

Reinforcement Learning: Principles and Applications






Credit: Aldebaran

Devdatt Dubhashi
AI and Data Science Division
Chalmers University

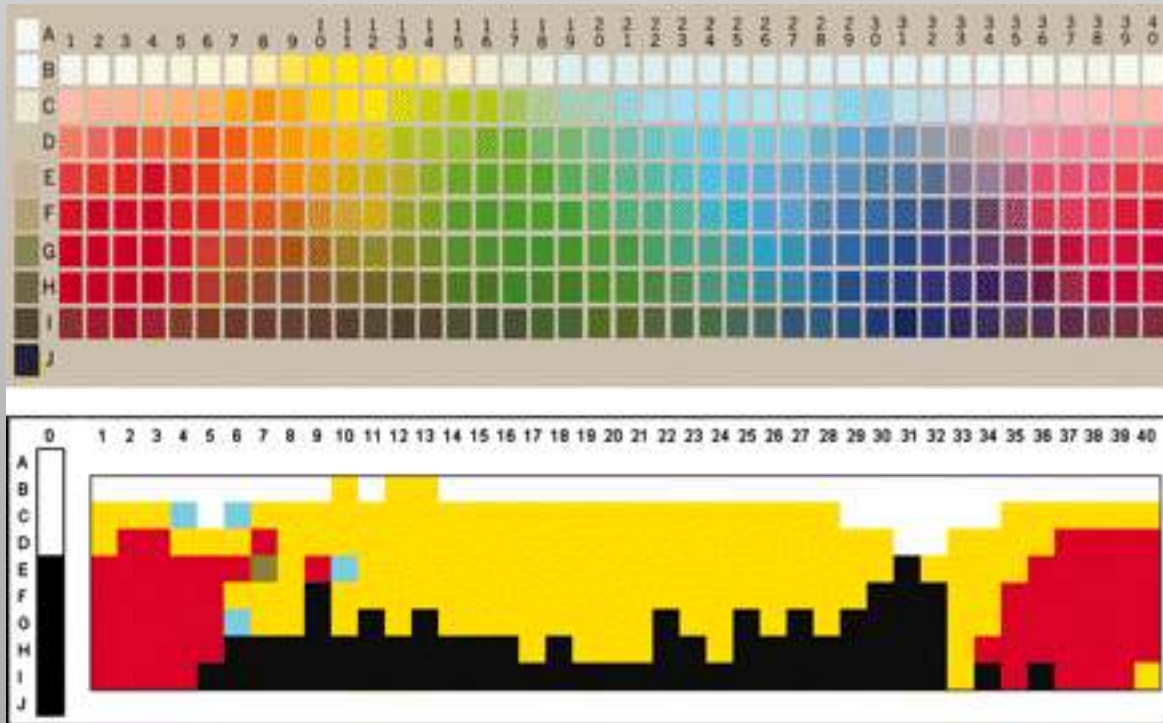
Are Colours Universal?



