Learning and Maintaining a Lexicon for Situated Interaction

David Schlangen
Bielefeld University

http://www.dsg-bielefeld.de/talks/gothenburg-2017

Universität Bielefeld

dialogue systems group [unibi]

CITEC
Cognitive Interaction Technology
Cluster of Excellence
Bielefeld University

Situated Interaction

… agents are co-present,

Situated Interaction

… agents are co-present, can make their physical environment the topic,

Situated Interaction

… agents are co-present, can make their physical environment the topic, can make use of variety of resources (language, body, environment).

A: Was there a Rottweiler?
B: Yes.

knowledge from testimony

- U: *Find the Rottweiler.*
- S: Picture 4.
- U: *Explain*.
- S: I have never seen a Rottweiler, but
  I know that it is a type of dog.
  4 is the only dog.

- U: *Find the Rottweiler.*
- S: Picture 4.
- U: *Explain*.
- S: I have never seen a Rottweiler, but
  I know that it is a type of bicycle.
  4 is the only bicycle.

- U: *Find the Rottweiler.*
- S: Picture 4.
- U: *Explain*.
- S: I have never seen a Rottweiler, but I know that it is a type of dog. 4 is the only dog.

language / language relation

language / world relation

Diego Marconi 1997, Lexical Competence

Desiderata:

* A lexicon that provides these *referential* and *inferential* links.
* A way to use it to resolve and generate references, and to generate "meta-conceptual" interaction.
* A plausible story on how it can be learned.

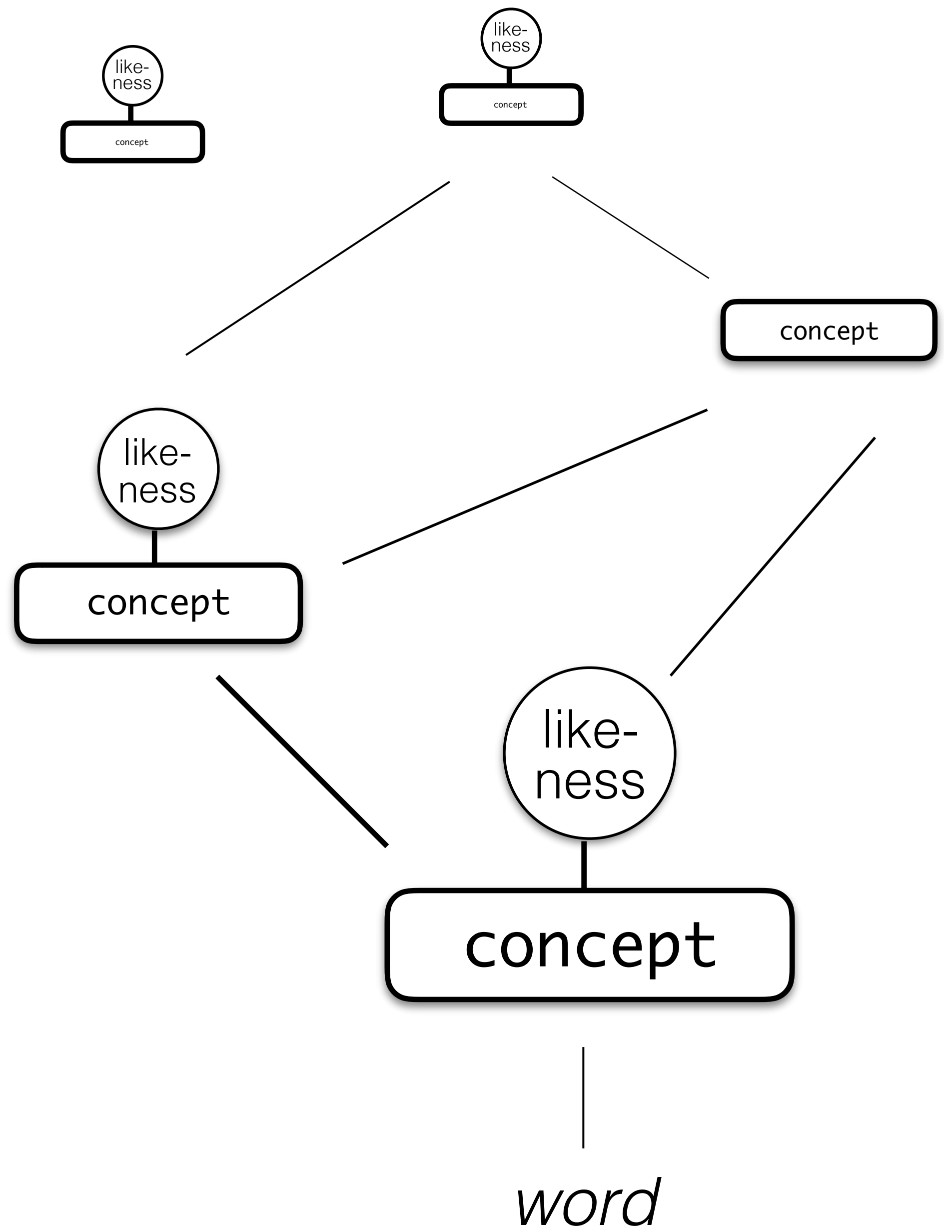
like-
ness
concept
like-
ness
concept
concept
like-
ness
concept
like-
ness
concept
word

demonstration:

“*This* is a wampimuk.”

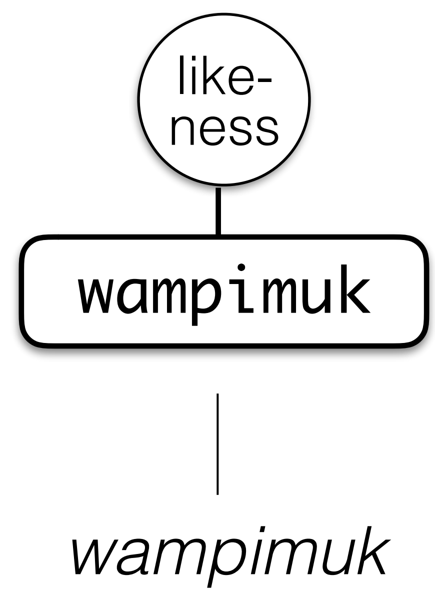

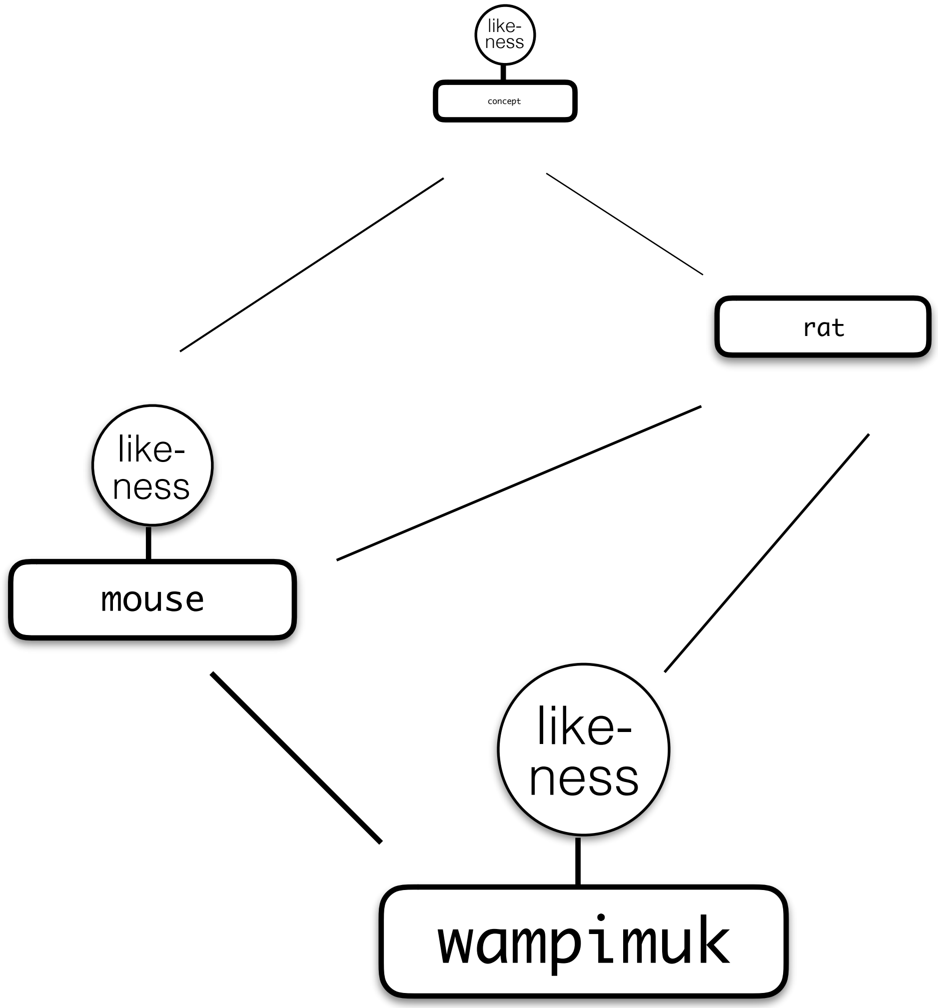

demonstration:

"*This* is a wampimuk."

expl. definition:

"The wampimuk is a small, mouse-like mammal native to the Trentino area."

```
is_a(w, m)
lives_in(w, T)
…
```

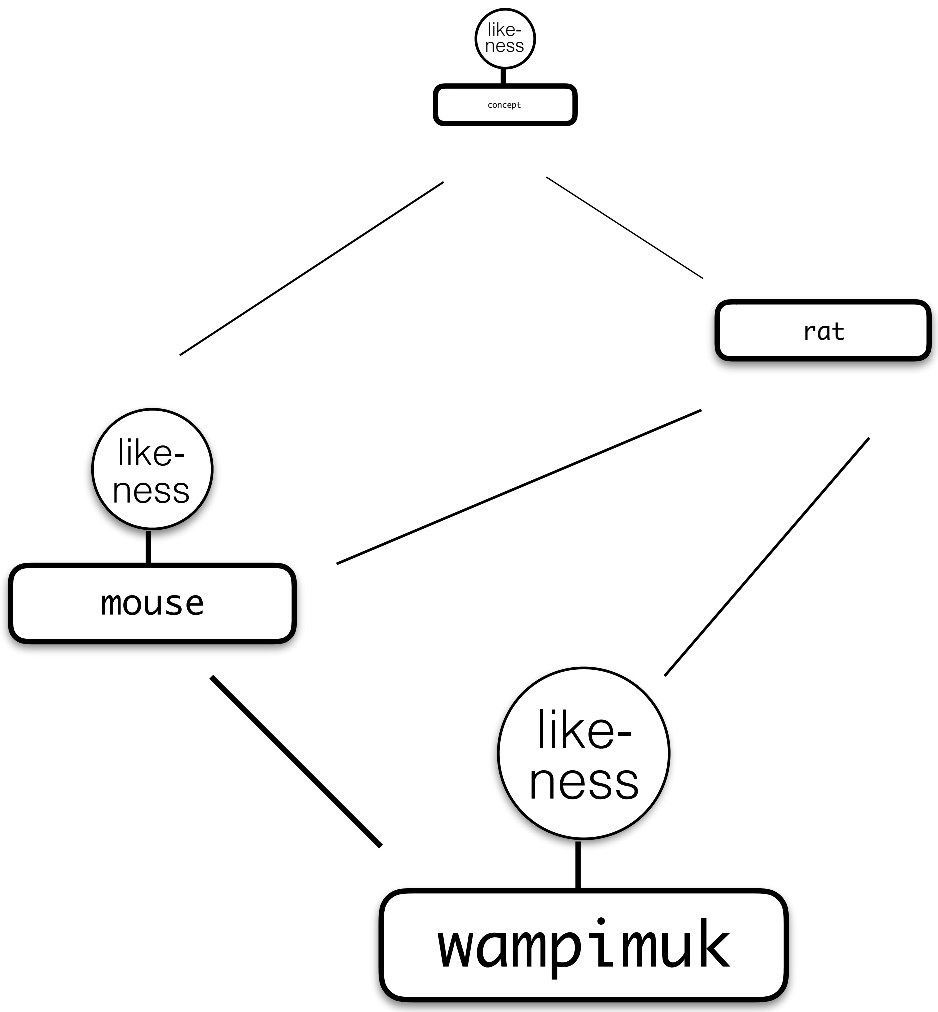

demonstration:

