A Bayesian model of grounded color semantics†

Matthew Stone

Rutgers University

Oct 18, 2016

†Joint work with Brian McMahan and Timothy Meo

Grounding Semantics

Building AI systems that connect word meanings to experience

- Practical interest:
  Such systems can use language in powerful new ways
- Theoretical interest:
  Only such systems could have meanings intrinsically

See e.g., Harnad 1990; Roy & Pentland 2002; Cohen et al 2002;
Yu & Ballard 2004; Steels & Belpaeme 2005.

Recent Progress

Captioning images (e.g., Young et al 2014)

- Given a picture
- Give a natural language summary of what's happening

a brown dog and a tan one

Image from flickr user Joanne Bacon (jlbacon) via Young et al's data set.

How Grounding Works

By combining visual classifiers, compositional semantics, and models of natural communication

- For example, you need to associate color terms like *tan* with appropriate regions of color space

Key issue: How to you handle vagueness

Vagueness

Vague predicates have borderline cases:

definitely blue definitely not green	borderline blue borderline green	definitely not blue definitely green

Challenges of Borderline Cases

Within discourse, you need to use logic.

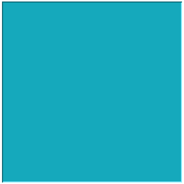

A:	*It's blue or it's green.*
B:	*It's not blue.*
A:	*It must be green.*

This is true even for borderline cases.

Challenges of Borderline Cases

Across discourses, you need to use common sense.

- People would be very likely to call this blue:

- People would be very reluctant to call this blue:

One Interpretation

Language users respect **logic**

- they reason consistently and coherently
- they don't say the same thing is blue at one moment, not blue the next

Language users respect **probability**, to prefer likely judgments

- they track others' likely meanings
- they describe things in ways others expect

Challenge of Vagueness

Not clear that logic and probability are really compatible.

- As illustrated by Sorites paradox

Sorites, I

This color is definitely blue:

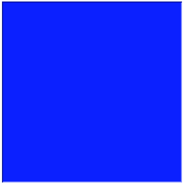

Sorites, II

If two colors are very similar, if one is blue then so is the other:

Sorites, II

If two colors are very similar, if one is blue then so is the other:

Sorites, II

If two colors are very similar, if one is blue then so is the other:

Sorites, II

If two colors are very similar, if one is blue then so is the other:

Sorites, II

If two colors are very similar, if one is blue then so is the other:

Sorites, III

Therefore, this color is definitely blue:

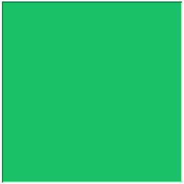

But this is absurd!

Outline of the Rest of the Talk

Thre main parts

- An emerging theory of vagueness
  Context is uncertain
  Meaning is fixed in context
- An implementation
  Describing colors in English with learned models
- Important open problems

But first, some background about color

Color is a sensation

A sophisticated abstraction of the world

- summarizes complexity and ambiguity of lighting
- product of the human visual system
- exploits human sensory apparatus and experience

See Fairchild 2013, Hughes et al 2013

Illumination

Everyday lighting involves varying distributions of energy

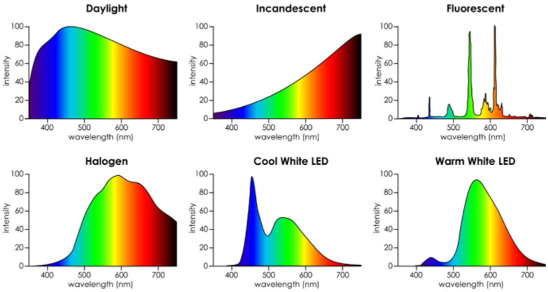

When we look at an object, we have no way to sense how it was illuminated

Reflectance

Light hits objects and scatters off

- The surface absorbs some wavelengths and emits others
- Depends on angle of incoming light and viewing angle
- Also depends on material: is it shiny, glossy, dull?

