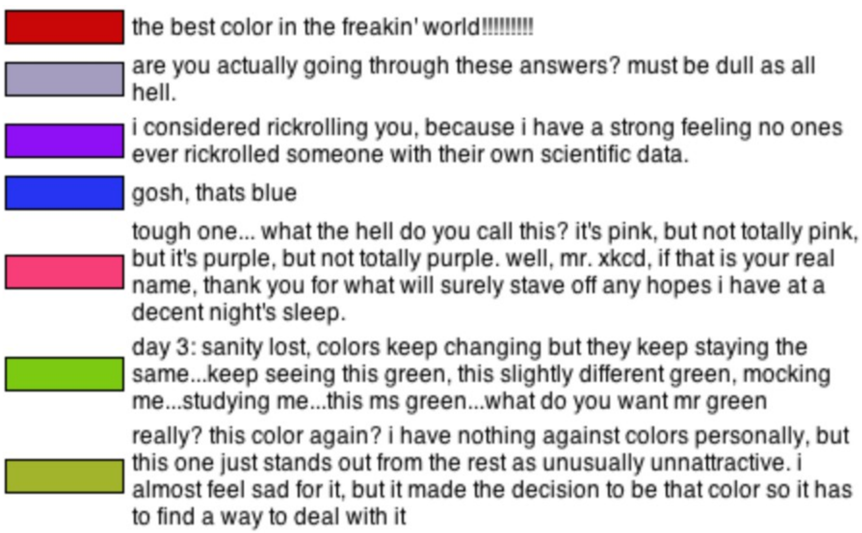

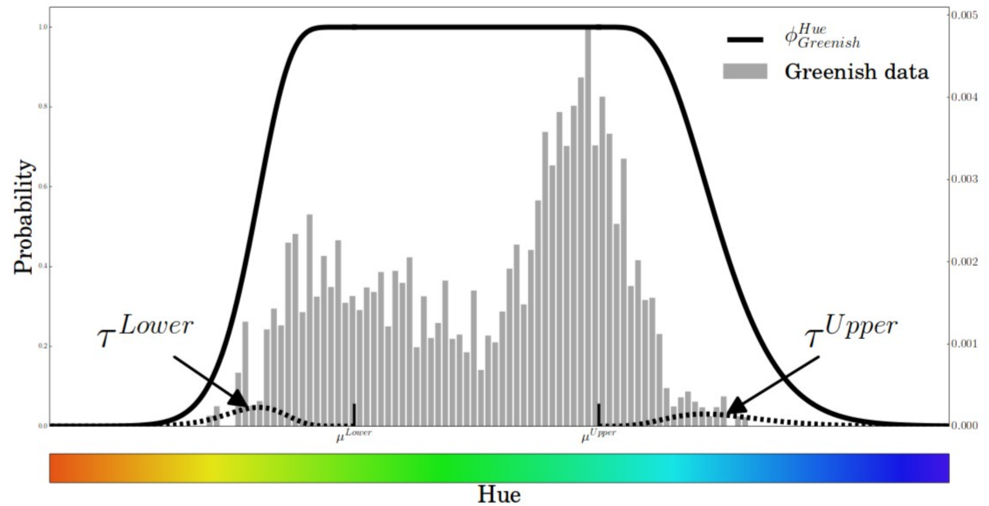

(no relation)

(Munroe, 2010; McMahan and Stone, 2015)