"*This* is a wampimuk."

expl. definition:

"The wampimuk is a small, mouse-like mammal native to the Trentino area."

impl. definition:

"… the cute wampimuk squeaked…" "… a mouse, a wampimuk and a …" "… she saw a wampimuk sitting on…" …

```
is_a(w, m)
lives_in(w, T)
…
```

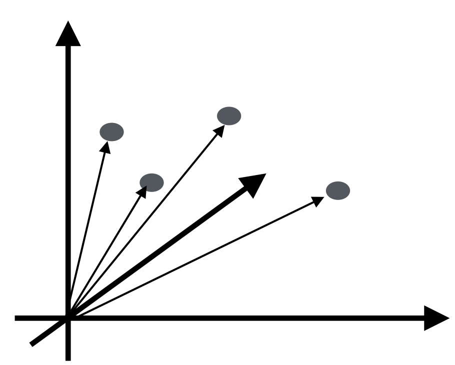

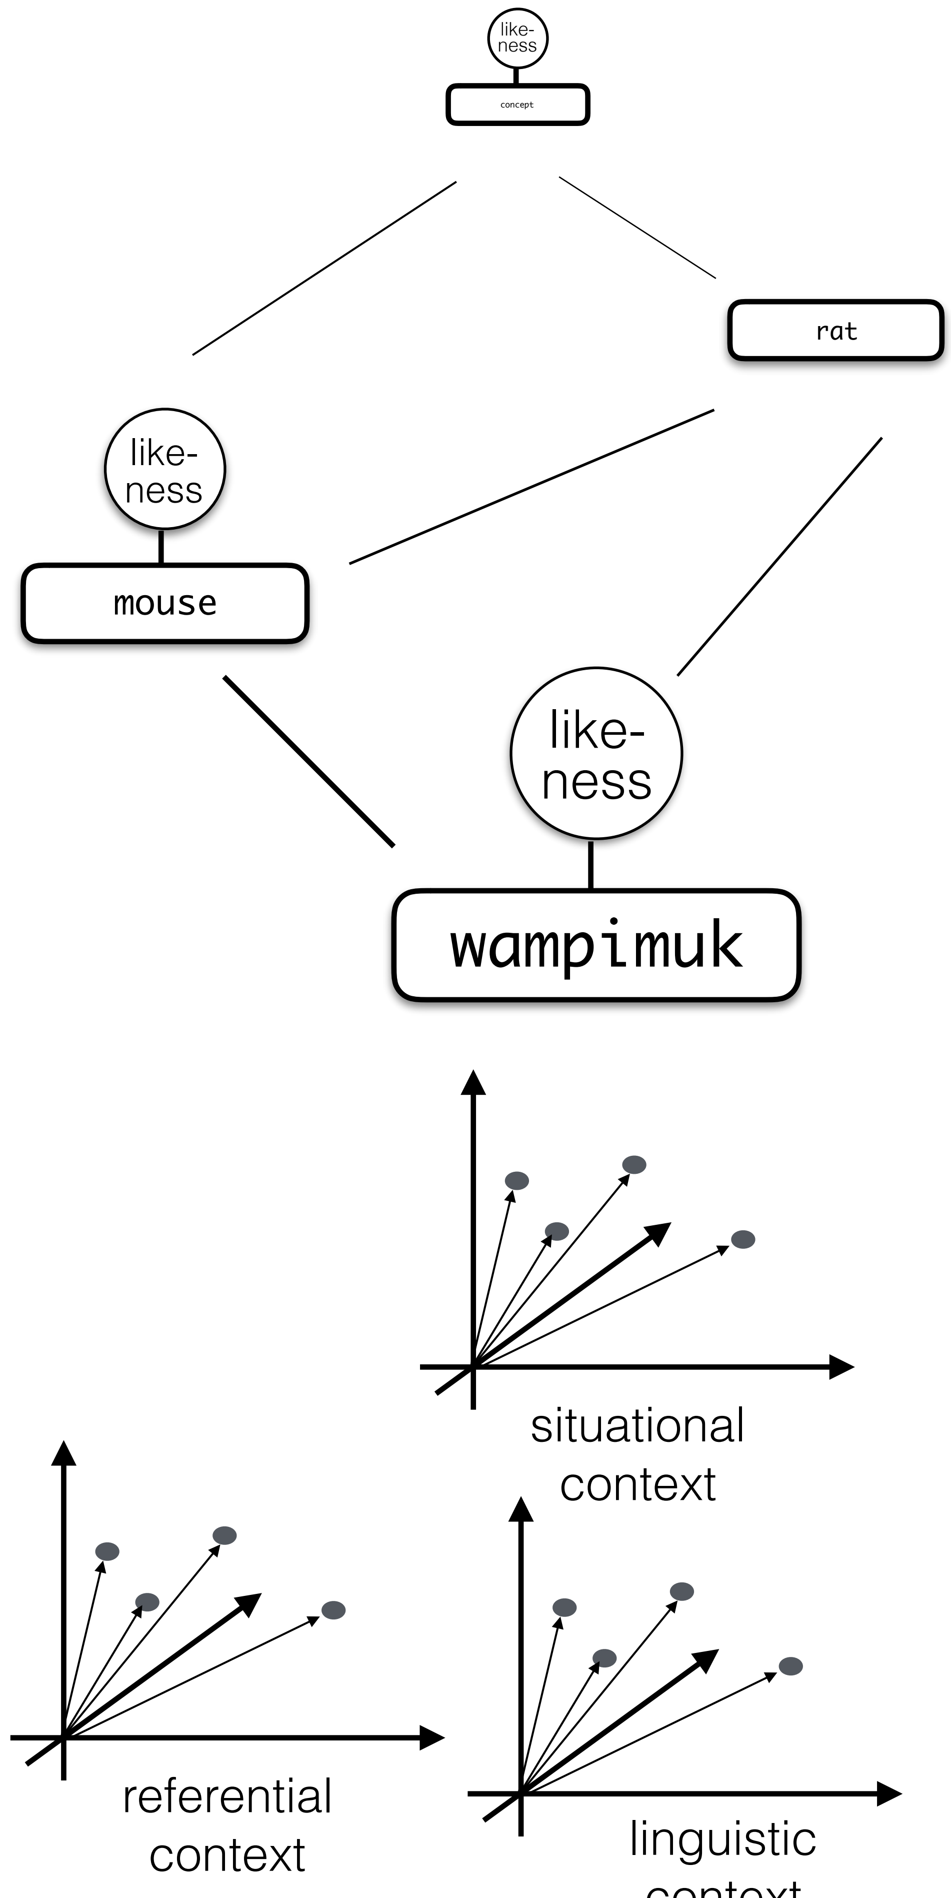


```
is_a(w, m)
lives_in(w, T)
…
```

__demonstration:__

"*This* is a wampimuk."

__expl. definition:__

"The wampimuk is a small, mouse-like mammal native to the Trentino area."

__impl. definition:__

"… the cute wampimuk squeaked…" "… a mouse, a wampimuk and a …" "… she saw a wampimuk sitting on…" …

Overview

- Motivation: Knowledge from Testimony
- The Lexicon: Referential & Inferential Knowledge
- **Referential Knowledge: *Likeness***
  - Acquisition from Referential Interaction
  - Application in Reference Resolution
  - Application in Reference Generation
- **Inferential Knowledge**
  - … from Referential Knowledge / Referential Interaction
  - … from Definitions
- **Towards Justifying Concepts**

Overview

- Motivation: Knowledge from Testimony
- The Lexicon: Referential & Inferential Knowledge
- **Referential Knowledge: *Likeness***
  - Acquisition from Referential Interaction
  - Application in Reference Resolution
  - Application in Reference Generation
- **Inferential Knowledge**
  - … from Referential Knowledge / Referential Interaction
  - … from Definitions
- **Towards Justifying Concepts**

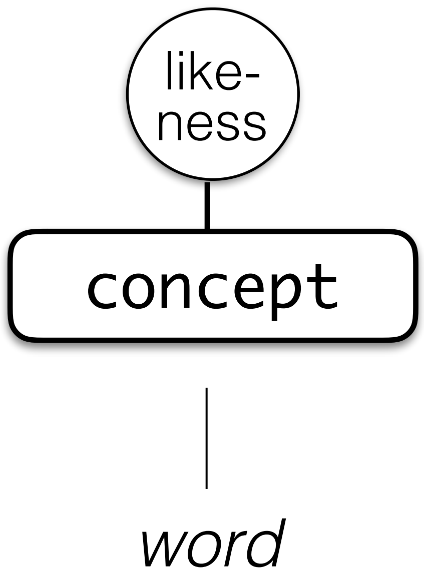

http://www.dsg-bielefeld.de/talks/gothenburg-2017

Overview

- Motivation: Knowledge from Testimony
- The Lexicon: Referential & Inferential Knowledge
- **Referential Knowledge: *Likeness***
    - **Acquisition from Referential Interaction**
    - Application in Reference Resolution
    - Application in Reference Generation
- **Inferential Knowledge**
    - … from Referential Knowledge / Referential Interaction
    - … from Definitions
- **Towards Justifying Concepts**

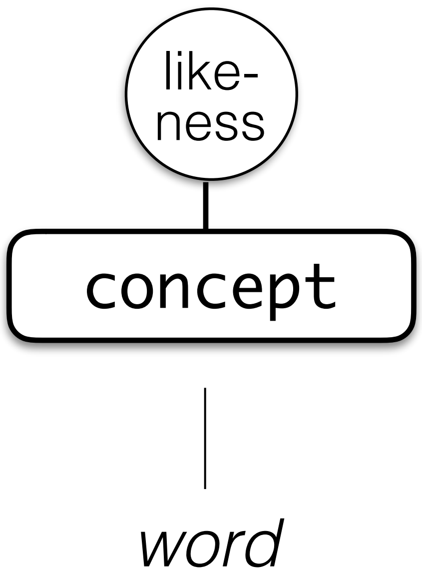

http://www.dsg-bielefeld.de/talks/gothenburg-2017

Referential Interaction

primarily: ReferIt corpus (Berg *et al.*)

A and **B** play a game.
A sees image with highlight on object, **B** without.
A says: "person left".
B clicks on object.

Result: pairs of object in scene and ref-exp, filtered for success.

- Referring expressions, not labels!
  - No closed-world assumption.
  - No pre-conceived tagset.

Referential Interaction

primarily: ReferIt corpus (Berg *et al.)*

A and **B** play a game.
A sees image with highlight on object, **B** without.
A says: “person left”.
B clicks on object.