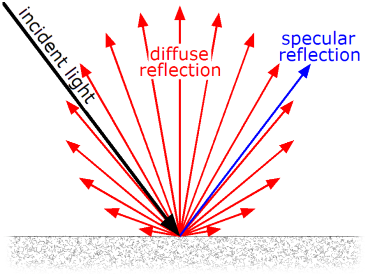

Image by GianniG46 via Wikimedia Commons

Eventually light hits your eyes

Three kinds of cones in the retina measure intensity

- Measure energy profiles across spectral intervals for red, green, and blue
- Values are products of illumination and reflectance

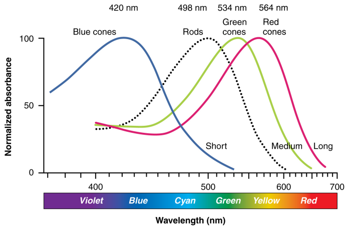

Image from OpenStax Human Anatomy and Physiology

From retina to color

Color of objects gives brain's estimate of diffuse reflectance

- Inferred by complex processing, not read off retina

Visual system can "unmultiply" numbers

Example

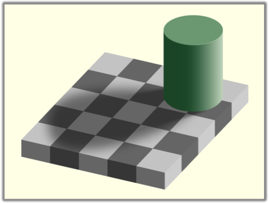

Image by Doug DeCarlo after Ted Adelson

Example

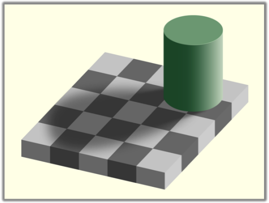

=

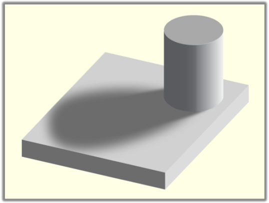

illumination

×

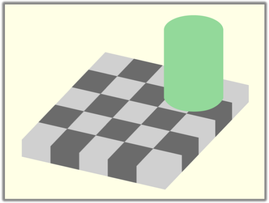

reflectance

Image by Doug DeCarlo after Ted Adelson

From retina to color

Color of objects gives brain's estimate of diffuse reflectance

- Approximately invariant to ambiguities from illumination, material, noise
- Draws on heuristics, experience (prior expectations), and statistical inference
- Describes how object looks across the visible spectrum

See e.g., Barrow and Tannenbaum 1978; Tappen et al 2005; Sharan et al 2008.

Color as spectrum

Colors summarize spectral profiles, **not** wavelengths of light

- We have pure primaries like red, green and blue

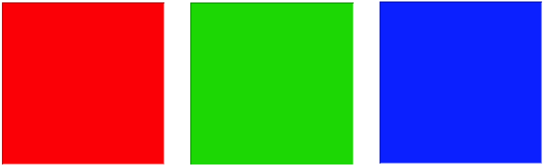

- But also saturated mixtures like yellow, cyan and magenta

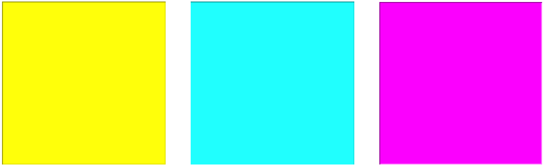

Physically, this is where the color wheel comes from

Color as 3D Space

Color reflects not only differences in hue

- Differences of saturation: muted to bright colors
- Differences of value: dark to light intensity

3D HSV space

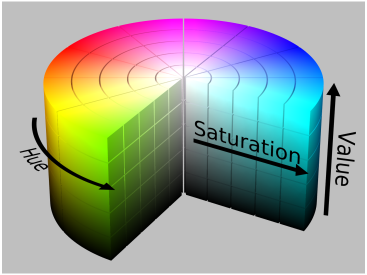

Image by Michael Horvath

Describing colors

We can distinguish millions of color values, but we typically group them into categories

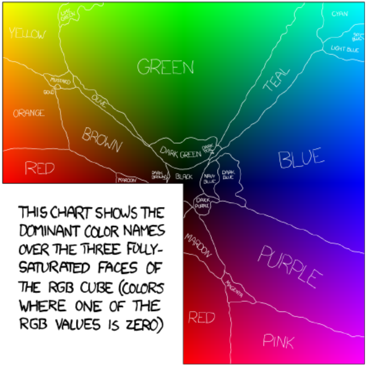

Image from Randall Munroe

Basic color terms

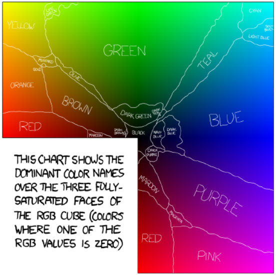

- Basic categories are convex regions
- They are subject to strong perceptual constraints
- There are strong crosslinguistic universals

See e.g., Berlin 1991; Regier et al 2005, 2007; Kay et al 2009

Describing colors

Not just **basic color terms**

- We also have **subordinate categories** (beige, lavender), **named subcategories** (royal blue, olive green), **modified ones** (dark blue, bright purple), and **blends** (reddish-orange).
- Some of these turn out to be very frequent.