Task: modeling color descriptions

Speaker 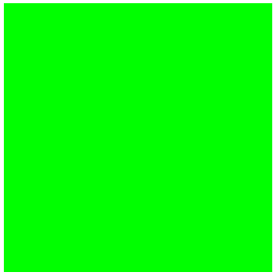 *"neon green"*

Listener *"navy blue"* 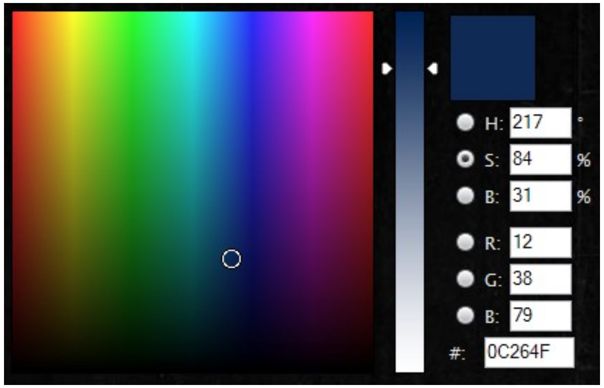

Sequence modeling with RNNs

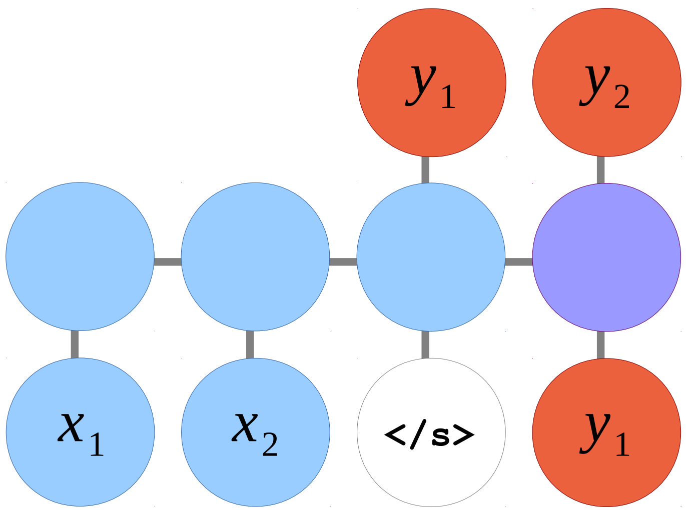

Sequence modeling with RNNs

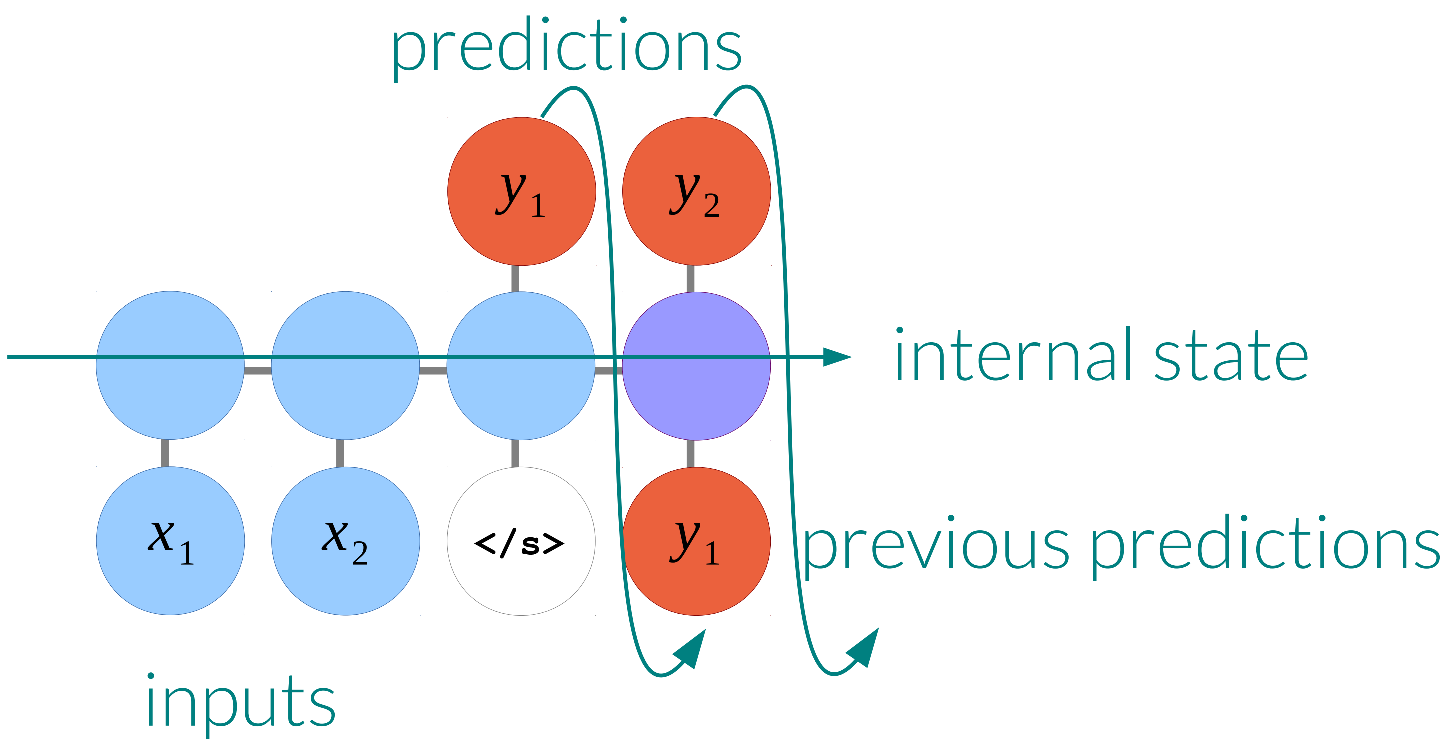

Sequence modeling with RNNs