Result: pairs of object in scene and ref-exp, filtered for success.

- Referring expressions, not captions!
    - Discriminative, not exhaustive.
    - Minimal, not exhaustive.

Referential Interaction

primarily: ReferIt corpus (Berg *et al.)*

A and **B** play a game.
A sees image with highlight on object, **B** without.
A says: “person left”.
B clicks on object.

Result: pairs of object in scene and ref-exp, filtered for success.

- Demonstration: “*this* is a [person left]”

Referential Interaction

primarily: ReferIt corpus (Berg *et al.)*

A and **B** play a game.
A sees image with highlight on object, **B** without.
A says: “person left”.
B clicks on object.

Result: pairs of object in scene and ref-exp, filtered for success.

- ReferIt corpus (Kazemzadeh *et al.* 2014): 20k images (SAIAPR, [Escalante *et al.* 2010]), 120k referring expressions

- MSCOCO (Lin *et al.* 2014): 27k images, 100k region descriptions (Mao *et al.* 2015) + 140k referring expressions (Berg *et al.* 2015) + 140k (non-positional) ref exp (Yu *et al.* 2016)

like-
ness

wrench

https://it.wikipedia.org/wiki/File:Alan_Key_2.jpg
https://commons.wikimedia.org/wiki/File:Claw-hammer.jpg
https://commons.wikimedia.org/wiki/File:Screw_Driver_display.jpg

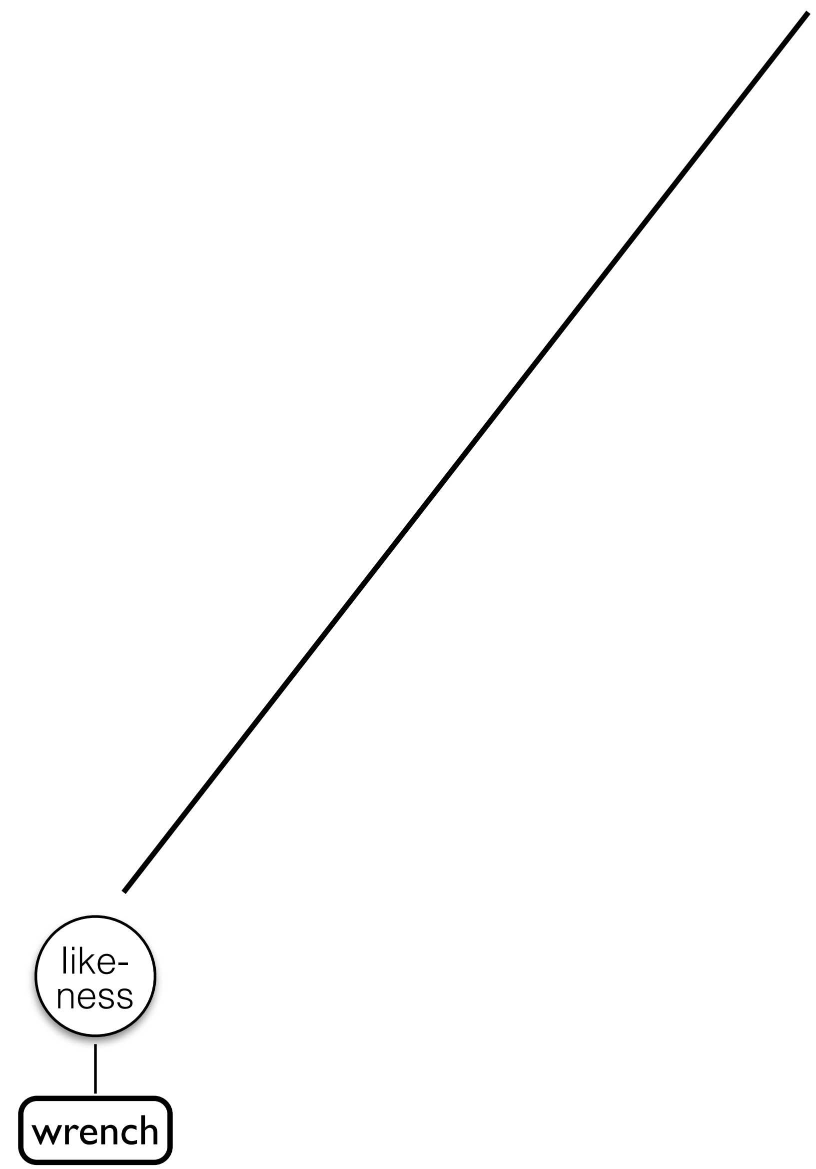

https://it.wikipedia.org/wiki/File:Alan_Key_2.jpg
https://commons.wikimedia.org/wiki/File:Claw-hammer.jpg
https://commons.wikimedia.org/wiki/File:Screw_Driver_display.jpg

The “words as classifiers” approach

(Harnad 1990), The Symbol Grounding Problem:
“[H]ow can the semantic interpretation of a formal symbol system be made intrinsic to the system, rather than just parasitic on the meanings in our heads?
*“[…] invariant features […] that will reliably distinguish a member of a category from any nonmembers […] Let us call the output of this **category-specific feature detector** the* categorical reprs.*”*

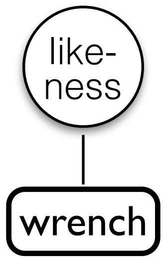

Deb Roy (Roy et al. 2002, 2005),
Siebert & Schlangen (2008),

Larsson (2013 / ’15),

Kennington & Schlangen (2015),
Schlangen et al. (2016)

https://it.wikipedia.org/wiki/File:Alan_Key_2.jpg
https://commons.wikimedia.org/wiki/File:Claw-hammer.jpg
https://commons.wikimedia.org/wiki/File:Screw_Driver_display.jpg

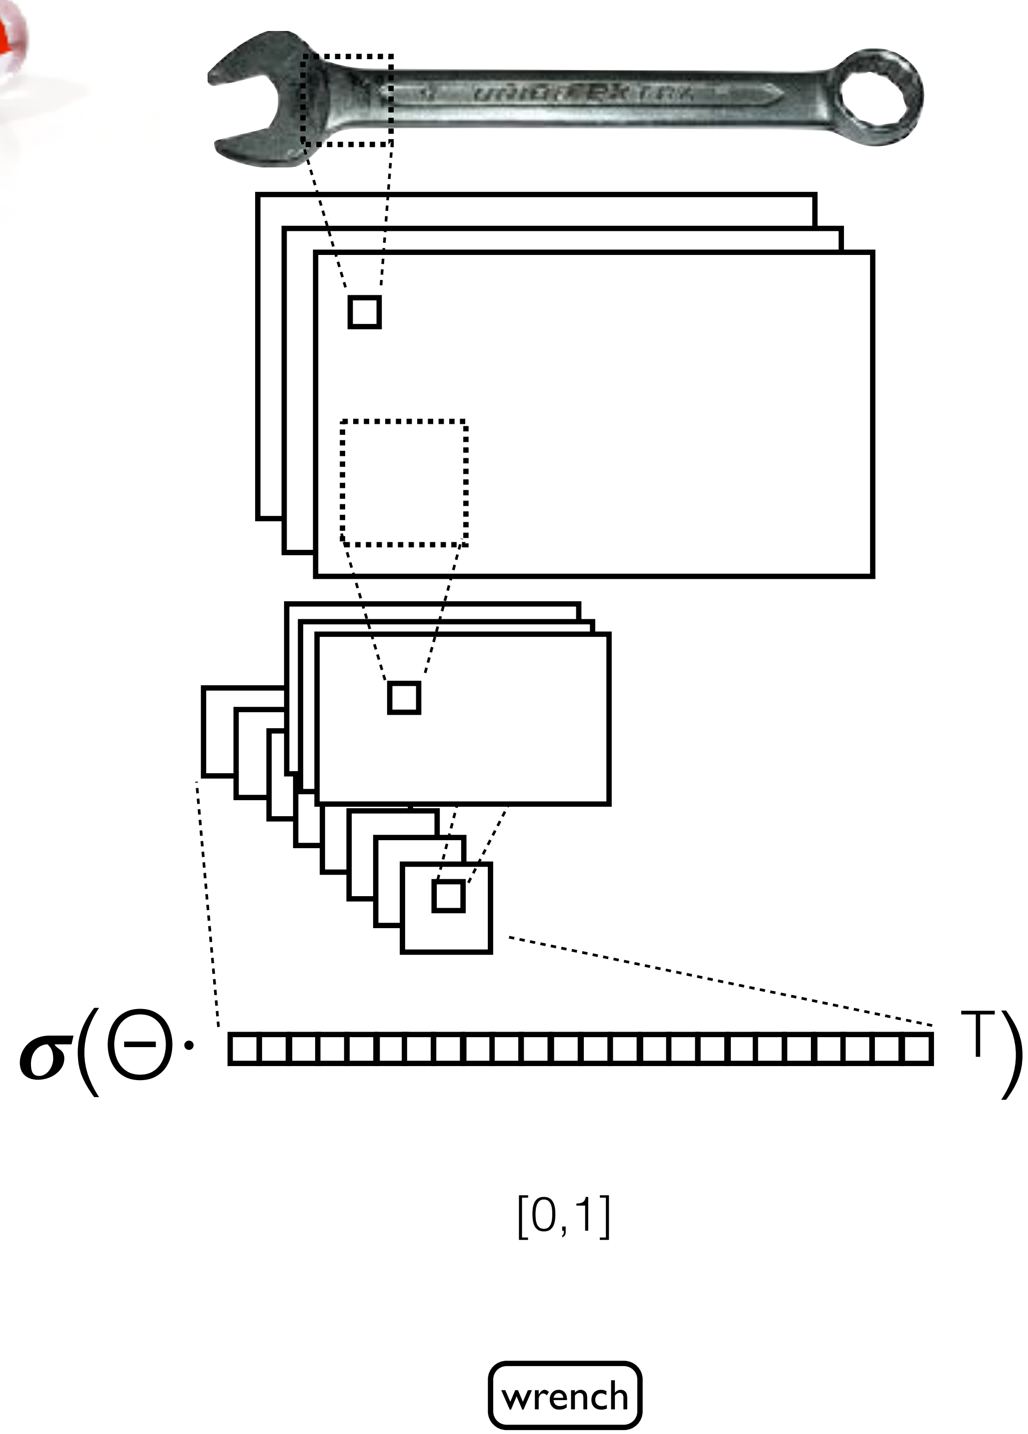

GoogLeNet; deep convolutional neural network (Szegedy *et al.* 2015)

1024 + 7 positional features

in humans, learned over phylogenetic time?