Empirical Method

Crowdsourcing Color Descriptions

- In 2010, Randall Munroe (xkcd.com) conducted an online survey, gathering 3.4 million color judgments
- Our subset: 2.17 million data points across 829 color descriptions
- Pruned spam words, non-English speakers, and infrequent descriptions

We can use this data to explore color meanings at large scale

Vagueness

1. Context dependent meanings
2. Unceratin contexts
3. Continuous modeling

Context-sensitivity

The meaning of a vague word depends on context

- Words target distinctions that matter to us now
- The goal is efficient communication

Kamp 1975; Lewis 1979; Kyburg and Morreau 200; Graff Fara 2000

Context-sensitivity

The meaning of a vague word depends on context

- Words target distinctions that matter to us now
- The goal is efficient communication

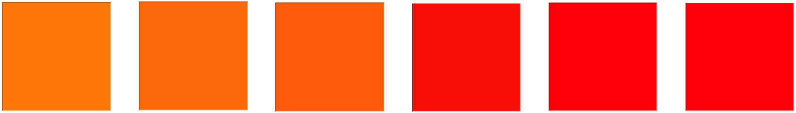

Kamp 1975; Lewis 1979; Kyburg and Morreau 200; Graff Fara 2000

Context-sensitivity for color

Context associates colors with a specific range of hues

blue range

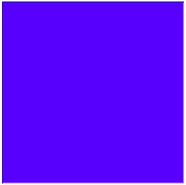

greenest blue

purplest blue

Meanings are defined by thresholds or boundaries in color space

Kennedy 2007

Probabilistic Context

We have only partial information about these boundaries

- A range of possible meanings may be in play
- But the way we use words commits us to certain distinctions
- As conversation unfolds, we accumulate constraints
- We eventually discover what distinctions to make

Modeled via probability distributions

Barker 2002; Lassiter 2009

Illustration for Color

Sample of possible boundaries for blue

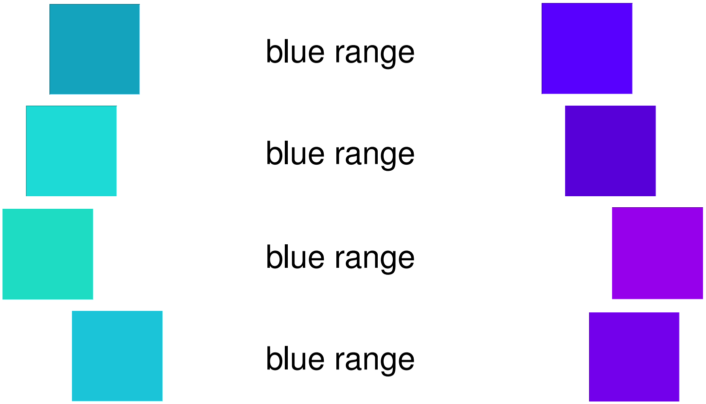

Illustration for Blue

Suppose we agree that this color counts as blue:

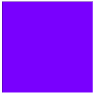

Then we rule out these meanings:

 blue range 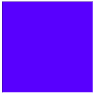

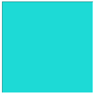 blue range 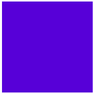

Illustration for Blue

Suppose we agree that this color counts as blue:

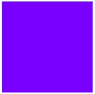

And narrow down to these ones:

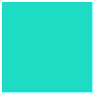 blue range 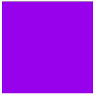

 blue range 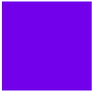

Formal Model

Use probability density functions to capture this uncertainty

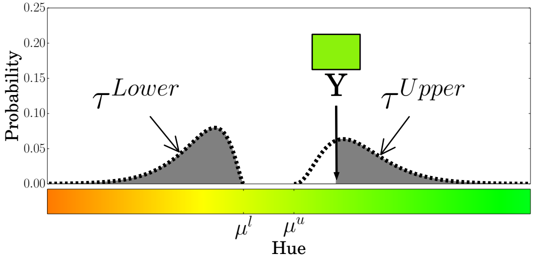

Formal Model

Use probability density functions to capture this uncertainty

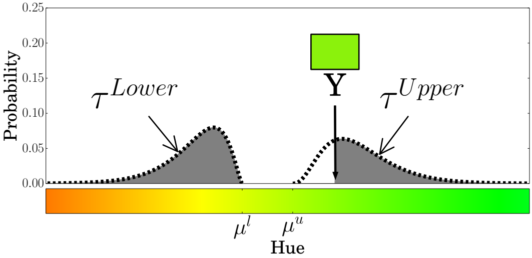

- Random variables τ^{Lower} and τ^{Upper}
- In any context, category fits value y if $\tau^{Lower} \leq y \leq \tau^{Upper}$