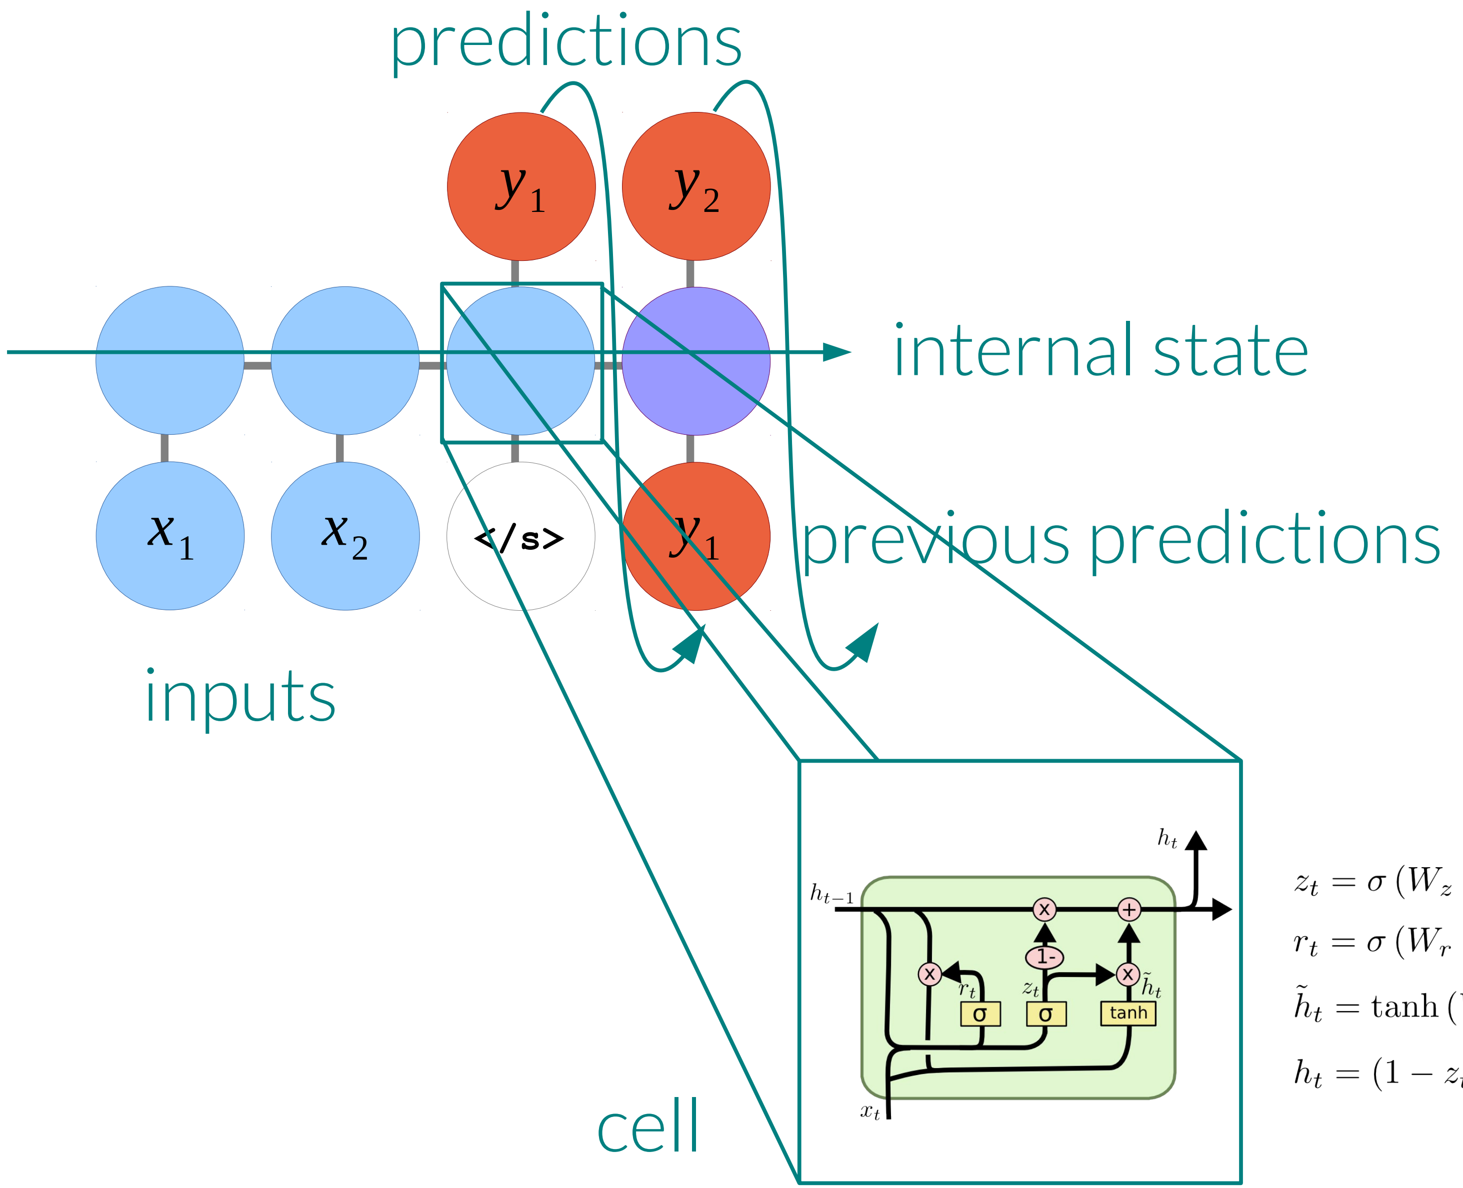

$$z_t = \sigma\left(W_z \cdot [h_{t-1}, x_t]\right)$$

$$r_t = \sigma\left(W_r \cdot [h_{t-1}, x_t]\right)$$

$$\tilde{h}_t = \tanh\left(W \cdot [r_t * h_{t-1}, x_t]\right)$$

$$h_t = (1 - z_t) * h_{t-1} + z_t * \tilde{h}_t$$

(GRU diagram by Christopher Olah)

Speaker and listener RNNs