WCS Stimulus Palette



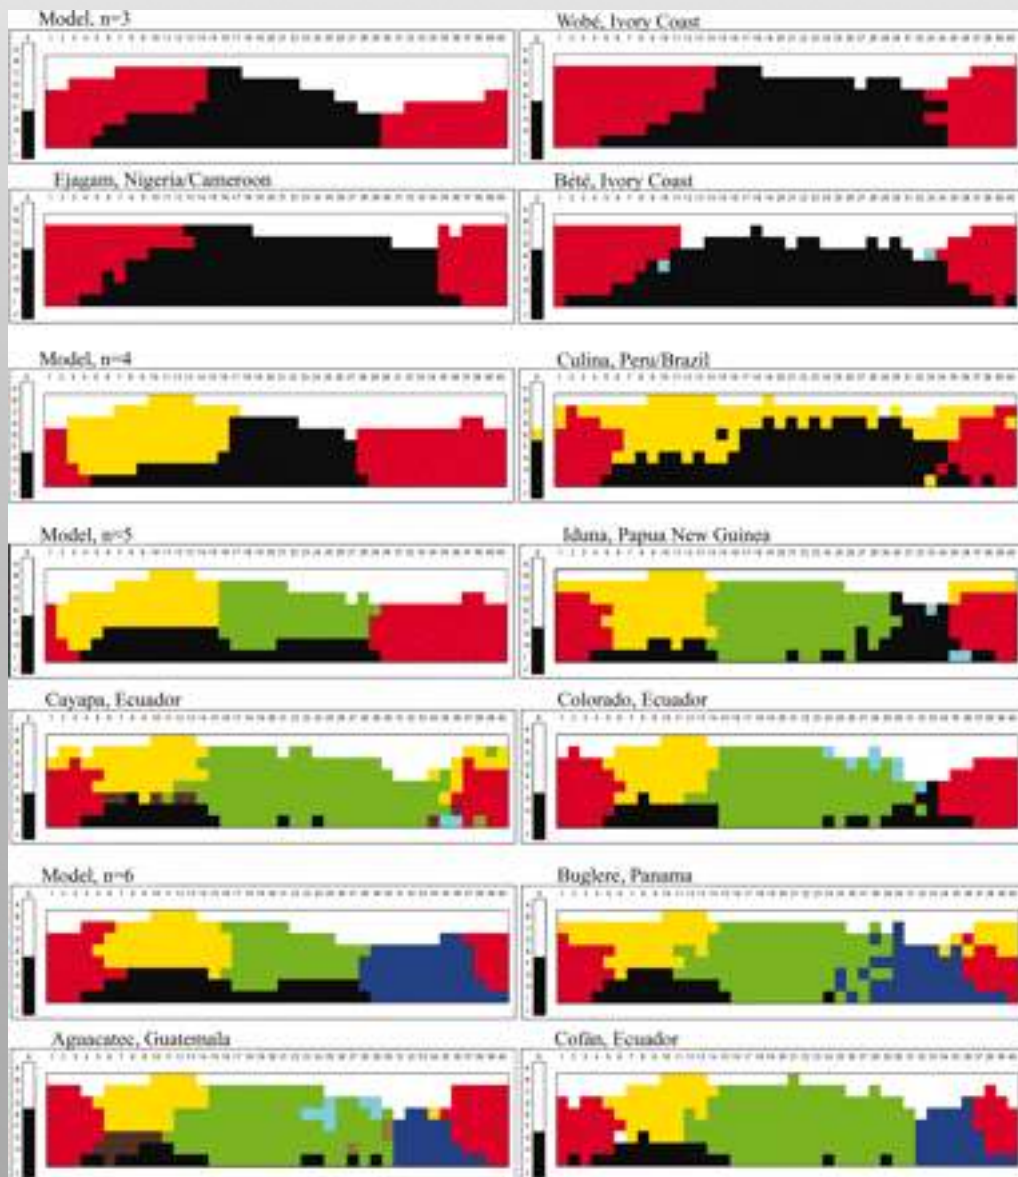
World Colour Survey



Language as Efficient Communication

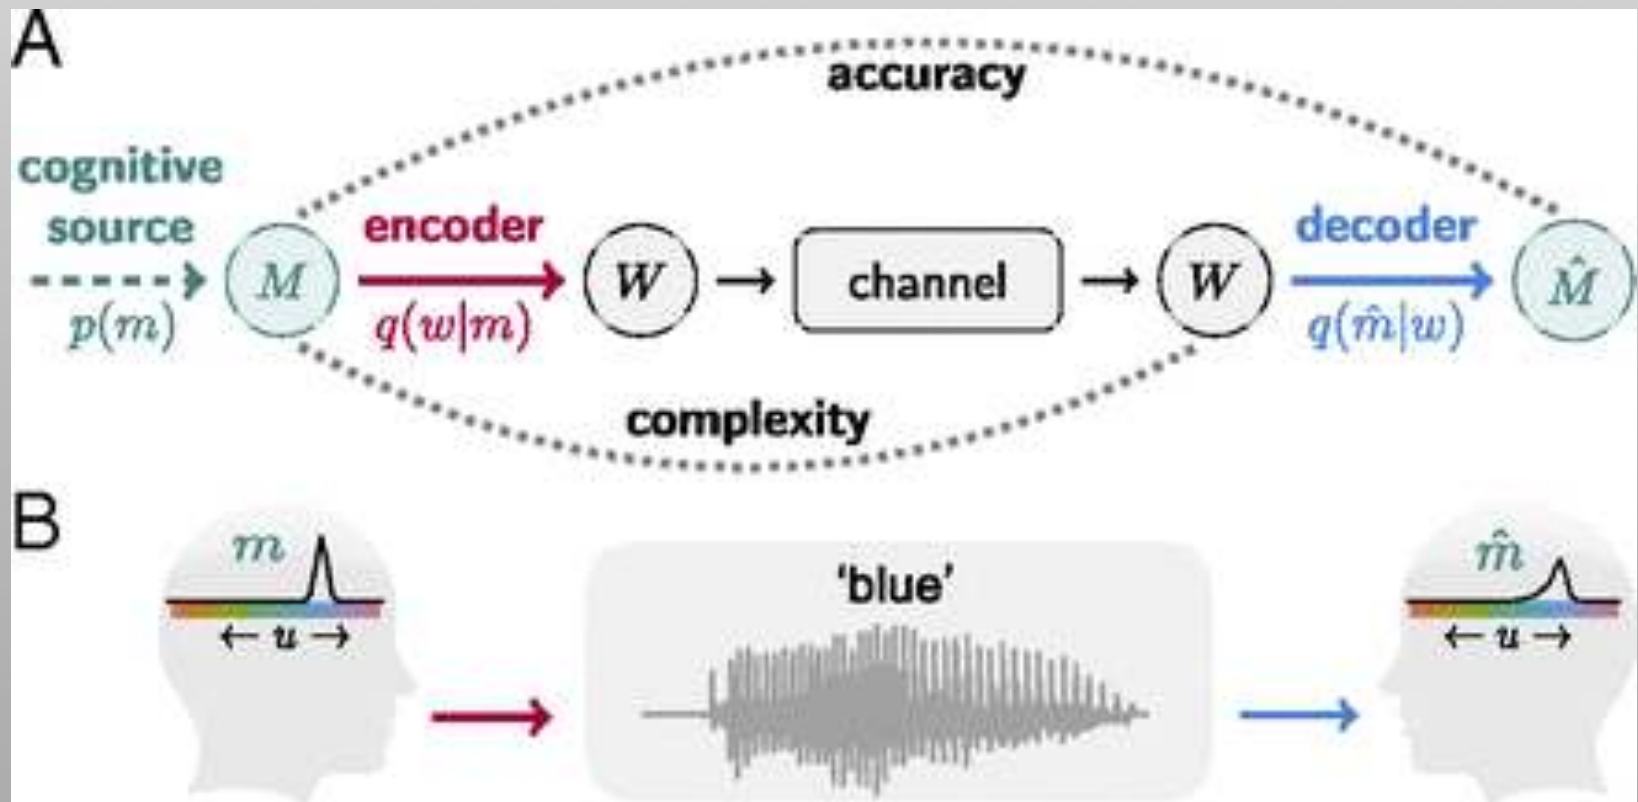
- Communication over a noisy channel (Shannon)
- Informativeness
- Communication Efficiency versus
- Cognitive Load
- Rosch (1999, p. 190): semantic categories tend to “provide maximum information with the least cognitive effort.”