Formal Model

Use probability density functions to capture this uncertainty

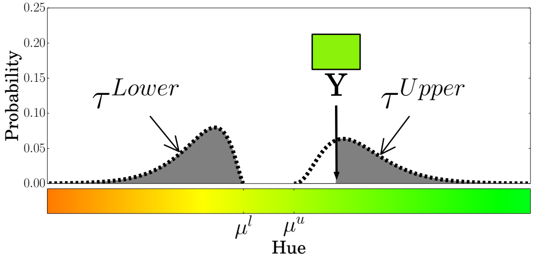

- Track a posterior over τ^{Lower} and τ^{Upper} in context
- Using prior and observed categorization examples

Formal Model

Use probability density functions to capture this uncertainty

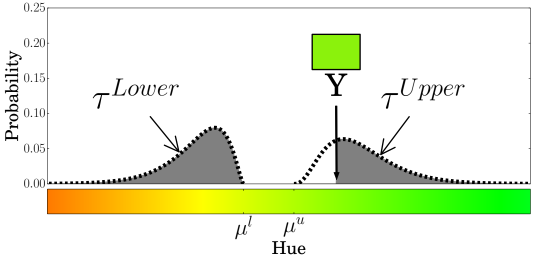

- Prior over τ^{Lower} and τ^{Upper} is estimated from data
- in terms of parameters μ^{Lower} and μ^{Upper} that give limits for τ^{Lower} and τ^{Upper}

Formal Model

Use probability density functions to capture this uncertainty

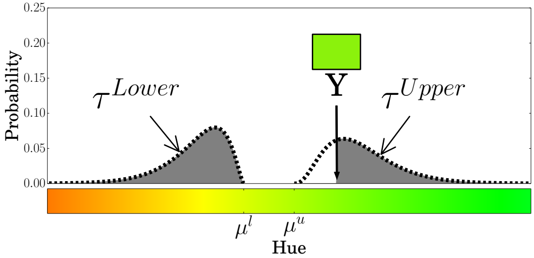

- Prior over τ^{Lower} and τ^{Upper} is estimated from data
- in terms of a Γ distribution (standard statistical tool) that models how flexible thresholds are

Formal Model

Mathematical summary—Applicability of a category *k* to a color *x* in context, written

$$P(k^{true}|\mathbf{x})$$

takes hue, saturation and value into account simultaneously:

$$= \prod_{d \in (H,S,V)} P(\tau_k^{L,d} < x_i^d < \tau_k^{U,d}) = \prod_{d \in (H,S,V)} \phi_k^d(x^d)$$

$$\phi_k^d(x^d) = \begin{cases} P(x^d > \tau_k^{L,d}), & x^d \leq \mu_k^{L,d} \\ P(x^d < \tau_k^{U,d}), & x^d \geq \mu_k^{U,d} \\ 1, & otherwise \end{cases}$$

Computational Experiments

1. Learning a color lexicon
2. Using a color lexicon

Model Challenges

Munroe's data associates color values with descriptions

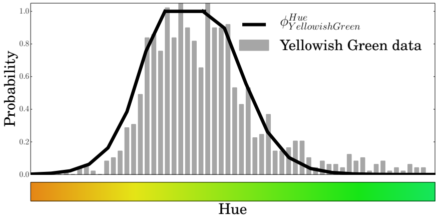

But we need a linking assumption to infer meaning

- Potential problem: pragmatics intervenes
- Choices could involve implicature and "theory of mind"

Anderson 1991; Smith et al 2013; Goodman and Lassiter 2013

Model Challenges

Munroe's data associates color values with descriptions

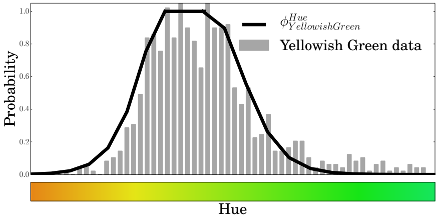

But we need a linking assumption to infer meaning

- Key question: what are speaker's communicative goals and assumptions about the listener