Speaker

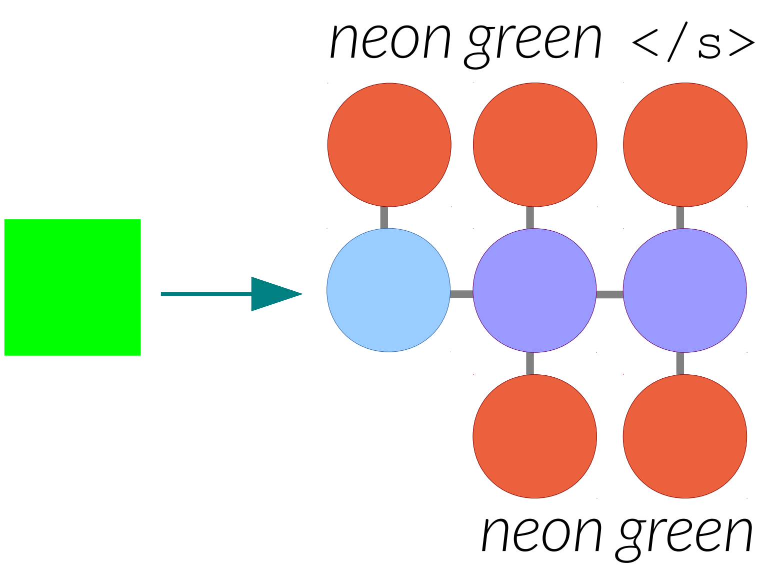

Listener

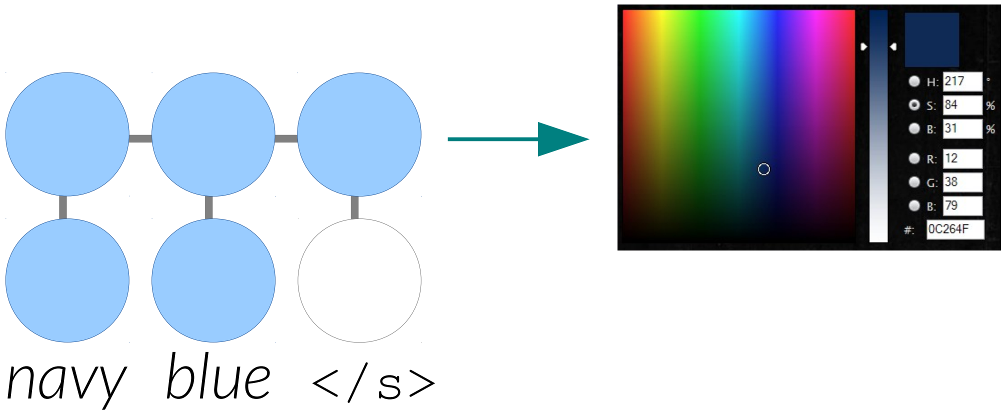

Results: modeling color descriptions

model	**features**	**perplexity**	**AIC** ($\times 10^6$)	**rec@1**
LSTM	Fourier	**12.86**	**4.07**	**39.76%**
M&S Lux	Gaussian	13.49	4.12	39.69
LSTM	buckets-4	17.83	4.58	34.96