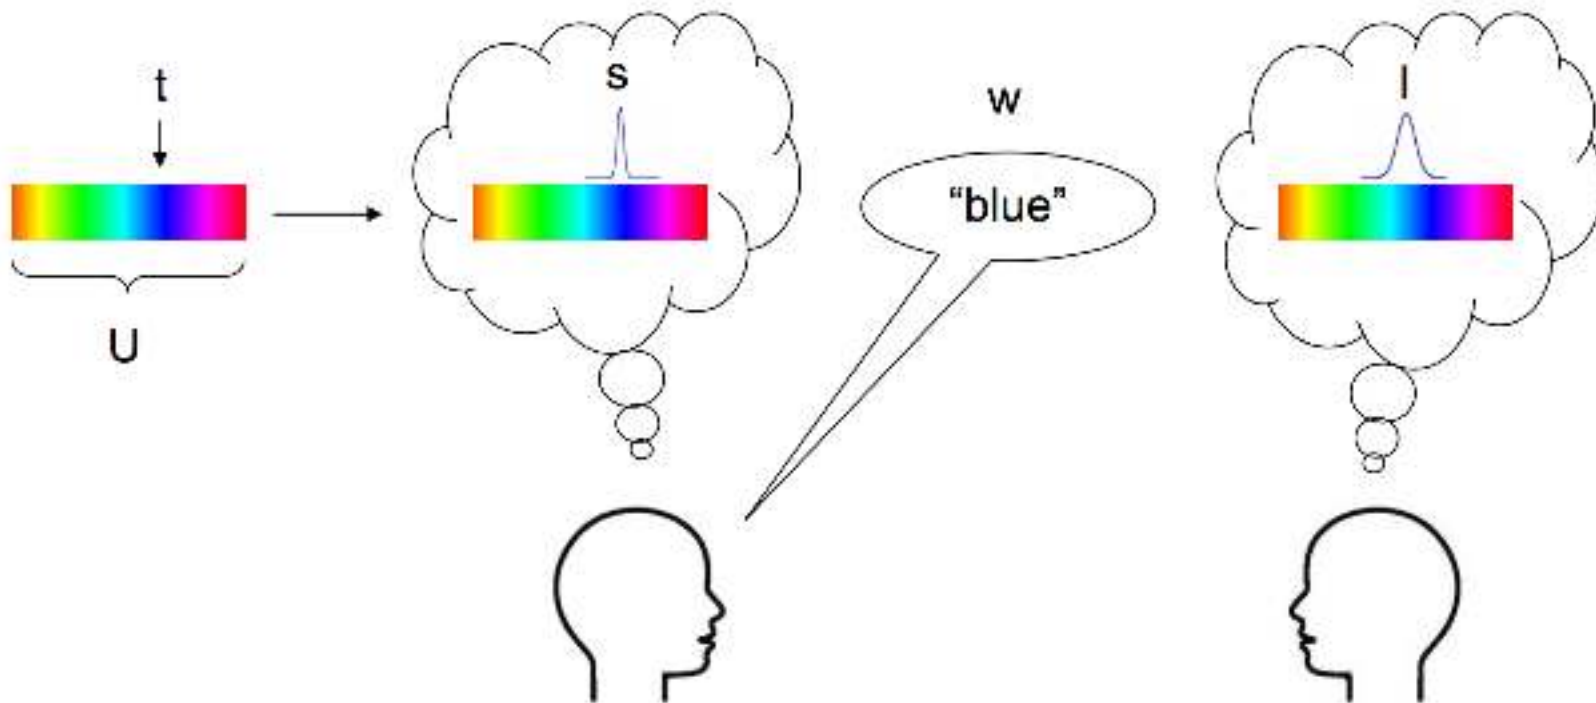


Regier T, Kemp C, Kay P. 2015.
 Word meanings across languages
 support efficient communication.
 In
*The Handbook of Language
 Emergence*, ed. B MacWhinney, W
 O'Grady


Shannon Information Theory



Information Theoretic View



$$D(s||l) = \sum_{i \in U} s(i) \log_2 \frac{s(i)}{l(i)} = \log_2 \frac{1}{l(t)} = -\log_2(l(t))$$



Tsimane' people, an indigenous nonindustrialized Amazonian group consisting of about 6,000 people from lowland Bolivia who live by farming, hunting, and foraging for subsistence English speakers in the United States; and Bolivian-Spanish speakers in Bolivia, neighboring the Tsimane. (Gibson et al, PNAS 2017.)