Anderson 1991; Smith et al 2013; Goodman and Lassiter 2013

Model Challenges

Munroe's data associates color values with descriptions

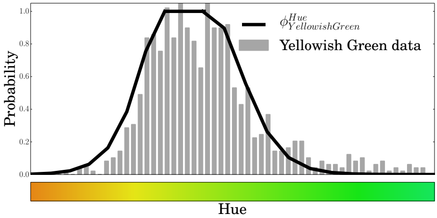

But we need a linking assumption to infer meaning

- First stab: no obvious signature of implicature in the histograms

Anderson 1991; Smith et al 2013; Goodman and Lassiter 2013

Our Model: Rational Observer

Our assumptions about crowd contributors

- They want to say things they expect to be true
- They want match baseline rates of word use

They **don't** have a communicative goal of informativeness

- So listeners can't second-guess
  why speakers said as much (or as little) as they did

Formal Details

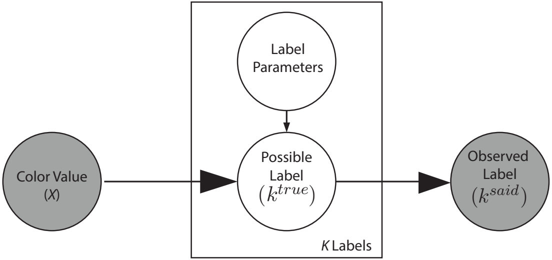

$$P(k^{said}, k^{true}, \mathbf{x}) = P(k^{said}|k^{true})P(k^{true}|\mathbf{x})P(\mathbf{x})$$

Formal Details

$$P(k^{said}, k^{true}, \mathbf{x}) = P(k^{said}|k^{true})P(k^{true}|\mathbf{x})P(\mathbf{x})$$

- **Availability**: $P(k^{said}|k^{true})$
  - How often is it used when it is true
  - Contrast *yellow-green* (high availability) and *chartreuse* (low availability)
- **Subjective Likelihood**: $P(k^{true}|\mathbf{x})$
  - Fit of the description given vague semantics
- **Prior**: $P(\mathbf{x})$
  - Munroe's sampling rate, unfortunately not uniform

Deriving the Lexicon

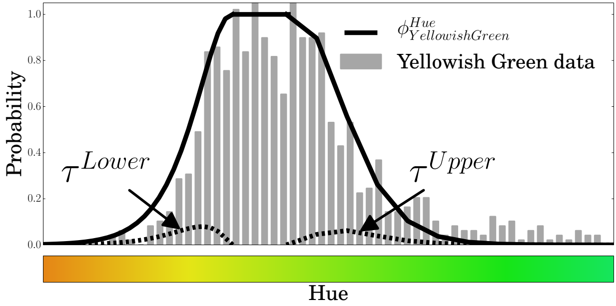

- Assume data is derived from such a process
- Fit to corpus using Markov Chain Monte Carlo

Result

The Lexicon of Uncertain Color Standards

- 829 color descriptions and their parameters
- Flexible semantic representations divorced from the methods to use them
- Happy to share this!