L1-regulated logistic regression, cross entropy loss function, SGD

Training

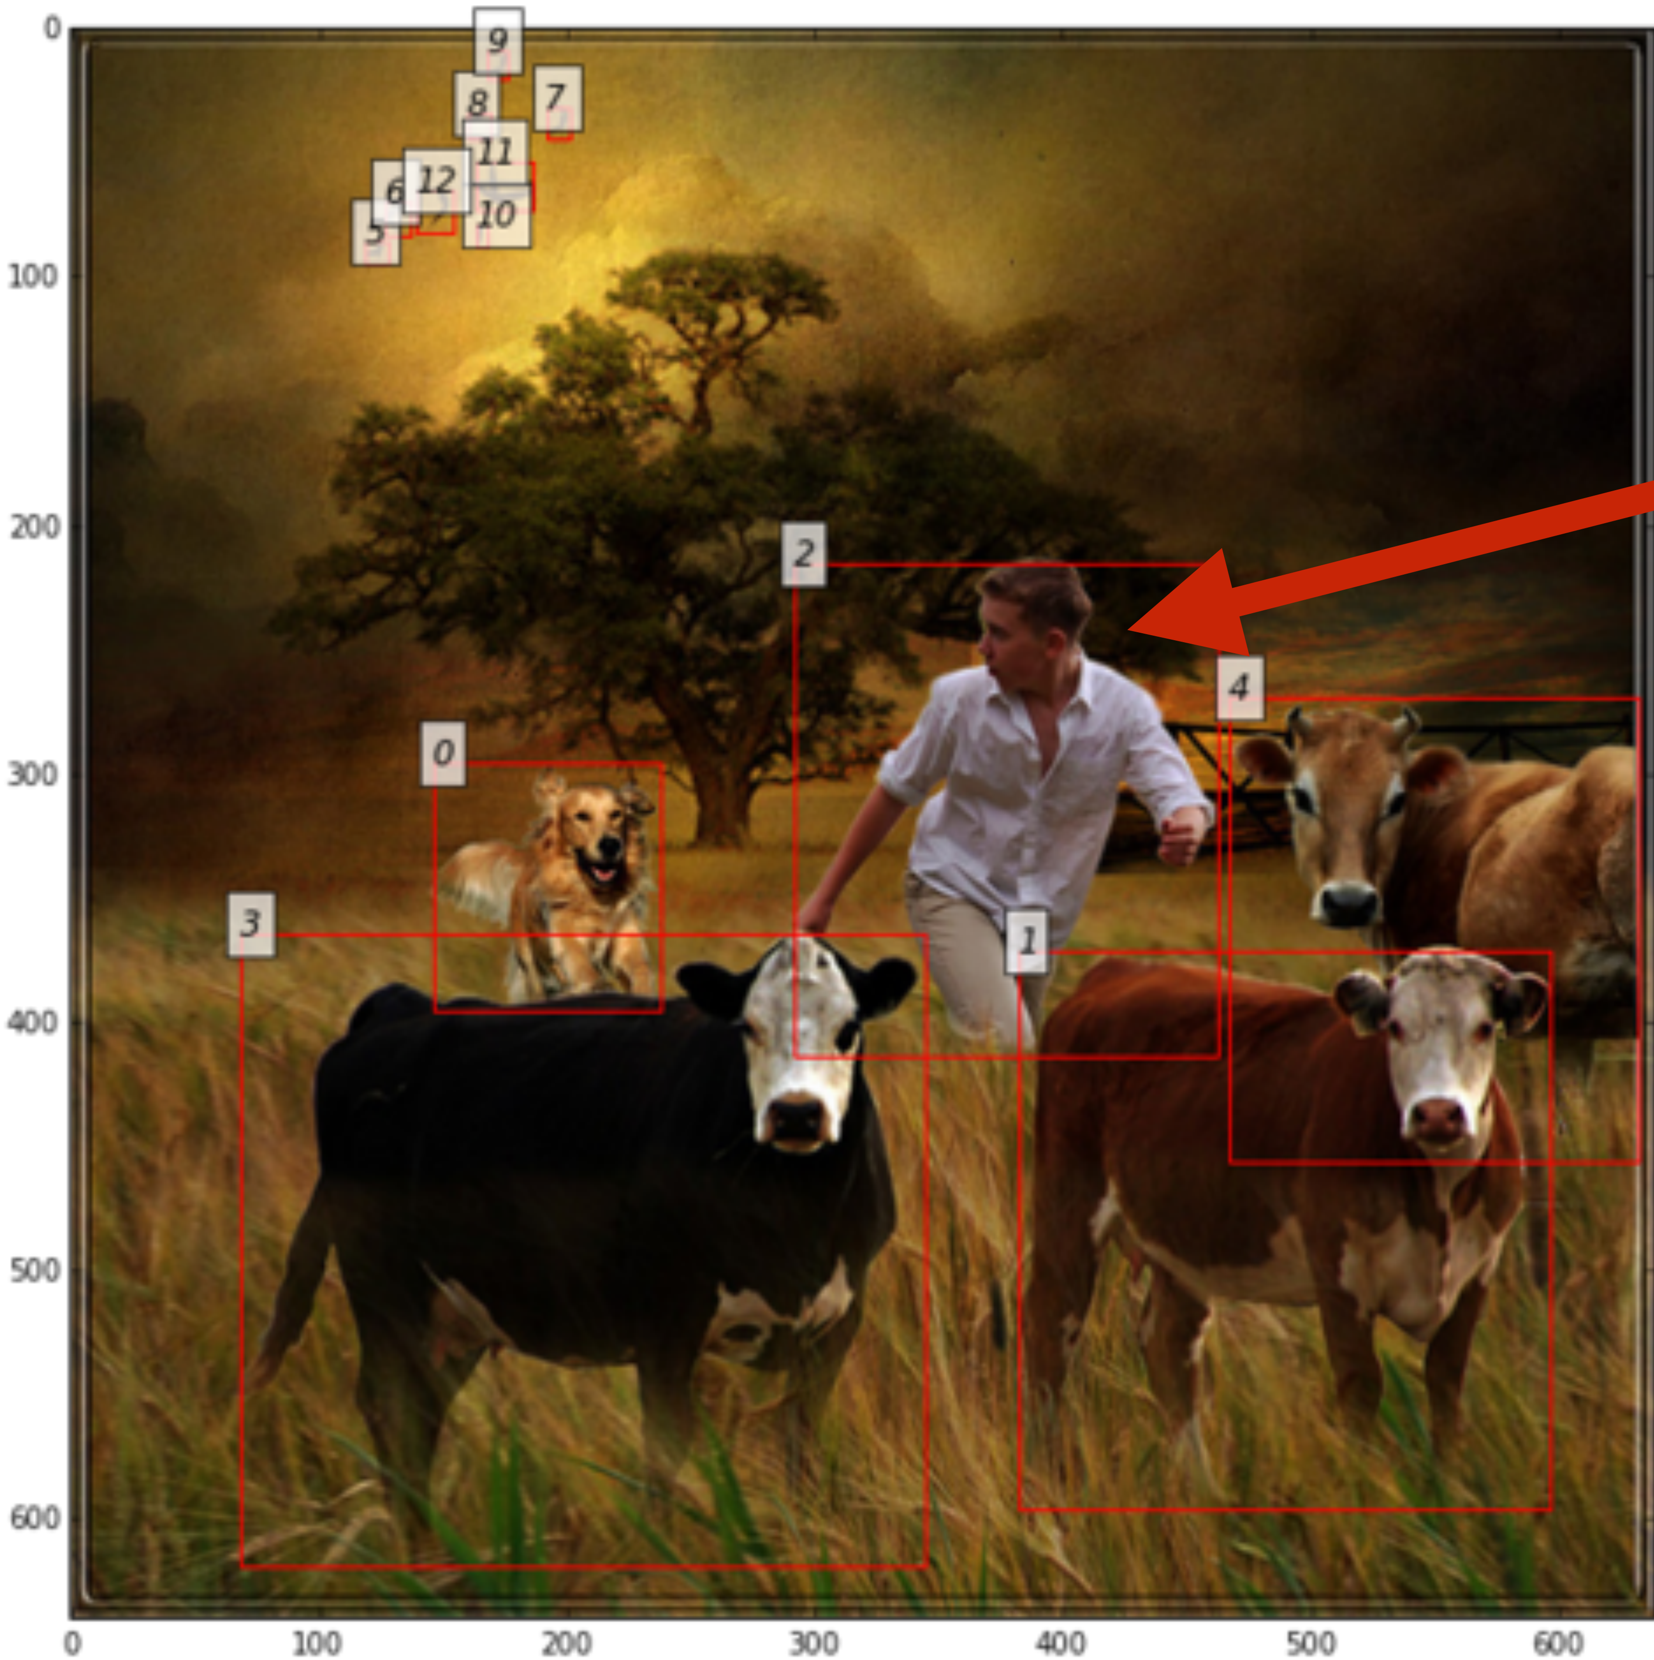

Guy with white shirt

Training

Guy
with
white
shirt

¬Guy
¬with
¬white
¬shirt

Training

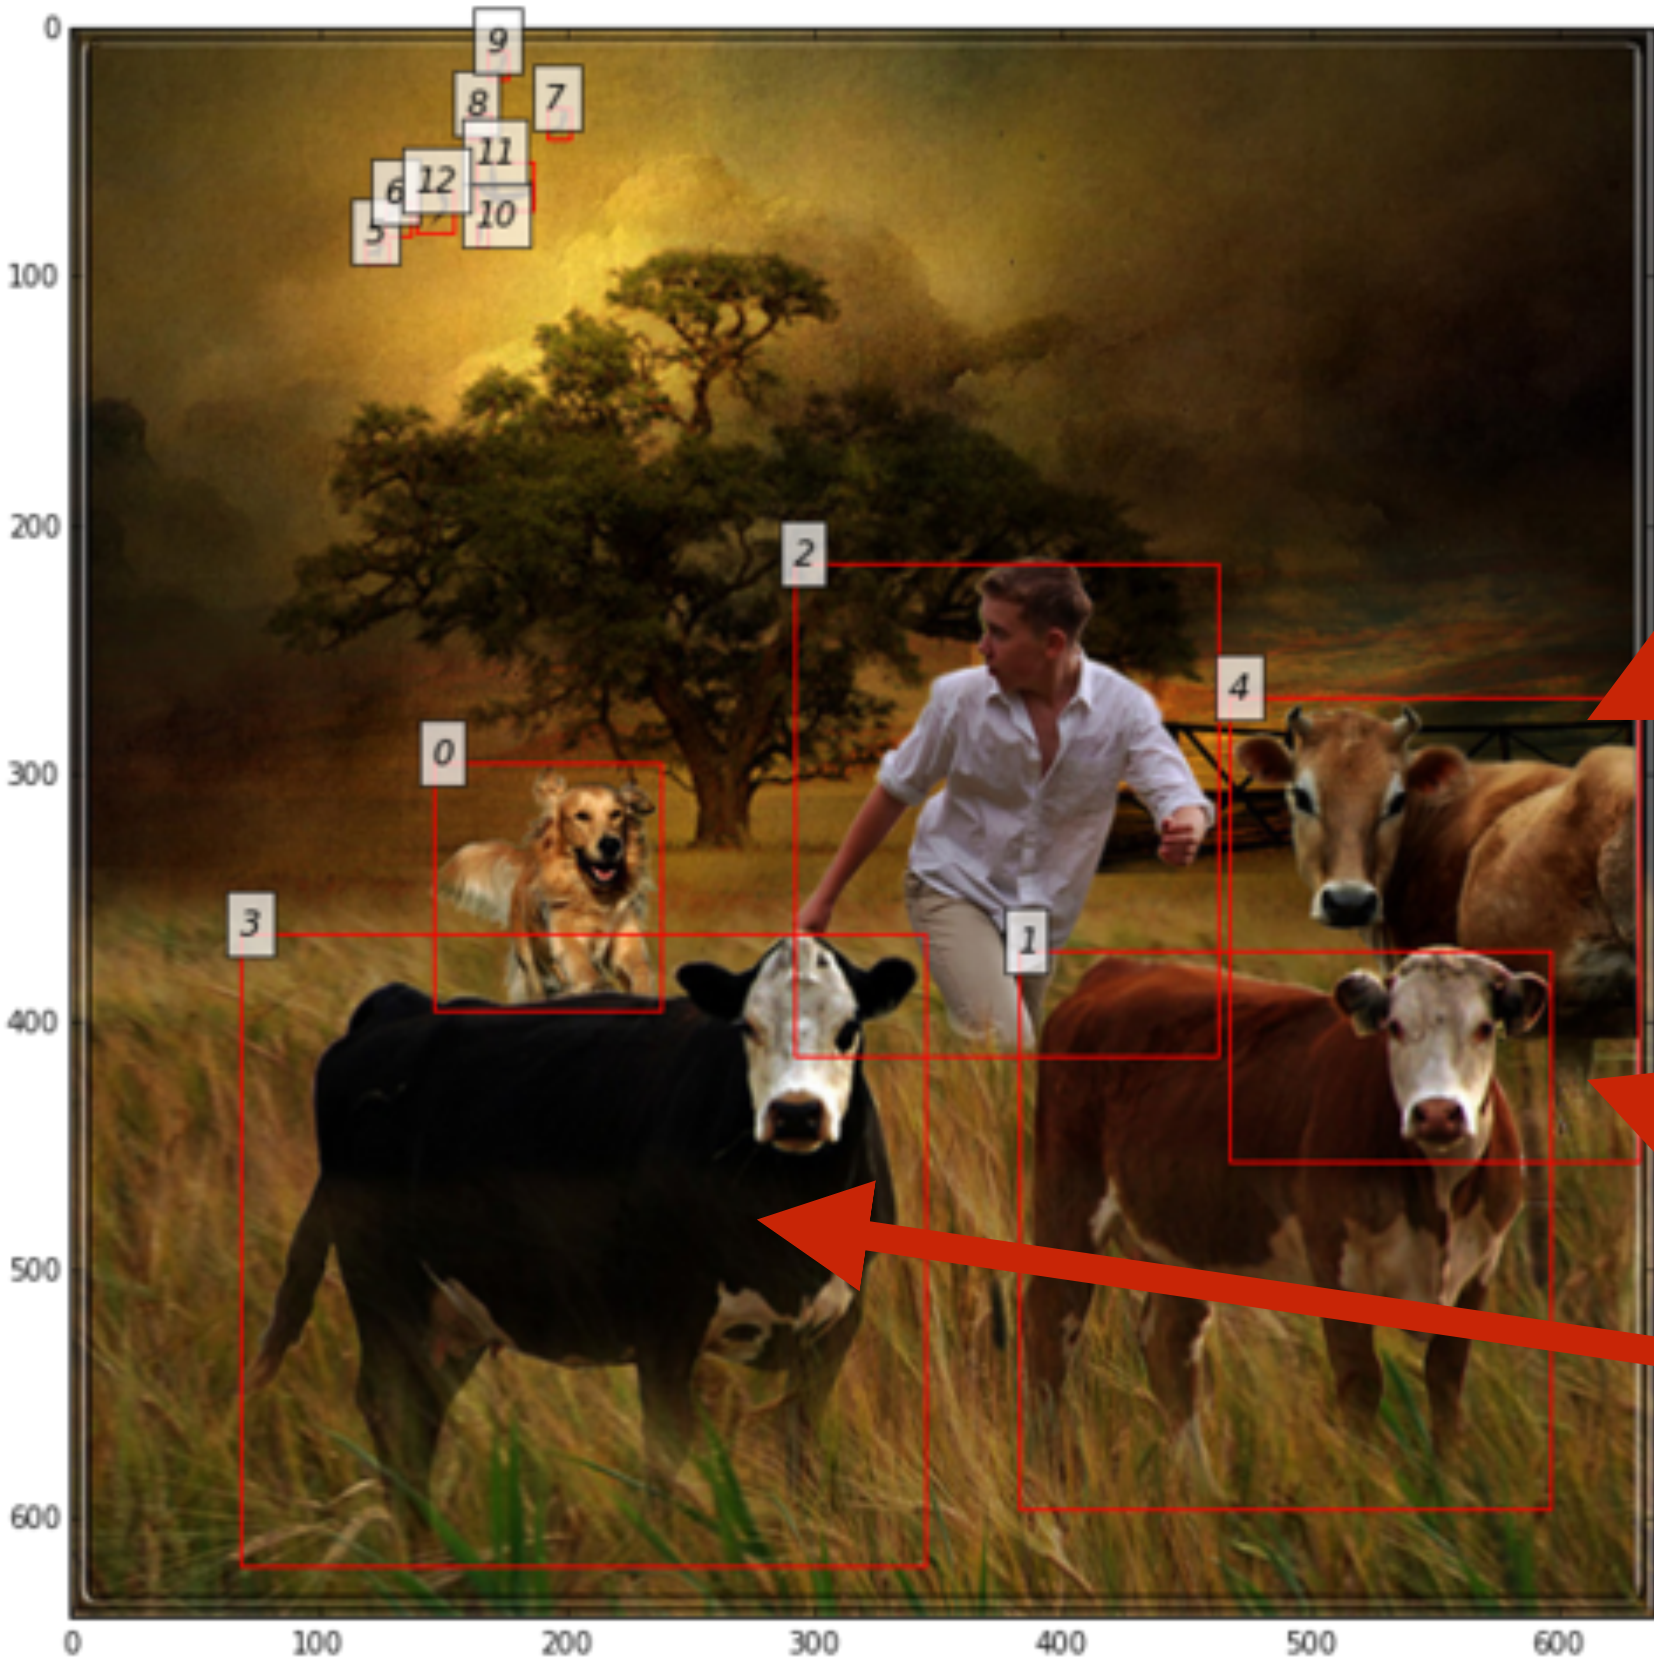

Cow
right

¬Cow
¬right