Unsupervised Learning

No
feedback



Images sampled
from learnt distribution

Feedback on
outcome


Reinforcement Learning



Sometimes
delayed

Supervised Learning

Classification



Correct
answer given

CAT

Reinforcement Learning

internal state



environment



action

observation


No Labels
needed



Flexible

Biologically
plausible

Weak signal

Can be
unstable





w

The Color Game

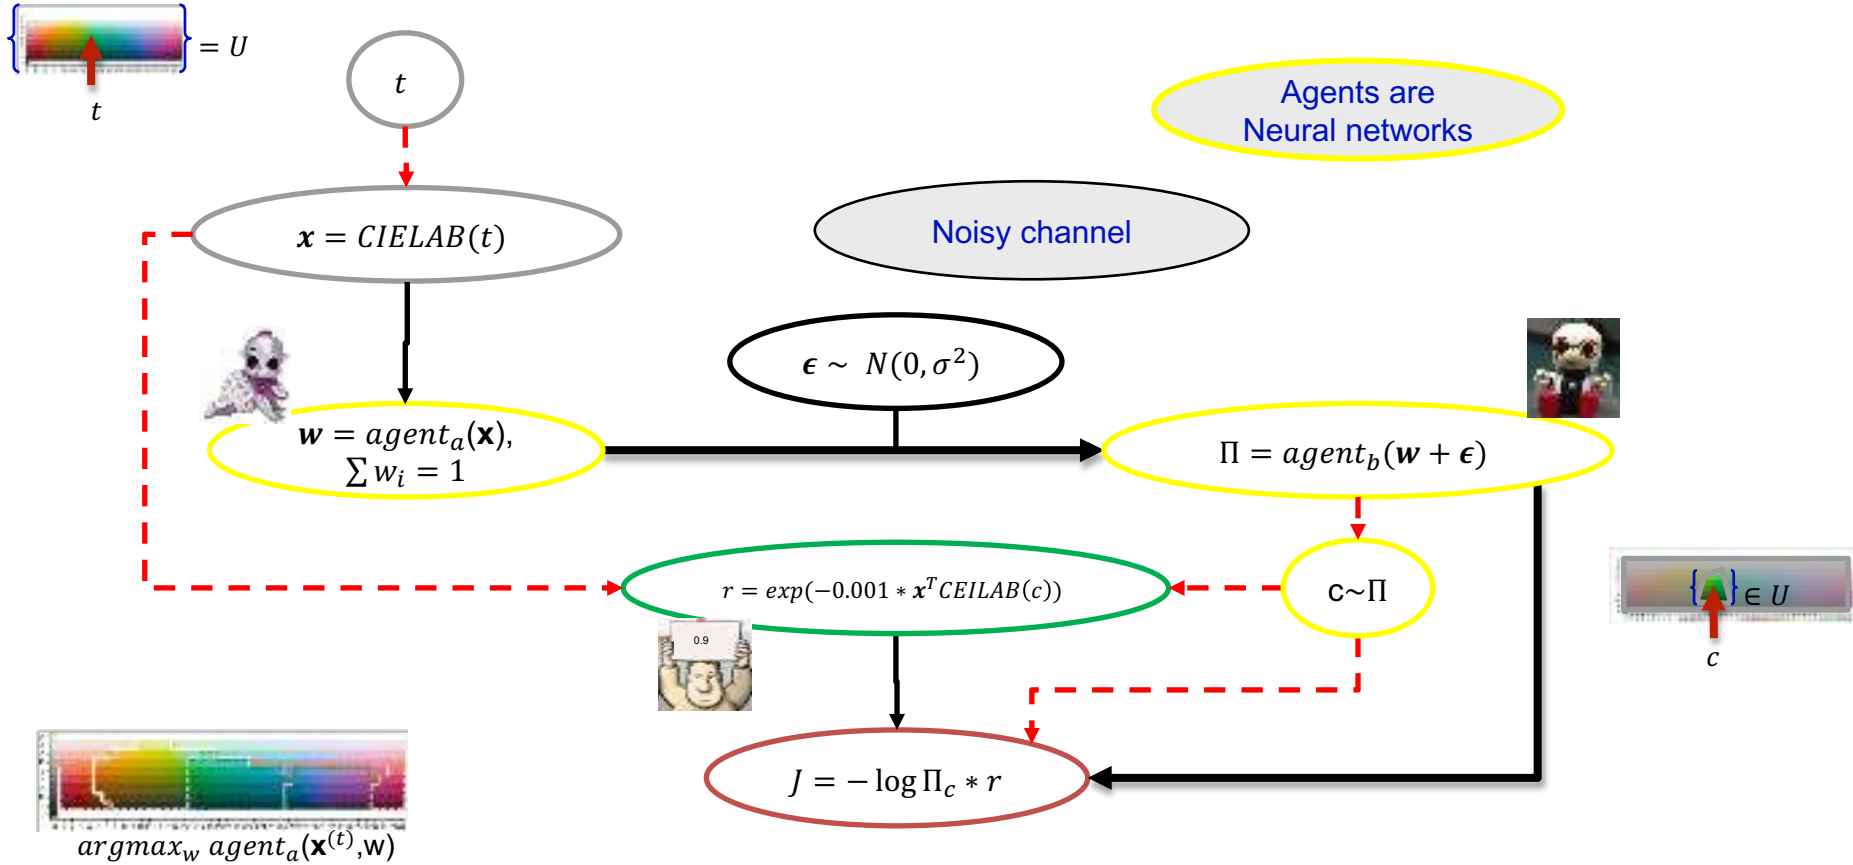
Speaking agent



Judge

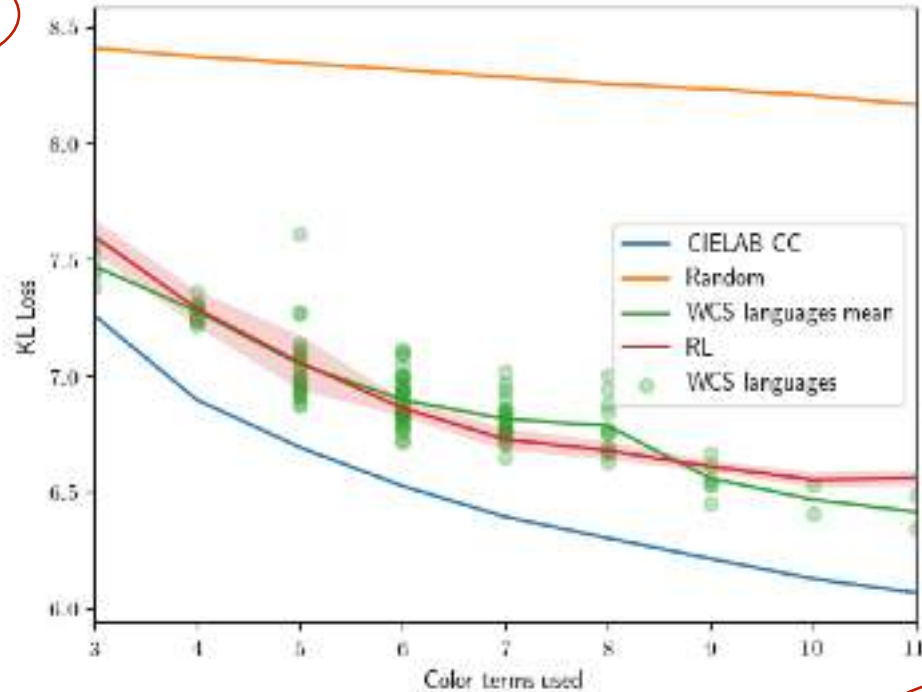


Listening agent



Information Loss