Demo

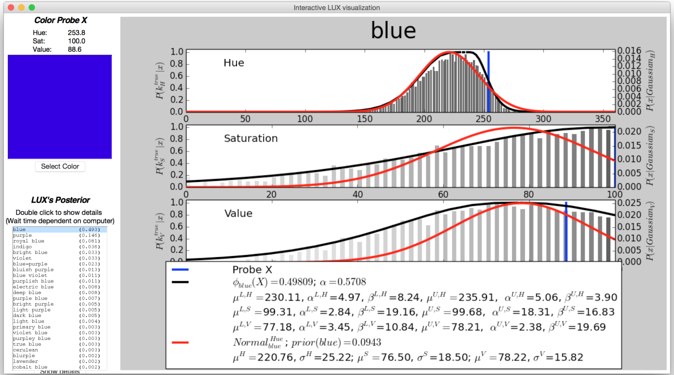
Interactive LUX visualization
Color Probe X
Hue: 253.8
Sat: 100.0
Value: 88.6
Select Color
LUX's Posterior
Double click to show details
(Wait time dependent on computer)
blue (0.493)
purple (0.146)
royal blue (0.081)
indigo (0.036)
bright blue (0.033)
violet (0.033)
blue-purple (0.023)
bluish purple (0.013)
blue violet (0.011)
purplish blue (0.011)
electric blue (0.008)
deep blue (0.008)
purple blue (0.007)
bright purple (0.005)
light purple (0.005)
dark blue (0.005)
light blue (0.004)
primary blue (0.003)
violet blue (0.003)
purpley blue (0.003)
true blue (0.003)
cerulean (0.003)
blurple (0.002)
lavender (0.002)
cobalt blue (0.002)
blue
Hue
Saturation
Value
$P(k_H^{true}|x)$
$P(k_S^{true}|x)$
$P(k_V^{true}|x)$
$P(x|Gaussian_H)$
$P(x|Gaussian_S)$
$P(x|Gaussian_V)$
Probe X
$\phi_{blue}(X) = 0.49809$; $\alpha = 0.5708$
$\mu^{L,H} = 230.11$, $\alpha^{L,H} = 4.97$, $\beta^{L,H} = 8.24$, $\mu^{U,H} = 235.91$, $\alpha^{U,H} = 5.06$, $\beta^{U,H} = 3.90$
$\mu^{L,S} = 99.31$, $\alpha^{L,S} = 2.84$, $\beta^{L,S} = 19.16$, $\mu^{U,S} = 99.68$, $\alpha^{U,S} = 18.31$, $\beta^{U,S} = 16.83$
$\mu^{L,V} = 77.18$, $\alpha^{L,V} = 3.45$, $\beta^{L,V} = 10.84$, $\mu^{U,V} = 78.21$, $\alpha^{U,V} = 2.38$, $\beta^{U,V} = 19.69$
$Normal_{blue}^{Hue}$; $prior(blue) = 0.0943$
$\mu^{H} = 220.76$, $\sigma^{H} = 25.22$; $\mu^{S} = 76.50$, $\sigma^{S} = 18.50$; $\mu^{V} = 78.22$, $\sigma^{V} = 15.82$

Evaluation

Train on 70% of target Munroe data; develop on 5%; test on 25%

Compare to two alternative models

- GM: Gaussian Model
  Color prototypes with category-specific distance measures
- HM: Histogram Model
  Divides HSV space into nested cells and predicts from smoothed counts

Measure

- Model complexity
- Precision of model fit
- Performance of model labeling

Performance

Model predictions

	TOP^1	TOP^5	TOP^{10}
LUX	39.55%	69.80%	80.46%
HM	39.40%	71.79%	82.53%
GM	39.05%	69.25%	79.99%

Model complexity and fit

	-LL	-LLV	AIC	Perplexity
LUX	$1.13*10^7$	$2.05*10^6$	$4.13*10^6$	13.61
HM	$1.13*10^7$	$2.05*10^6$	$4.82*10^6$	14.41
GM	$1.34*10^7$	$2.08\hat{}10^6$	$4.17*10^6$	14.14

Limitations

Doesn't handle ambiguity—as in case of *melon*

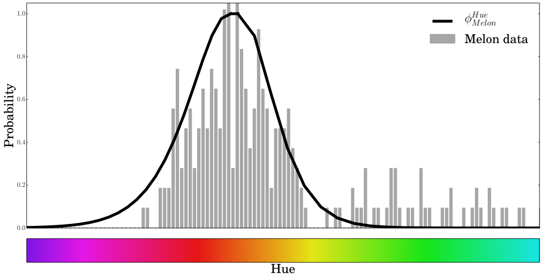

Limitations

Doesn't describe shape of all categories—as in case of *greenish*

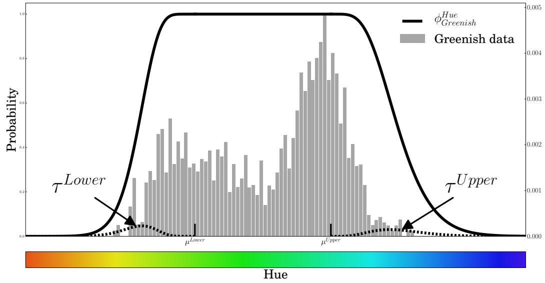

Generalizing Communicative Goals

New task (cf. van Deemter 2006):

- Distinguish one color Y from another color Z
- Two directions: Generation and understanding

Probabilistic and logical approach

- What's the probability in context
- that the meaning of a word applies to Y
- but does not apply to Z?