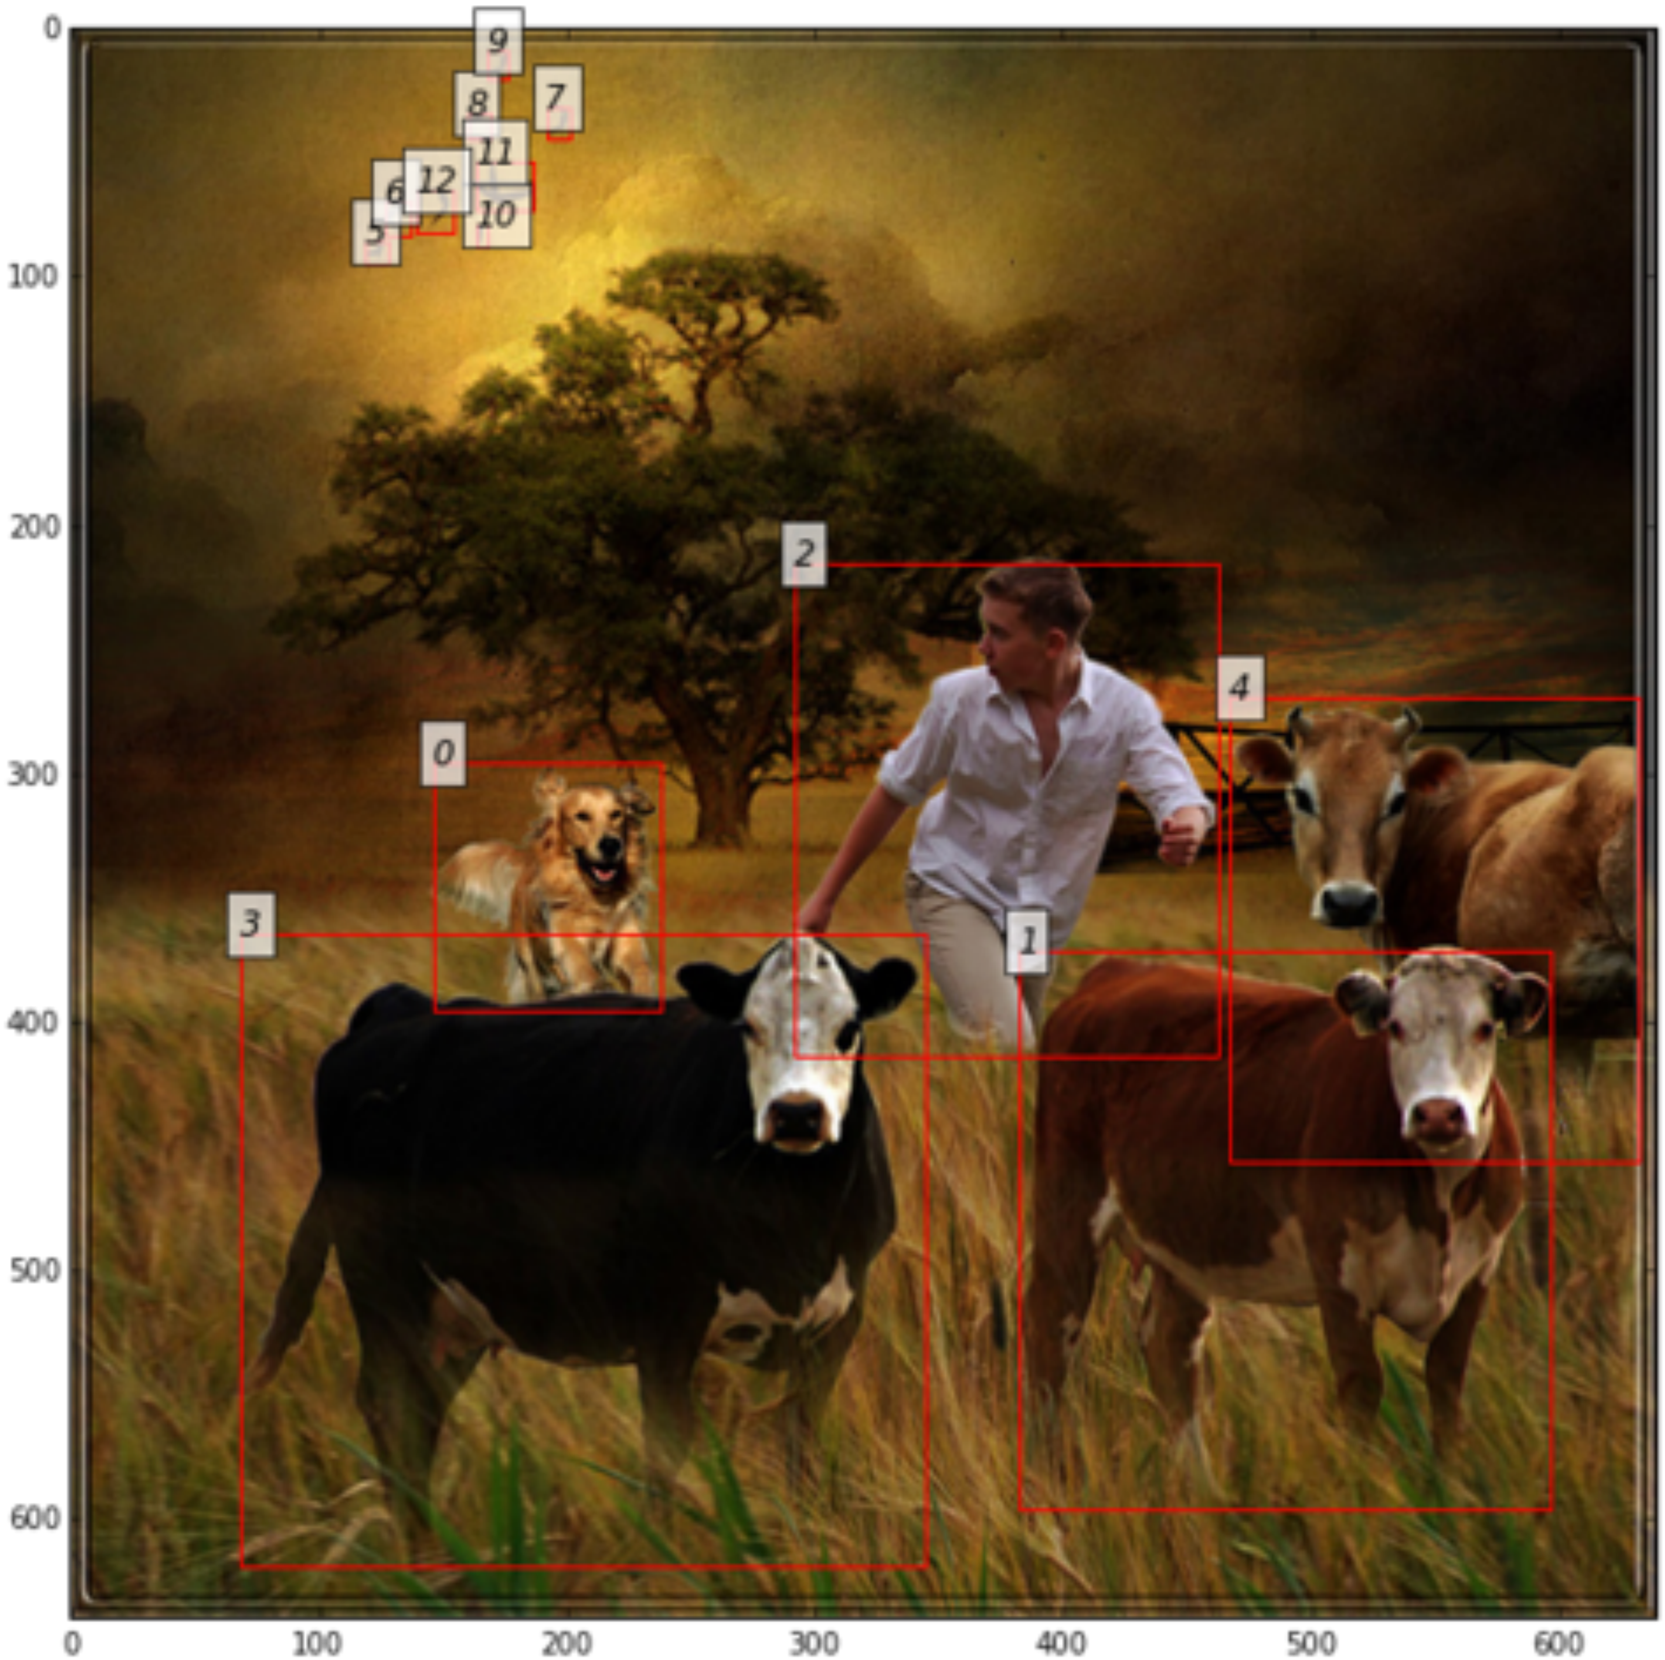

Training

All words separately.

One classifier per word.

Here, trained in batch mode, but could be done incrementally.

Tried taking neg inst. from same scene, and randomly from whole set.

Training

- condition: min. 40 positive training instances
- resulting vocab. size:
  - SAIAPR: 429
  - RefCoco: 503
  - SAIAPR + RefCoco: 783
  - SAIAPR + RefCoco + RefCoco_{+}: 1,174

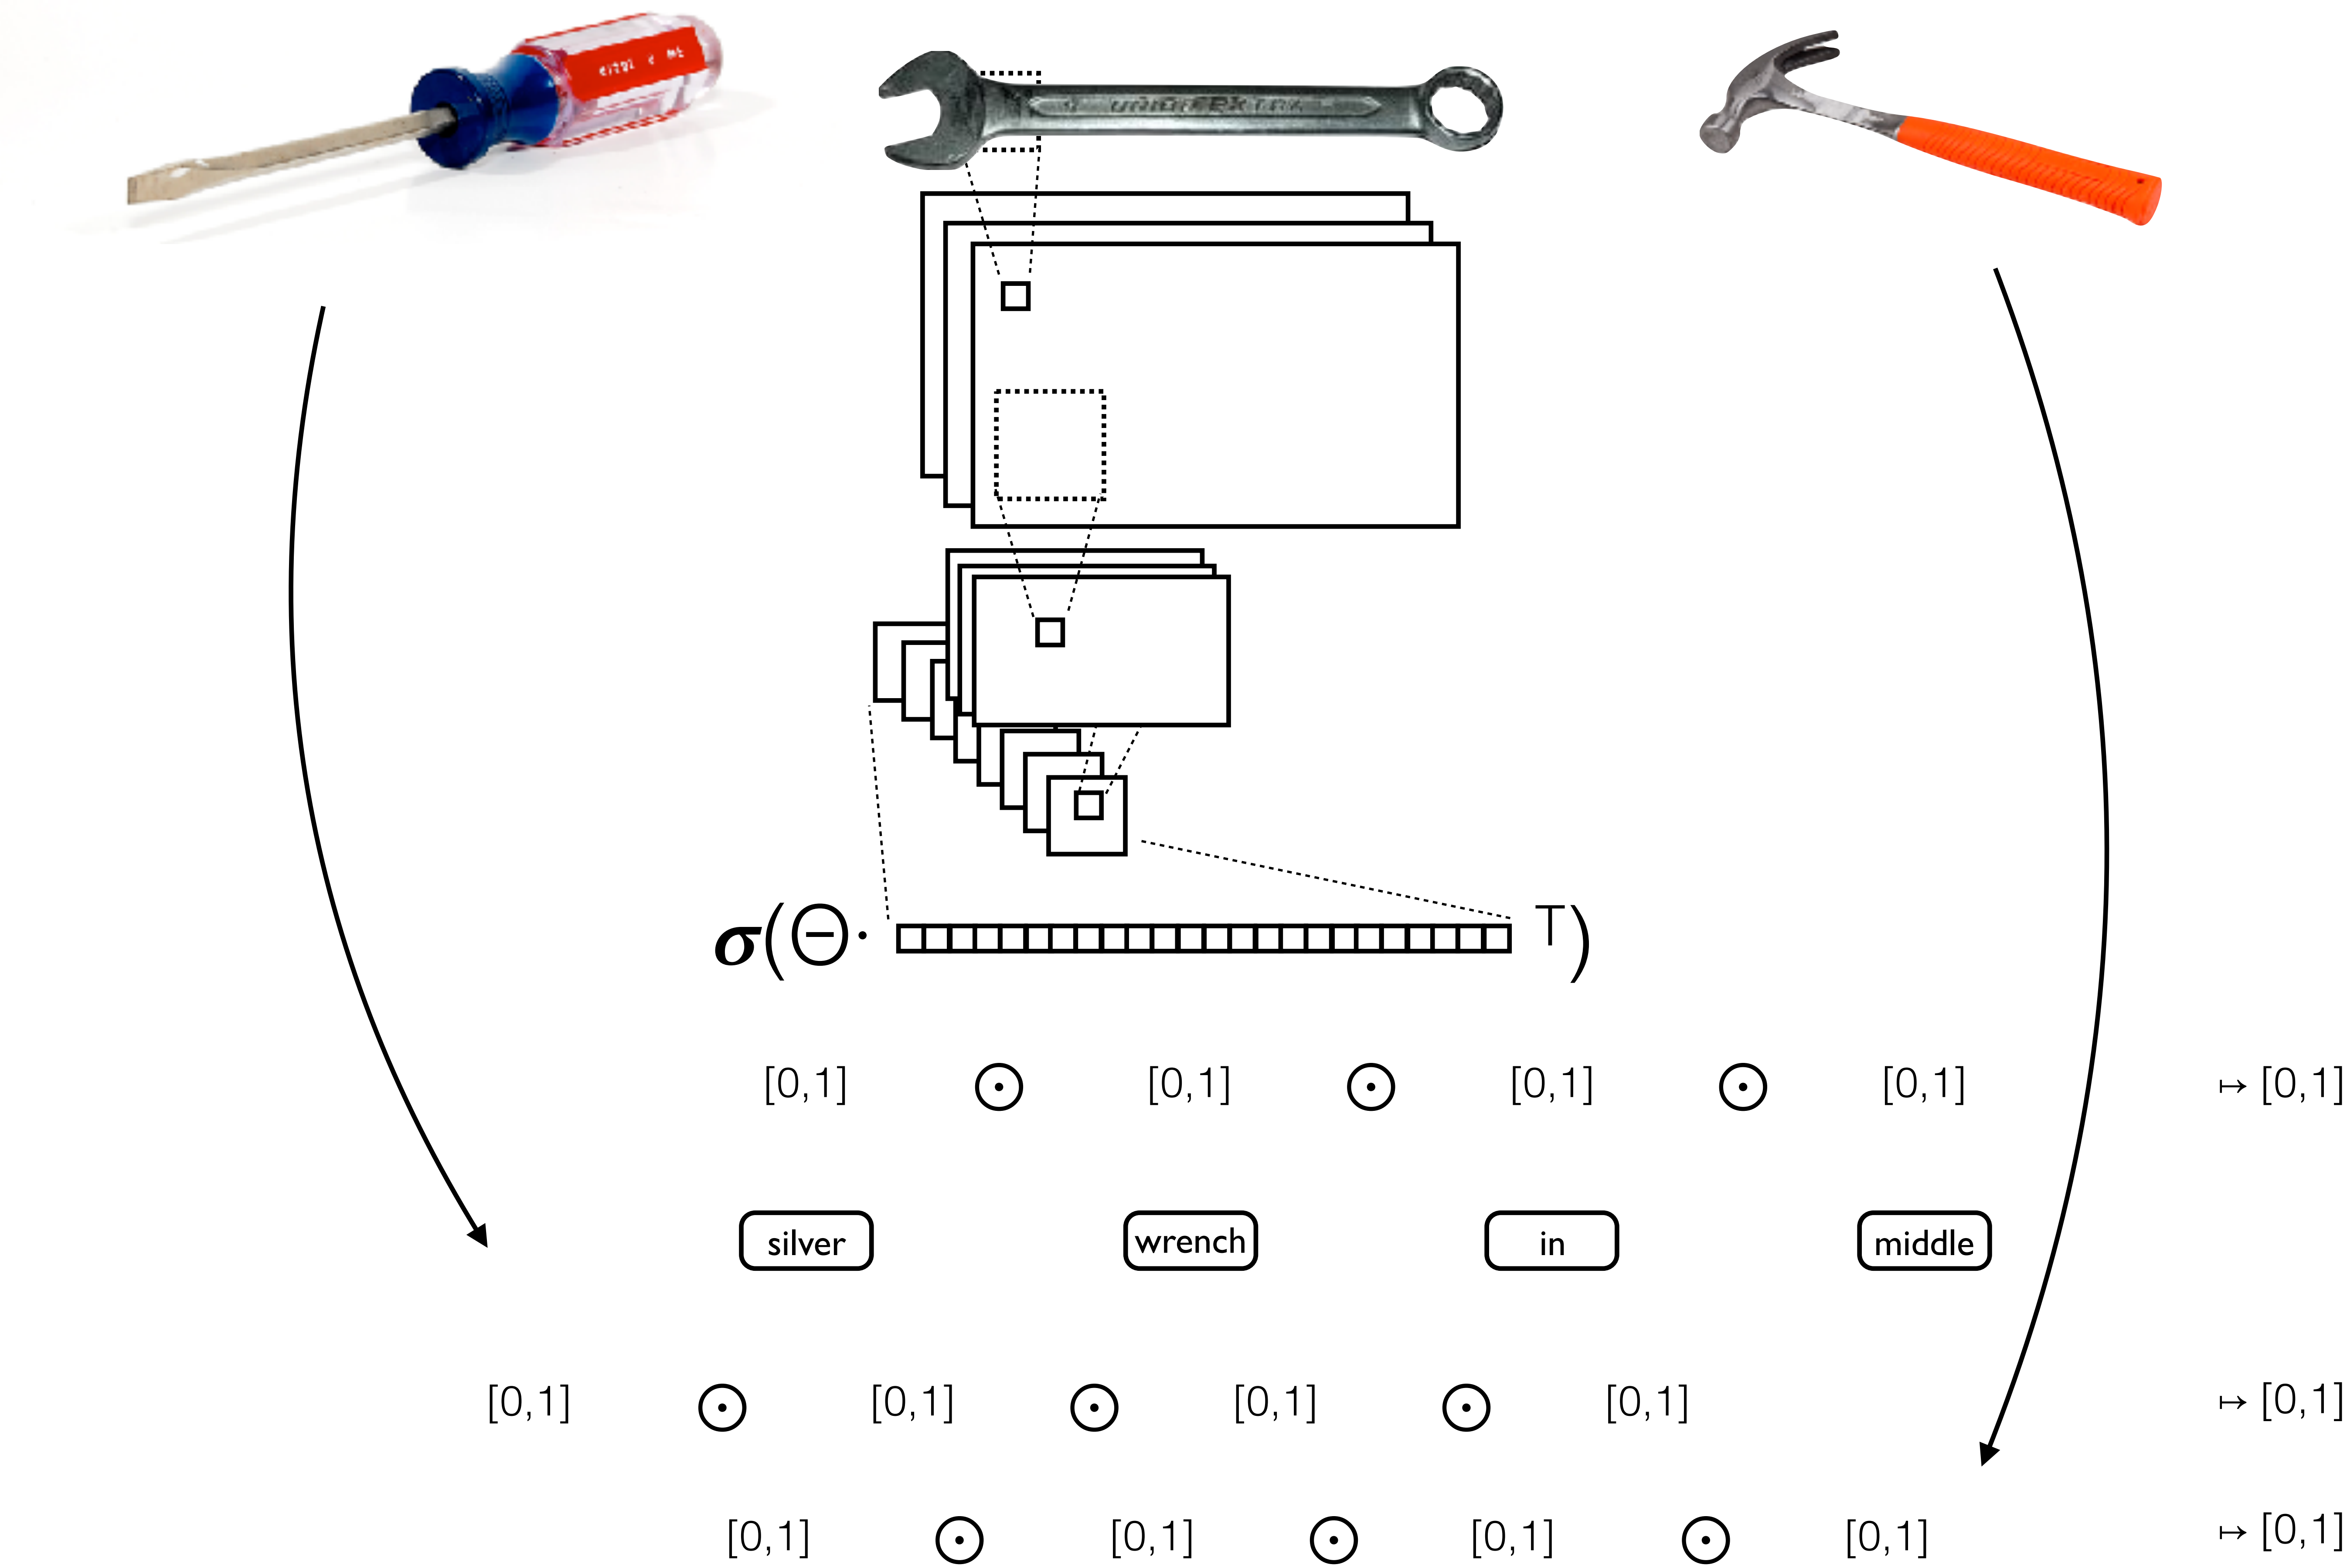
σ(Θ· T)
[0,1] ⊙ [0,1] ⊙ [0,1] ⊙ [0,1] ↦ [0,1]
silver
wrench
in
middle
[0,1] ⊙ [0,1] ⊙ [0,1] ⊙ [0,1] ↦ [0,1]
[0,1] ⊙ [0,1] ⊙ [0,1] ⊙ [0,1] ↦ [0,1]