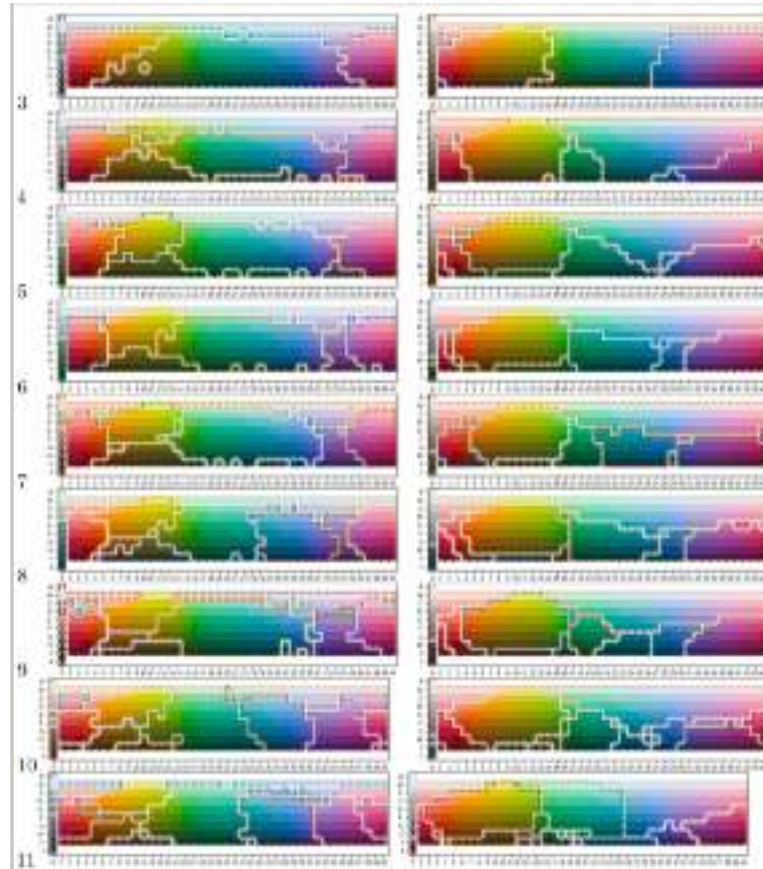
Efficient Communication




Complexity of language

Table 3.1.1: Comparison of the human languages in WCS to generated languages using Rand index. Abbreviations used in table column headers: H=human, RL=reinforcement learning, and R=random

Terms	H-H	RL-RL	H-RL	H-R
3	0.701	0.273	0.173	0.000
4	0.452	0.337	0.167	0.000
5	0.476	0.373	0.223	0.000
6	0.528	0.537	0.277	0.000
7	0.472	0.593	0.292	0.000
8	0.471	0.518	0.281	0.000
9	0.584	0.510	0.321	0.000
10	0.718	0.549	0.316	0.000
11	0.472	0.543	0.309	0.000