Recap: learning through RSA

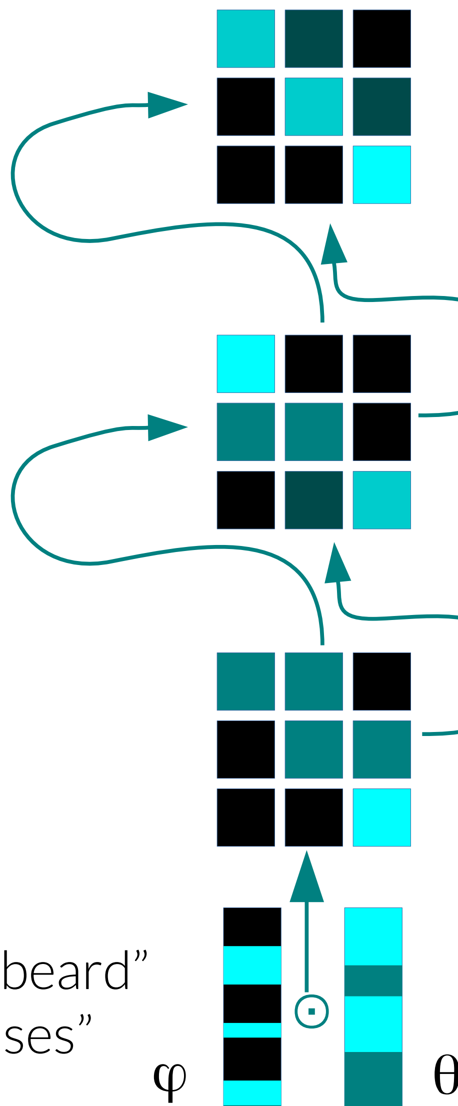

$$S_1(m|t,\theta) \propto L_1(t|m,\theta)$$

$$L_1(t|m,\theta) \propto S_0(m|t,\theta)$$

$$S_0(m|t,\theta) \propto \exp[\theta^T \varphi(t,m)]$$

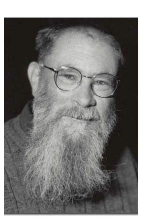

"beard"
"guy with the beard"
"guy with glasses"
...