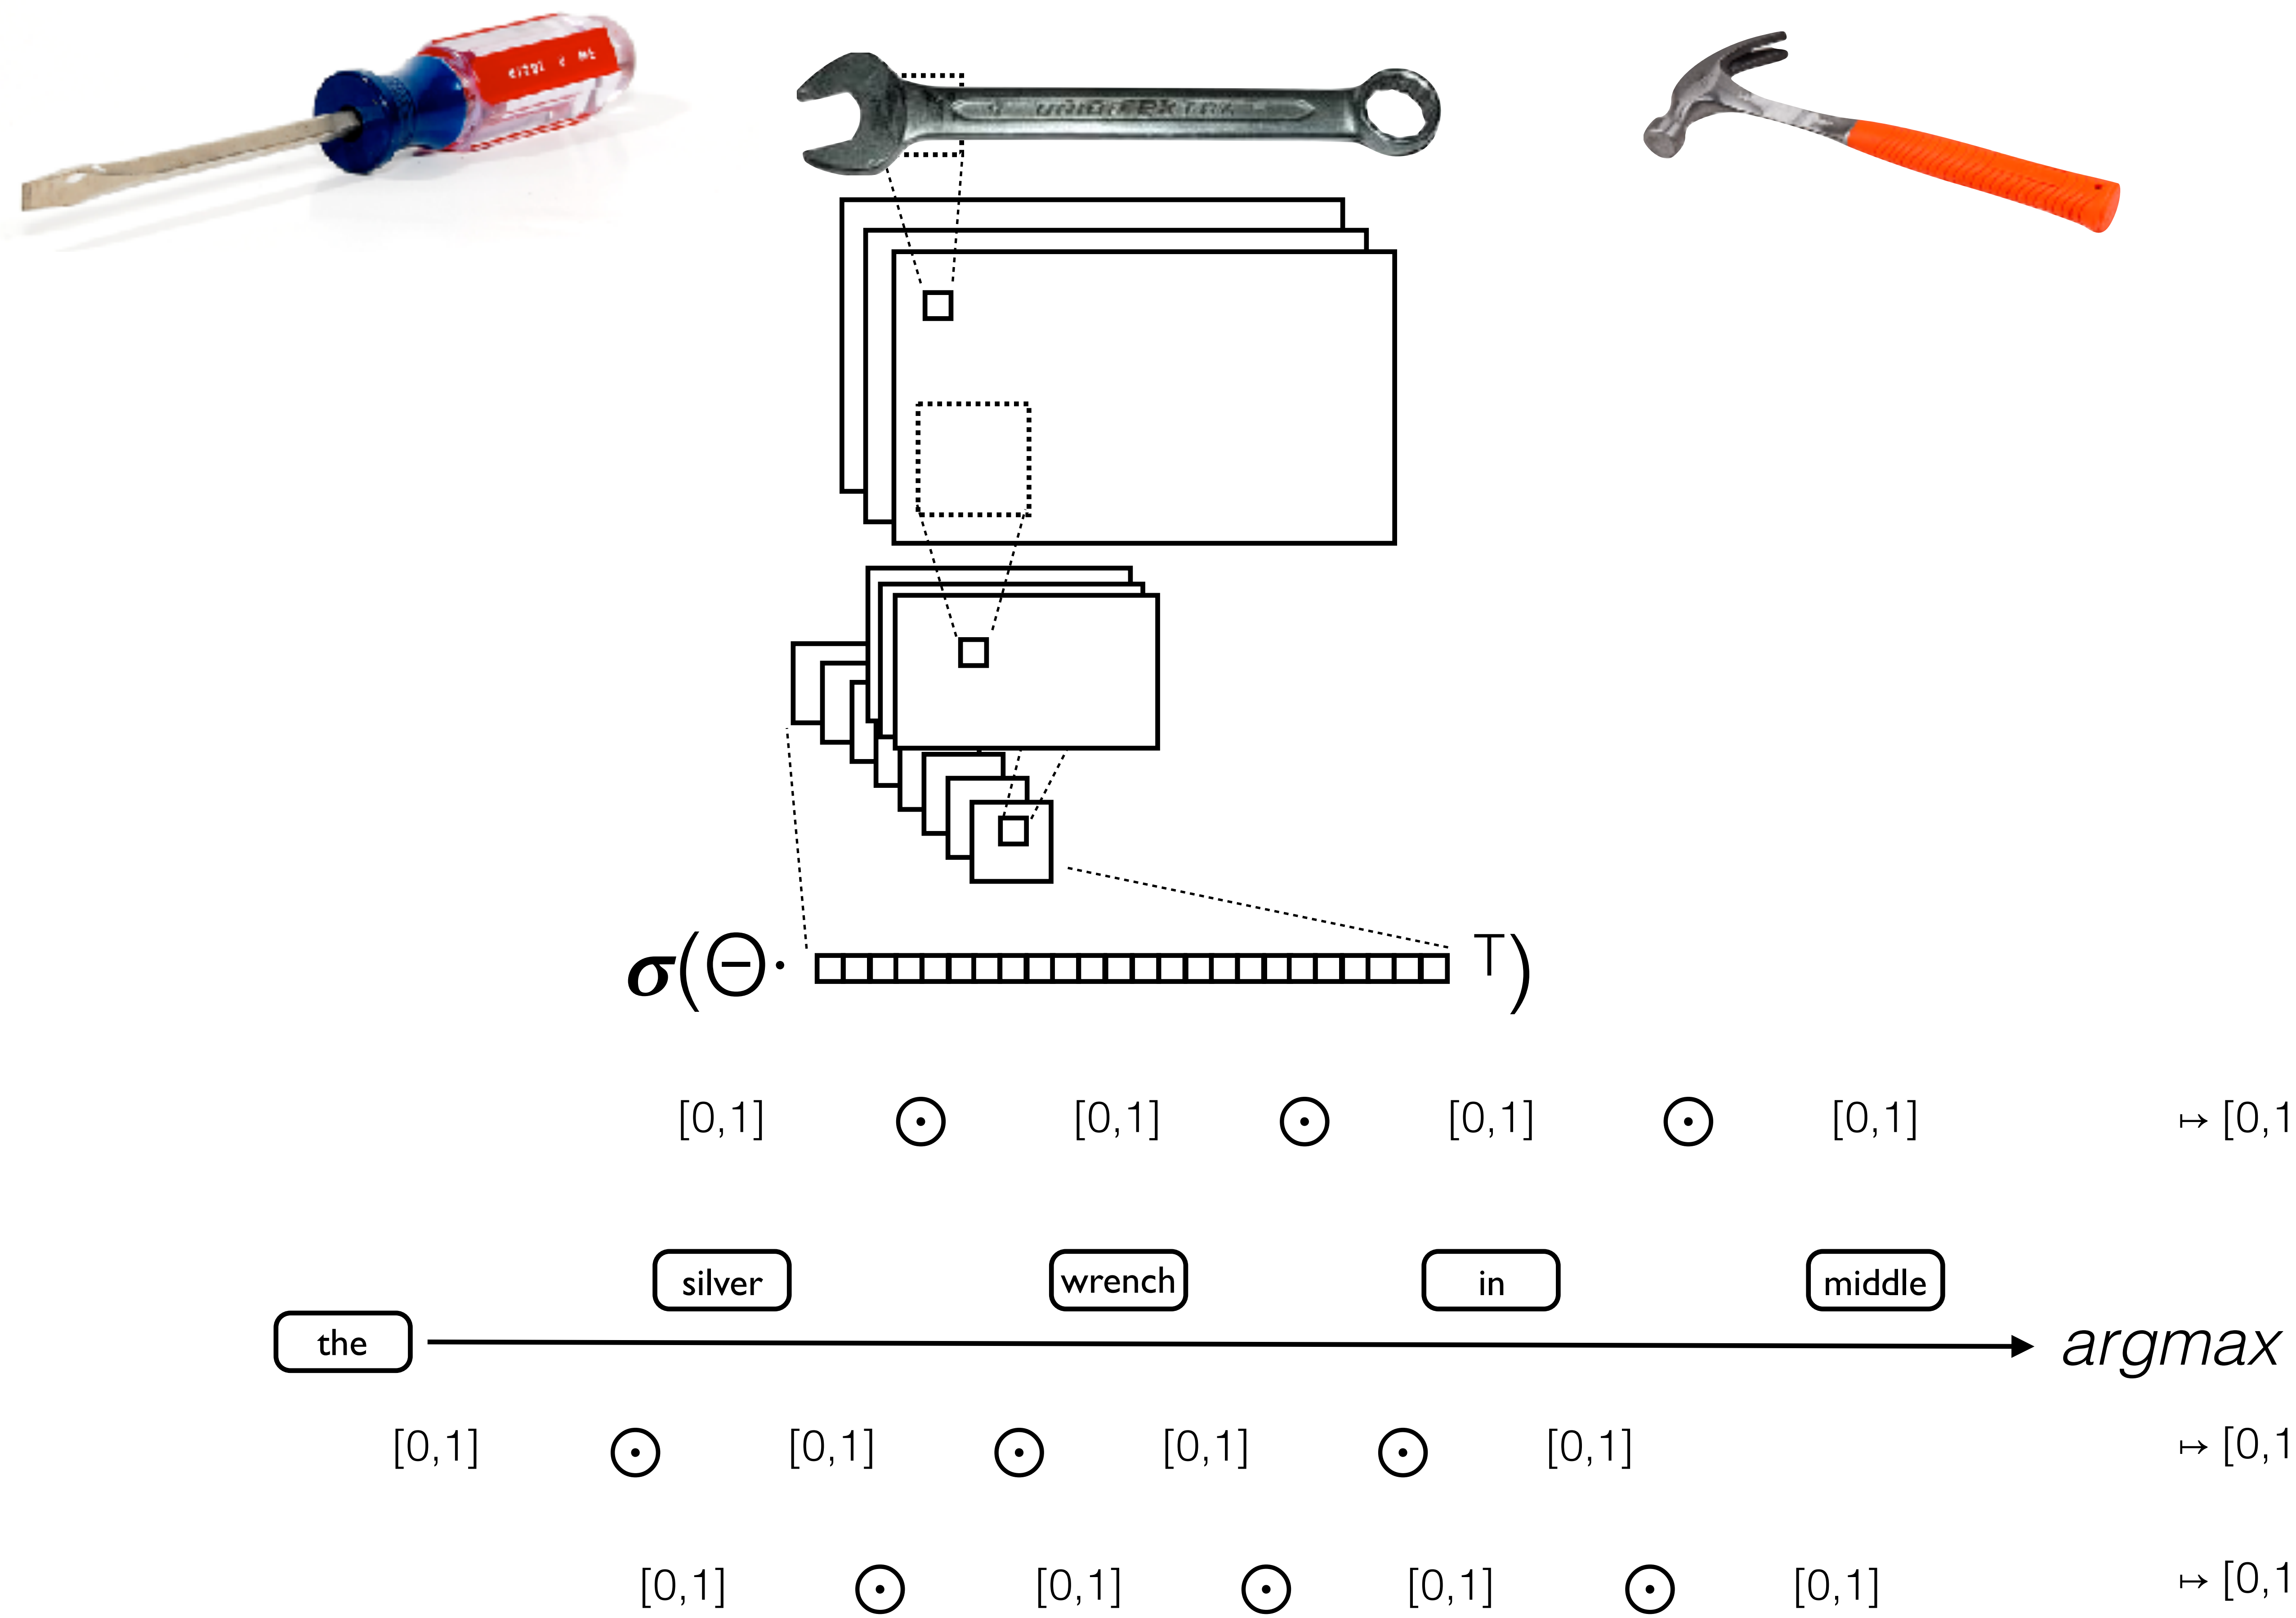
σ(Θ· T)
[0,1] ⊙ [0,1] ⊙ [0,1] ⊙ [0,1] ↦ [0,1]
the
silver
wrench
in
middle
argmax
[0,1] ⊙ [0,1] ⊙ [0,1] ⊙ [0,1] ↦ [0,1]
[0,1] ⊙ [0,1] ⊙ [0,1] ⊙ [0,1] ↦ [0,1]