Distinguishing Colors in Context

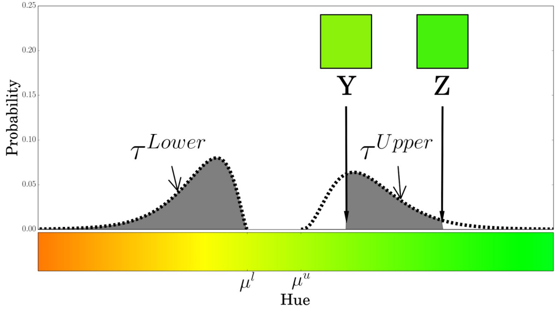

$$\phi(\mathbf{Y}) - \phi(\mathbf{Z})$$

Distinguishing Colors in Context

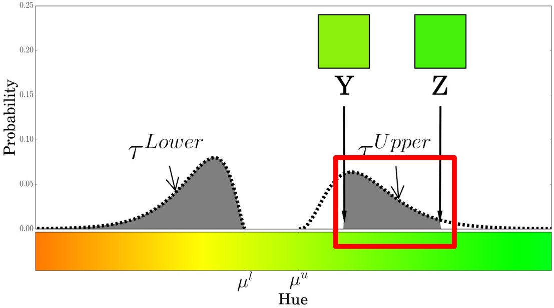

$$\phi(\mathbf{Y}) - \phi(\mathbf{Z})$$

Distinguishing Colors in Context

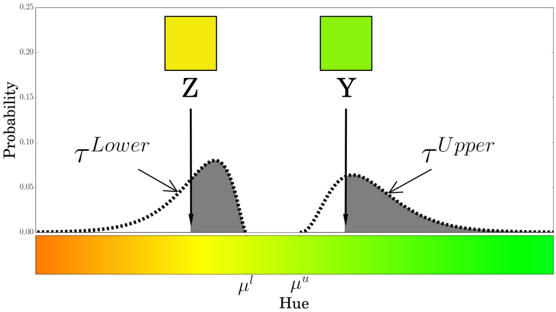

$$\phi(\mathbf{Y}) * (1 - \phi(\mathbf{Z}))$$

Distinguishing Colors in Context

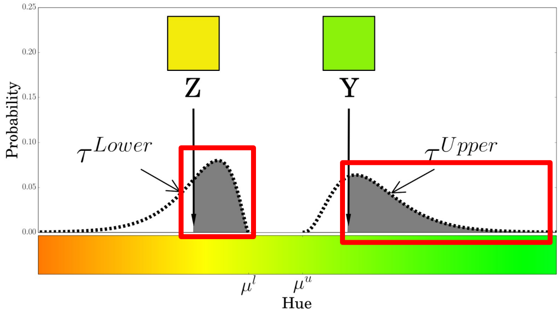

$$\phi(\mathbf{Y}) * (1 - \phi(\mathbf{Z}))$$

Three dimensional case

$$\phi_1(\mathbf{Y}) * (\phi_2(\mathbf{Y}) - \phi_3(\mathbf{Z}))$$

- $\phi_1(\mathbf{Y})$ is the product of $\phi(\mathbf{Y}^d)$ over dimensions d where Y and Z are contrasting borderline cases
- $\phi_2(\mathbf{Y})$ is the product of $\phi(\mathbf{Y}^d)$ over all other dimensions d
- $\phi_3(\mathbf{Z})$ is the product of $\phi(\mathbf{Z}^d)$ over all dimensions d

Demo

Interactive LUX visualization

Color 0.71 1.00 1.00 (hsv)

Associated Color Terms		Differentiating Color Terms	
blue	0.432	blue	0.669
purple	0.167	royal blue	0.090
royal blue	0.058	bright blue	0.071
violet	0.044	blue-purple	0.023
bright blue	0.041	electric blue	0.022
indigo	0.036	purplish blue	0.012
blue-purple	0.028	bluish purple	0.010
bluish purple	0.017	indigo	0.009
electric blue	0.017	blue violet	0.008
purplish blue	0.015	deep blue	0.007

Color 0.75 1.00 1.00 (hsv)

Associated Color Terms		Differentiating Color Terms	
purple	0.469	purple	0.710
blue	0.122	bright purple	0.103
violet	0.077	violet	0.057
bright purple	0.049	neon purple	0.020
indigo	0.042	fuschia	0.019
blue-purple	0.021	light purple	0.015
bluish purple	0.016	pink	0.012
royal blue	0.016	magenta	0.012
light purple	0.015	electric purple	0.006
blue violet	0.014	royal purple	0.006