Cooperative training with RNNs

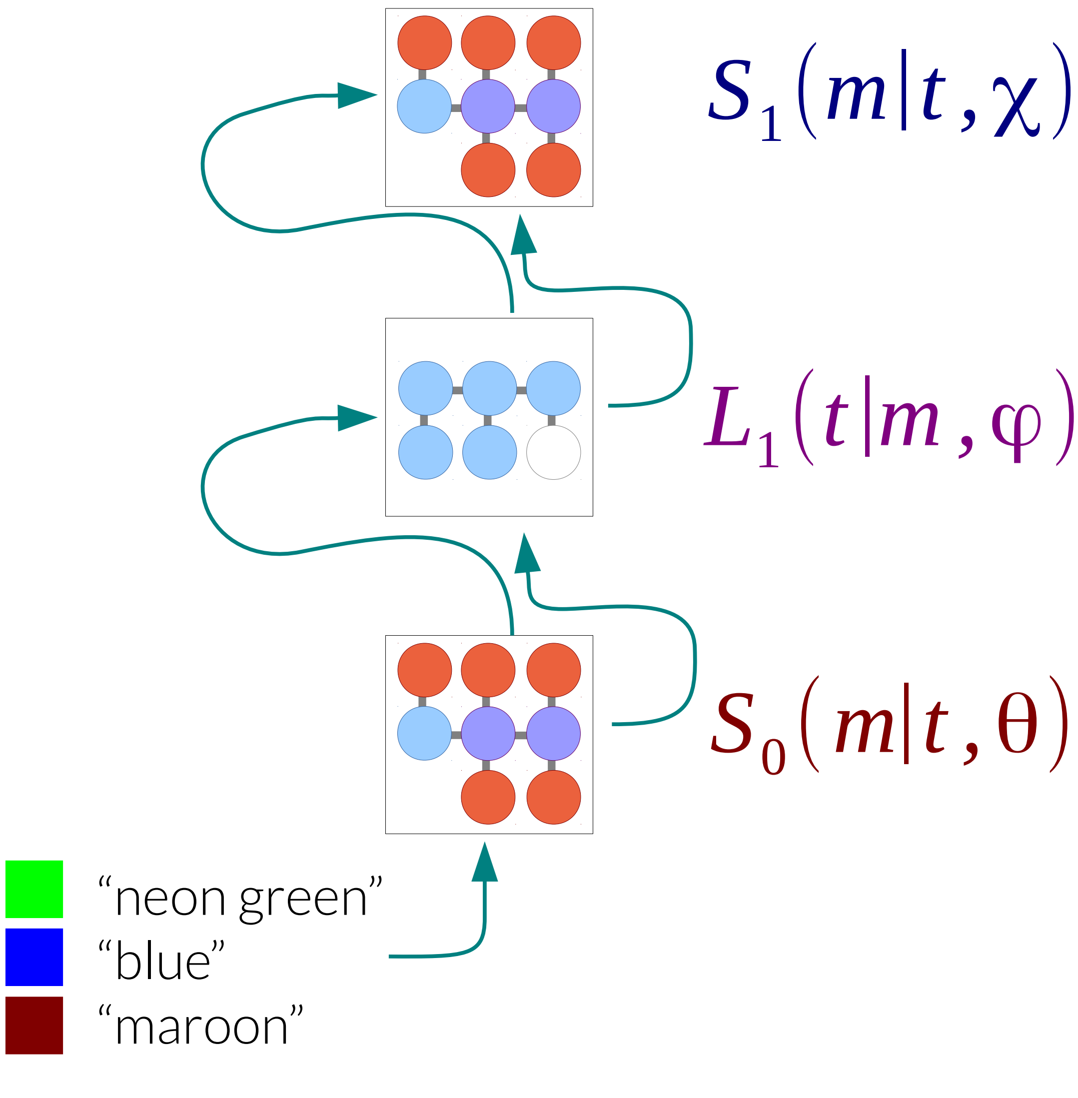

Cooperative training with RNNs

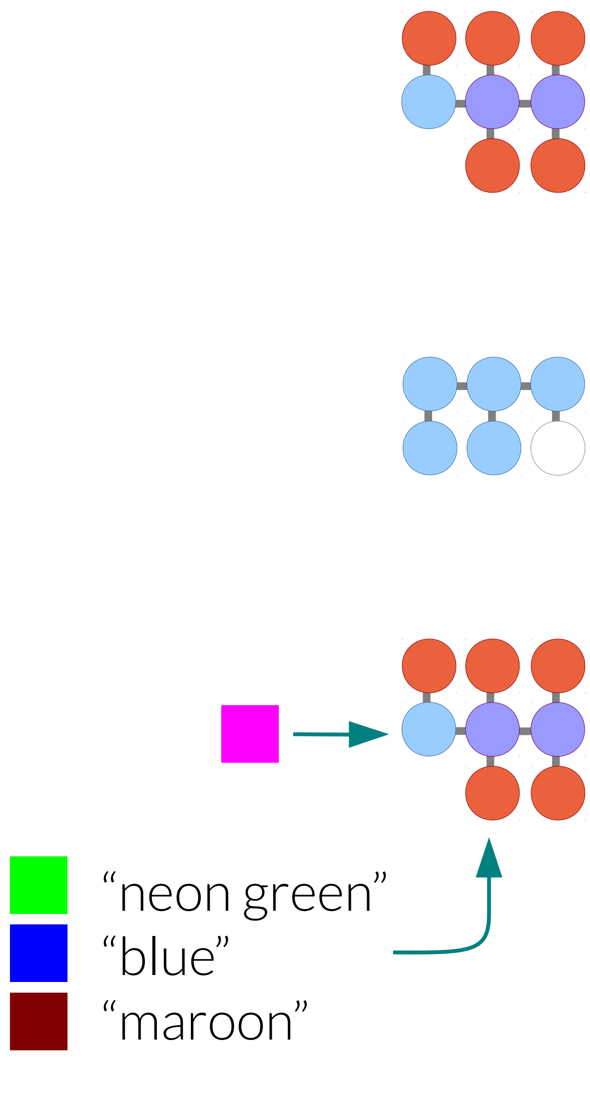

Cooperative training with RNNs

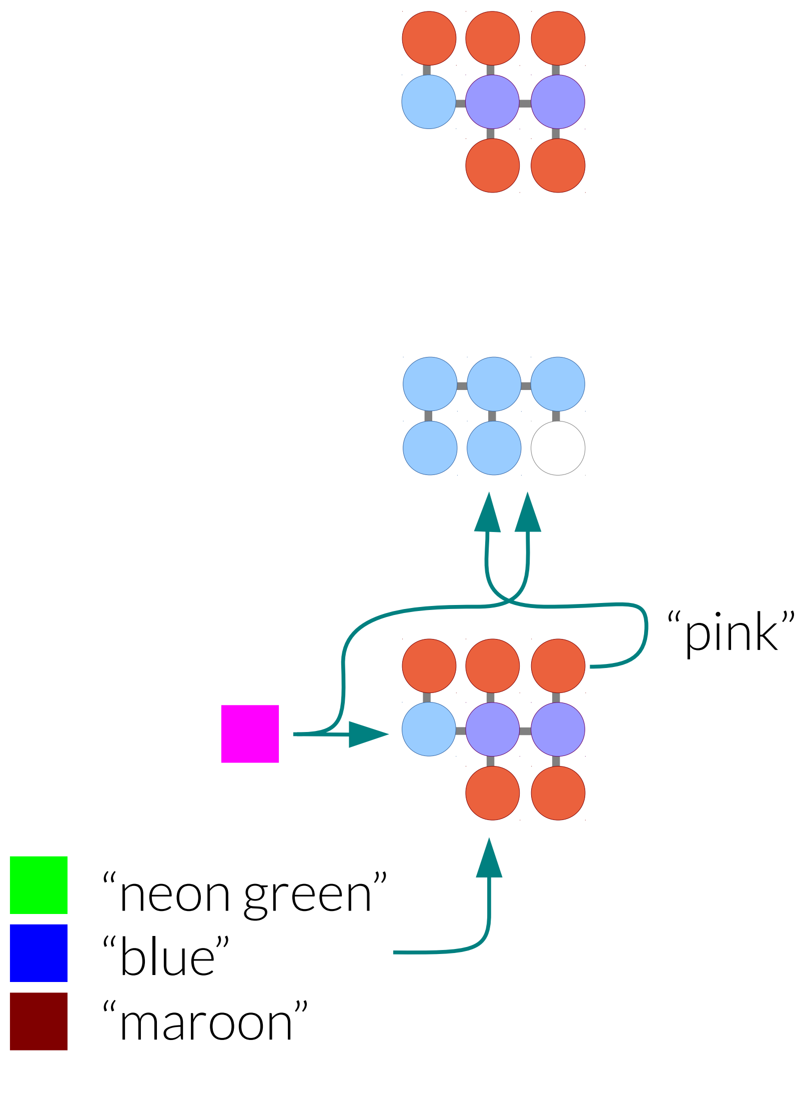

Cooperative training with RNNs

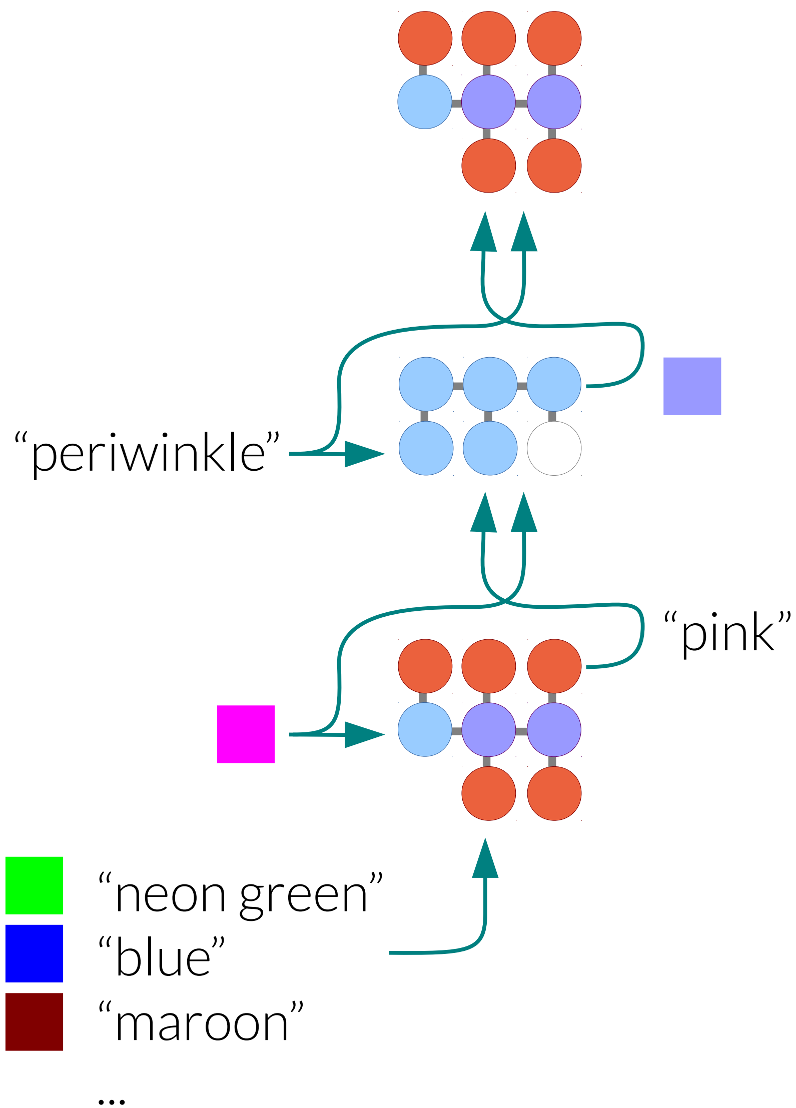

Cooperative training with RNNs

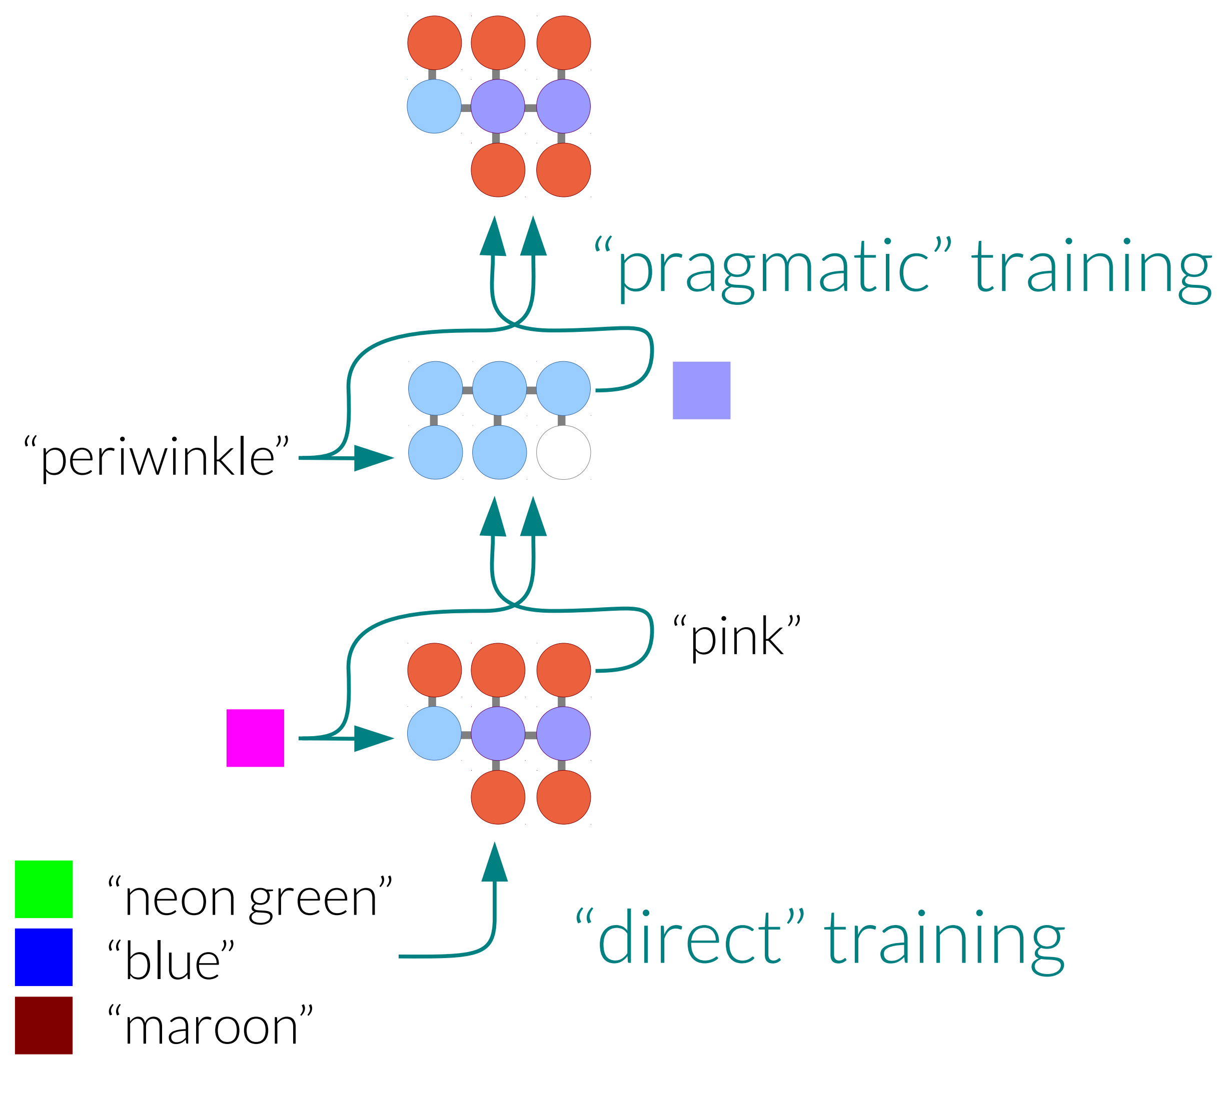

Comparing outputs

Human	Direct	Pragmatic
“marine blue”	“blue”	**“bright sky blue”**
“navy blue”	“navy blue”	**“almost black”**
“purple”	“purple”	**“faded purple”**
“purple”	“purple”	**“hot purple”**
“deep green”	“green”	**“true green”**
“green”	“green”	**“sap green”**
“olive”	“light green”	**“celery”**
“olive green”	“brown”	**“mustard brown”**
“mauve”	“peach”	**“peachy pink”**
“dark blue”	“blue”	**“marine blue”**

Summary

- Combining Bayesian pragmatics and learning:
  - context-dependent disambiguation
  - capturing oddities of human language use
  - avoiding the need for a hand-coded lexicon

- Making pragmatics scalable:
  - RNN-based sequence modeling
  - approximate optimization of RSA-based objective
  - bootstrapping a hyper-specific generation model

Thanks!