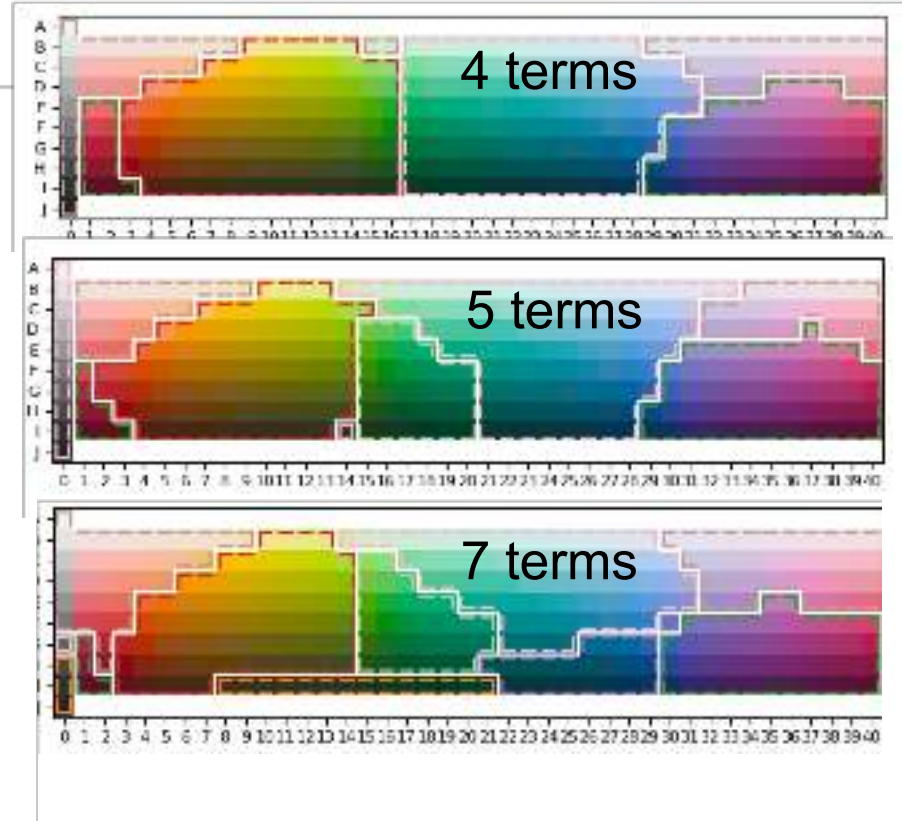



Complexity
of language



Variability in the
environment

Language evolution over time



Concluding remarks: learning (a language) to communicate efficiently

Mikael Kågebäck¹, Devdatt Dubhashi¹, Asad Sayeed²

¹Chalmers University of Technology, CSE

²University of Gothenburg, FLoV, CLASP



CHALMERS



GÖTEBORGS UNIVERSITET

Borges



An Argentinian philosopher and fiction writer. One of his stories mentions 'a certain Chinese Encyclopedia', the *Celestial Emporium of Benevolent knowledge*. It contains a classification of animals.

- those that belong to the emperor
- embalmed ones
- those that are trained
- suckling pigs
- mermaids
- fabulous ones
- stray dogs

Borges



... actually, it goes on.

- those that are included in the present classification
- those that tremble as if they are mad
- innumerable ones
- those drawn with a very fine camelhair brush
- others
- those that have just broken a flower vase
- those that from a long way off look like flies

Meaning

Think of lexical meaning as a kind of an amorphous blob.



(Lynda Benglis - New Museum)

Meaning

Think of lexical meaning as a kind of an amorphous blob.



(Lynda Benglis - New Museum)

Previous efforts to find a formal, universal ontology of features failed. . .

Meaning

... and yet human speech communities manage to come up with common lexical structure that everyone can more or less understand.



Meaning

... and yet human speech communities manage to come up with common lexical structure that everyone can more or less understand.



A partitioning of the space, if you will.

The case for colour



The case for colour



Do we have common partitionings across languages?

The case for colour

Colour is . . .

The case for colour

Colour is...

- ...literally an amorphous blob.

The case for colour

Colour is...

- ...literally an amorphous blob.
 - i.e., a space defined by a number of continuous variables with no obvious partitioning

The case for colour

Colour is...

- ...literally an amorphous blob.
 - i.e., a space defined by a number of continuous variables with no obvious partitioning
- ...subject to a linguistic universality claim.

The case for colour

Colour is...

- ...literally an amorphous blob.
 - i.e., a space defined by a number of continuous variables with no obvious partitioning
- ...subject to a linguistic universality claim.
 - Famous studies (e.g., Berlin and Kay, 1969) have held that there is a common colour hierarchy across languages.

The case for colour

Colour is...

- ...literally an amorphous blob.
 - i.e., a space defined by a number of continuous variables with no obvious partitioning
- ...subject to a linguistic universality claim.
 - Famous studies (e.g., Berlin and Kay, 1969) have held that there is a common colour hierarchy across languages.
 - A function of the number of “main” colour words.

The case for colour

Colour is...

- ...literally an amorphous blob.
 - i.e., a space defined by a number of continuous variables with no obvious partitioning
- ...subject to a linguistic universality claim.
 - Famous studies (e.g., Berlin and Kay, 1969) have held that there is a common colour hierarchy across languages.
 - A function of the number of “main” colour words.
 - i.e. languages with three colour words will have the same colour words, languages with four will have those three, etc.

Our research question:

How do speech communities
converge on a partition of semantic
spaces?

What we can show via colour



- RL results track human results in terms of KL-divergence loss and wellformedness.

What we can show via colour



- RL results track human results in terms of KL-divergence loss and wellformedness.
 - Our model has captured core aspects of human communicative behaviour.

What we can show via colour



- RL results track human results in terms of KL-divergence loss and wellformedness.
 - Our model has captured core aspects of human communicative behaviour.
- Human maps and RL maps are significantly similar, but human maps are still more similar to a CIELAB-based ideal partitioning.

What we can show via colour



- RL results track human results in terms of KL-divergence loss and wellformedness.
 - Our model has captured core aspects of human communicative behaviour.
- Human maps and RL maps are significantly similar, but human maps are still more similar to a CIELAB-based ideal partitioning.
 - CIELAB is perceptually uniform relative to human vision.