Test color term to interpret

Evaluation: Label Resolution

- Evaluated using human reference data

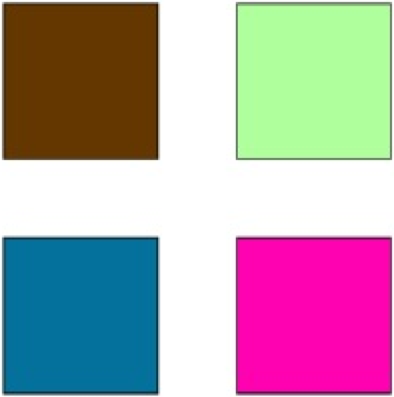

Baumgaertner et al 2012

Evaluation: Label Resolution

- Evaluated using human reference data

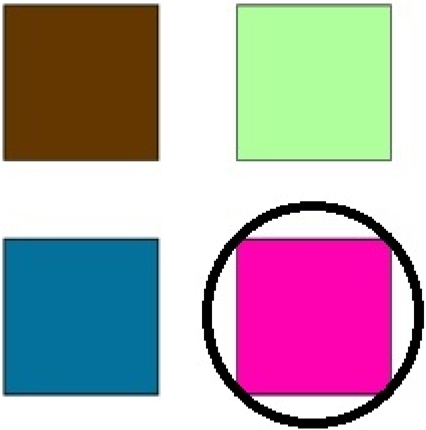

Baumgaertner et al 2012

Evaluation: Label Resolution

- Evaluated using human reference data

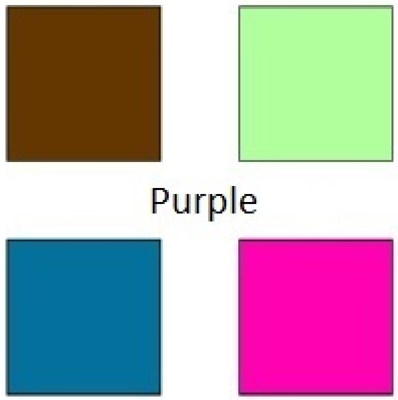

Baumgaertner et al 2012

Evaluation: Label Resolution

- Evaluated using human reference data

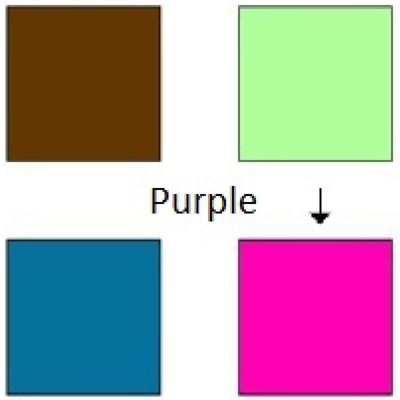

Baumgaertner et al 2012

Results: Label Resolution

- 196 tokens, varying difficulty levels
- Our system's performance:
  - 152 correct
  - 8 incorrect, several of which seem to be cases of human error, e.g. directors produced color labels for the wrong swatch
  - 28 out-of-vocabulary terms that are closely related to vocabulary in our system that would have been correct answers
  - 6 out-of-vocabulary terms that are closely related to vocabulary in our system that would have been incorrect answers
  - 2 which are completely different from any vocabulary in our system

Limitations

More logic is still needed to capture relationships among descriptions

An attested case

a tan dog and a white one, or
a brown dog and a tan one

Conclusions

Vagueness is dependence on an uncertain context

- A way to explain linguistic and philosphical intuitions
- A basis for learning and using grounded meanings

Focus on the case of color

- Learned models of vague meaning from multimodal corpus
- Applied the models for new tasks

Looking forward

Observed several limits on representations

- Handling ambiguity
- Modeling complex categories
- Interpreting extended descriptions

Important to address these limitations

Looking forward

More expressive language

- Visually-grounded categories
- Spatial language
- Verbs—change over time

Looking forward

Studying the semantics–pragmatics interface

- Our model gives a Nash equilibrium for semantics

  $$P(\mathbf{x}|k^{said}) = P(\mathbf{x}|k^{true})$$

  in keeping with Lewis 1969
- But different assumptions about goals lead to different strategic situations
- So what kind of goals and reasoning really fit conversation?

Looking forward

Richer interactions

- Empirically testing dynamics of vagueness in dialogue
- Using dynamic models to plan and track conversations
- Achieving grounding
  for example with situated interaction with robots