What we can show via colour



- RL results track human results in terms of KL-divergence loss and wellformedness.
 - Our model has captured core aspects of human communicative behaviour.
- Human maps and RL maps are significantly similar, but human maps are still more similar to a CIELAB-based ideal partitioning.
 - CIELAB is perceptually uniform relative to human vision.
 - Our agents do not explicitly represent human colour perception.

What we can show via colour

- RL results track human results in terms of KL-divergence loss and wellformedness.
 - Our model has captured core aspects of human communicative behaviour.
- Human maps and RL maps are significantly similar, but human maps are still more similar to a CIELAB-based ideal partitioning.
 - CIELAB is perceptually uniform relative to human vision.
 - Our agents do not explicitly represent human colour perception.
 - **Research question:** how much of the gap can we account by representing a more plausible perceptual system?

What we can show via colour



- Possible perceptual issue: light/dark distinctions.

What we can show via colour



- Possible perceptual issue: light/dark distinctions.
 - RL has difficulty replicating this distinction.

What we can show via colour

- Possible perceptual issue: light/dark distinctions.
 - RL has difficulty replicating this distinction.
 - **Model improvement?** Lightness variable in CIELAB space needs to be weighted differently?

What we can show via colour

- Possible perceptual issue: light/dark distinctions.
 - RL has difficulty replicating this distinction.
 - **Model improvement?** Lightness variable in CIELAB space needs to be weighted differently?
- RL maps share behaviour of human maps in terms of partition stability as colour term numbers increase.

What we can show via colour

- Possible perceptual issue: light/dark distinctions.
 - RL has difficulty replicating this distinction.
 - **Model improvement?** Lightness variable in CIELAB space needs to be weighted differently?
- RL maps share behaviour of human maps in terms of partition stability as colour term numbers increase.
 - RL agents tend to sub-partition existing partitions rather than rework the partitions as they use more words.

What our model does for studying semantic change

A framework for hypothesis generation and simulation.

What our model does for studying semantic change

A framework for hypothesis generation and simulation.

- 1 Exploring the trade-offs between communicative factors and cognitive architectures in semantic development.

What our model does for studying semantic change

A framework for hypothesis generation and simulation.

- 1 Exploring the trade-offs between communicative factors and cognitive architectures in semantic development.
- 2 Modeling semantic change: we expect change to be represented as motion along a frontier of optimally efficient systems.

What our model does for studying semantic change

A framework for hypothesis generation and simulation.

- 1 Exploring the trade-offs between communicative factors and cognitive architectures in semantic development.
- 2 Modeling semantic change: we expect change to be represented as motion along a frontier of optimally efficient systems.
- 3 Control over high-level variables: change noise parameter to represent environmental conditions.

Where we go from here



- Multi-agent and large population simulations—requires more elaborate model of communicative encounters.

Where we go from here



- Multi-agent and large population simulations—requires more elaborate model of communicative encounters.
- Other domains: approximate number, family relations—even large-scale semantic ontologies.

Where we go from here

- Multi-agent and large population simulations—requires more elaborate model of communicative encounters.
- Other domains: approximate number, family relations—even large-scale semantic ontologies.
- Complexifying the agents and environment:

Where we go from here

- Multi-agent and large population simulations—requires more elaborate model of communicative encounters.
- Other domains: approximate number, family relations—even large-scale semantic ontologies.
- Complexifying the agents and environment:
 - More realistic perceptual architecture (in colour domain).

Where we go from here



- Multi-agent and large population simulations—requires more elaborate model of communicative encounters.
- Other domains: approximate number, family relations—even large-scale semantic ontologies.
- Complexifying the agents and environment:
 - More realistic perceptual architecture (in colour domain).
 - Model environmentally-induced preferences (“warm/cold” colour distinction).

Where we go from here

- Multi-agent and large population simulations—requires more elaborate model of communicative encounters.
- Other domains: approximate number, family relations—even large-scale semantic ontologies.
- Complexifying the agents and environment:
 - More realistic perceptual architecture (in colour domain).
 - Model environmentally-induced preferences (“warm/cold” colour distinction).
- Connection to other research at CLASP?

Where we go from here



- Multi-agent and large population simulations—requires more elaborate model of communicative encounters.
- Other domains: approximate number, family relations—even large-scale semantic ontologies.
- Complexifying the agents and environment:
 - More realistic perceptual architecture (in colour domain).
 - Model environmentally-induced preferences (“warm/cold” colour distinction).
- Connection to other research at CLASP?
Hypothesizing expected behaviour in corpora.

Thanks!





CHALMERS
UNIVERSITY OF TECHNOLOGY