Results

	%tst	acc	mrr	arc	>0	acc
REFERIT	1.00	0.65	0.79	0.89	0.97	0.67
REFERIT; NR	0.86	0.68	0.82	0.91	0.97	**0.71**
(Hu et al., 2015)	–	0.73	–	–	–	–
REFCOCO	1.00	0.61	0.77	0.91	0.98	0.62
REFCOCO; NR	0.94	0.63	0.78	0.92	0.98	**0.64**
(Mao et al., 2015)	–	0.70	–	–	–	–
GREXP	1.00	0.43	0.65	0.86	1.00	0.43
GREXP; NR	0.82	0.45	0.67	0.88	1.00	**0.45**
(Mao et al., 2015)	–	0.61	–	–	–	–

Results, full model

	RP@1	RP@10	rnd
REFERIT	0.09	0.24	0.03
REFERIT; NR	**0.10**	0.26	0.03
(Hu et al., 2015)	0.18	0.45	
REFCOCO	0.52	–	0.17
REFCOCO; NR	**0.54**	–	0.17
(Mao et al., 2015)	0.52		
GREXP	0.36	–	0.16
GREXP; NR	**0.37**	–	0.17
(Mao et al., 2015)	0.45		

Region Proposals

	nopos	pos	full	top20
RI	0.53	0.60	0.65	0.46
RI; NR	0.56	0.62	0.68	0.48
RC	0.44	0.55	0.61	0.52
RC; NR	0.45	0.57	0.63	0.53

Feature Ablation

(Schlangen, Zarrieß, Kennington; ACL 2016)

Overview

- Motivation: Knowledge from Testimony
- The Lexicon: Referential & Inferential Knowledge
- **Referential Knowledge: *Likeness***
  - Acquisition from Referential Interaction ✔
  - Application in Reference Resolution ✔
  - Application in Reference Generation
- **Inferential Knowledge**
  - … from Referential Knowledge / Referential Interaction
  - … from Definitions
- **Towards Justifying Concepts**

Generation

- straightforward: Applying all classifiers to object imposes ranking on vocab. Select from that.

Generation

- straightforward: Applying all classifiers to object imposes ranking on vocab. Select from that.

(Zarrieß & Schlangen; ACL 2016; INLG 2016 [best paper]; ACL 2017)

Overview

- Motivation: Knowledge from Testimony
- The Lexicon: Referential & Inferential Knowledge
- **Referential Knowledge: *Likeness***
  - Acquisition from Referential Interaction ✔
  - Application in Reference Resolution ✔
  - Application in Reference Generation ✔
- **Inferential Knowledge**
  - … from Referential Knowledge / Referential Interaction
  - … from Definitions
- **Towards Justifying Concepts**

Overview

- Motivation: Knowledge from Testimony
- The Lexicon: Referential & Inferential Knowledge
- **Referential Knowledge: *Likeness***
  - Acquisition from Referential Interaction
  - Application in Reference Resolution
  - Application in Reference Generation
- **Inferential Knowledge**
  - … from Referential Knowledge / Referential Interaction
  - … from Definitions
- **Towards Justifying Concepts**

Overview

- Motivation: Knowledge from Testimony
- The Lexicon: Referential & Inferential Knowledge
- **Referential Knowledge: *Likeness***
  - Acquisition from Referential Interaction
  - Application in Reference Resolution
  - Application in Reference Generation
- **Inferential Knowledge**
  - … from Referential Knowledge / Referential Interaction
  - … from Definitions
- **Towards Justifying Concepts**

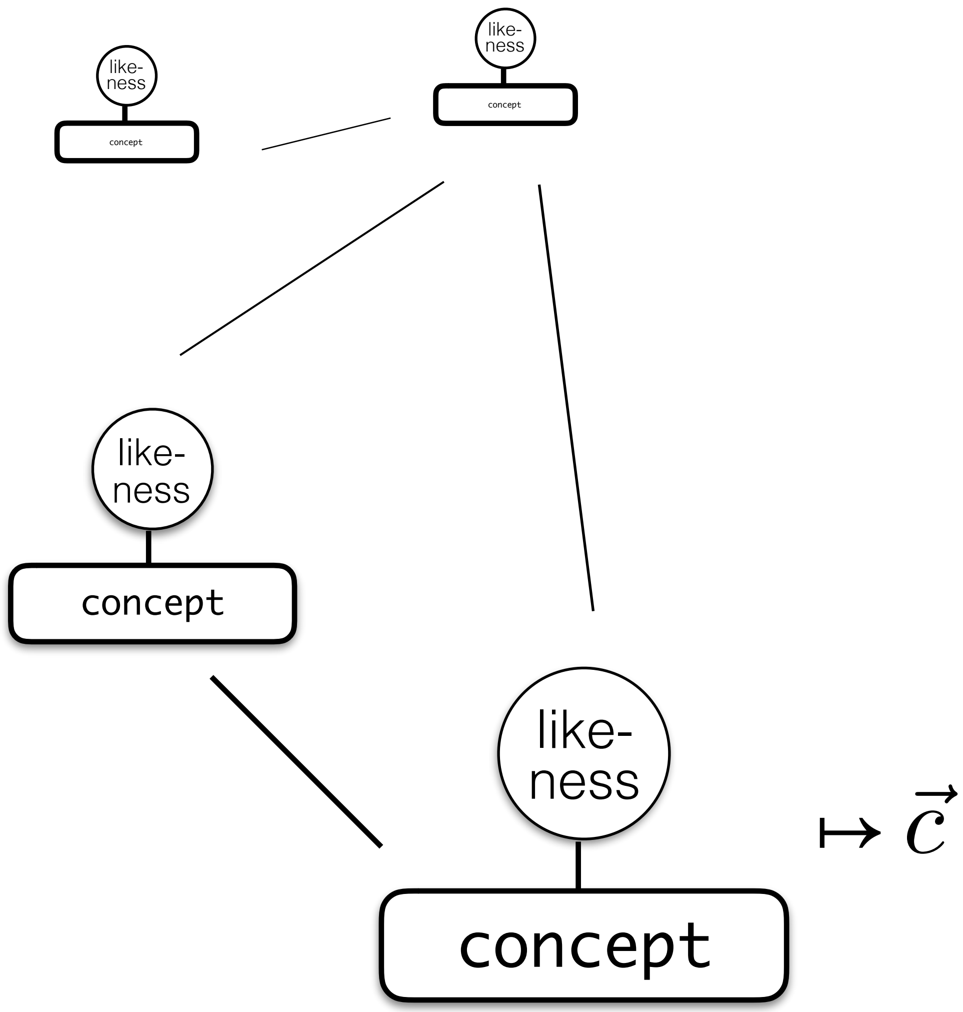

- *visual averages*: centroid of set of positive instances
- *weights / intensional:* weight vectors of classifier
- *representative responses / denotational:* vector of responses to randomly selected set of objects
- *full response signature:* vector of avg. responses of this classifier to positive instances of other categories

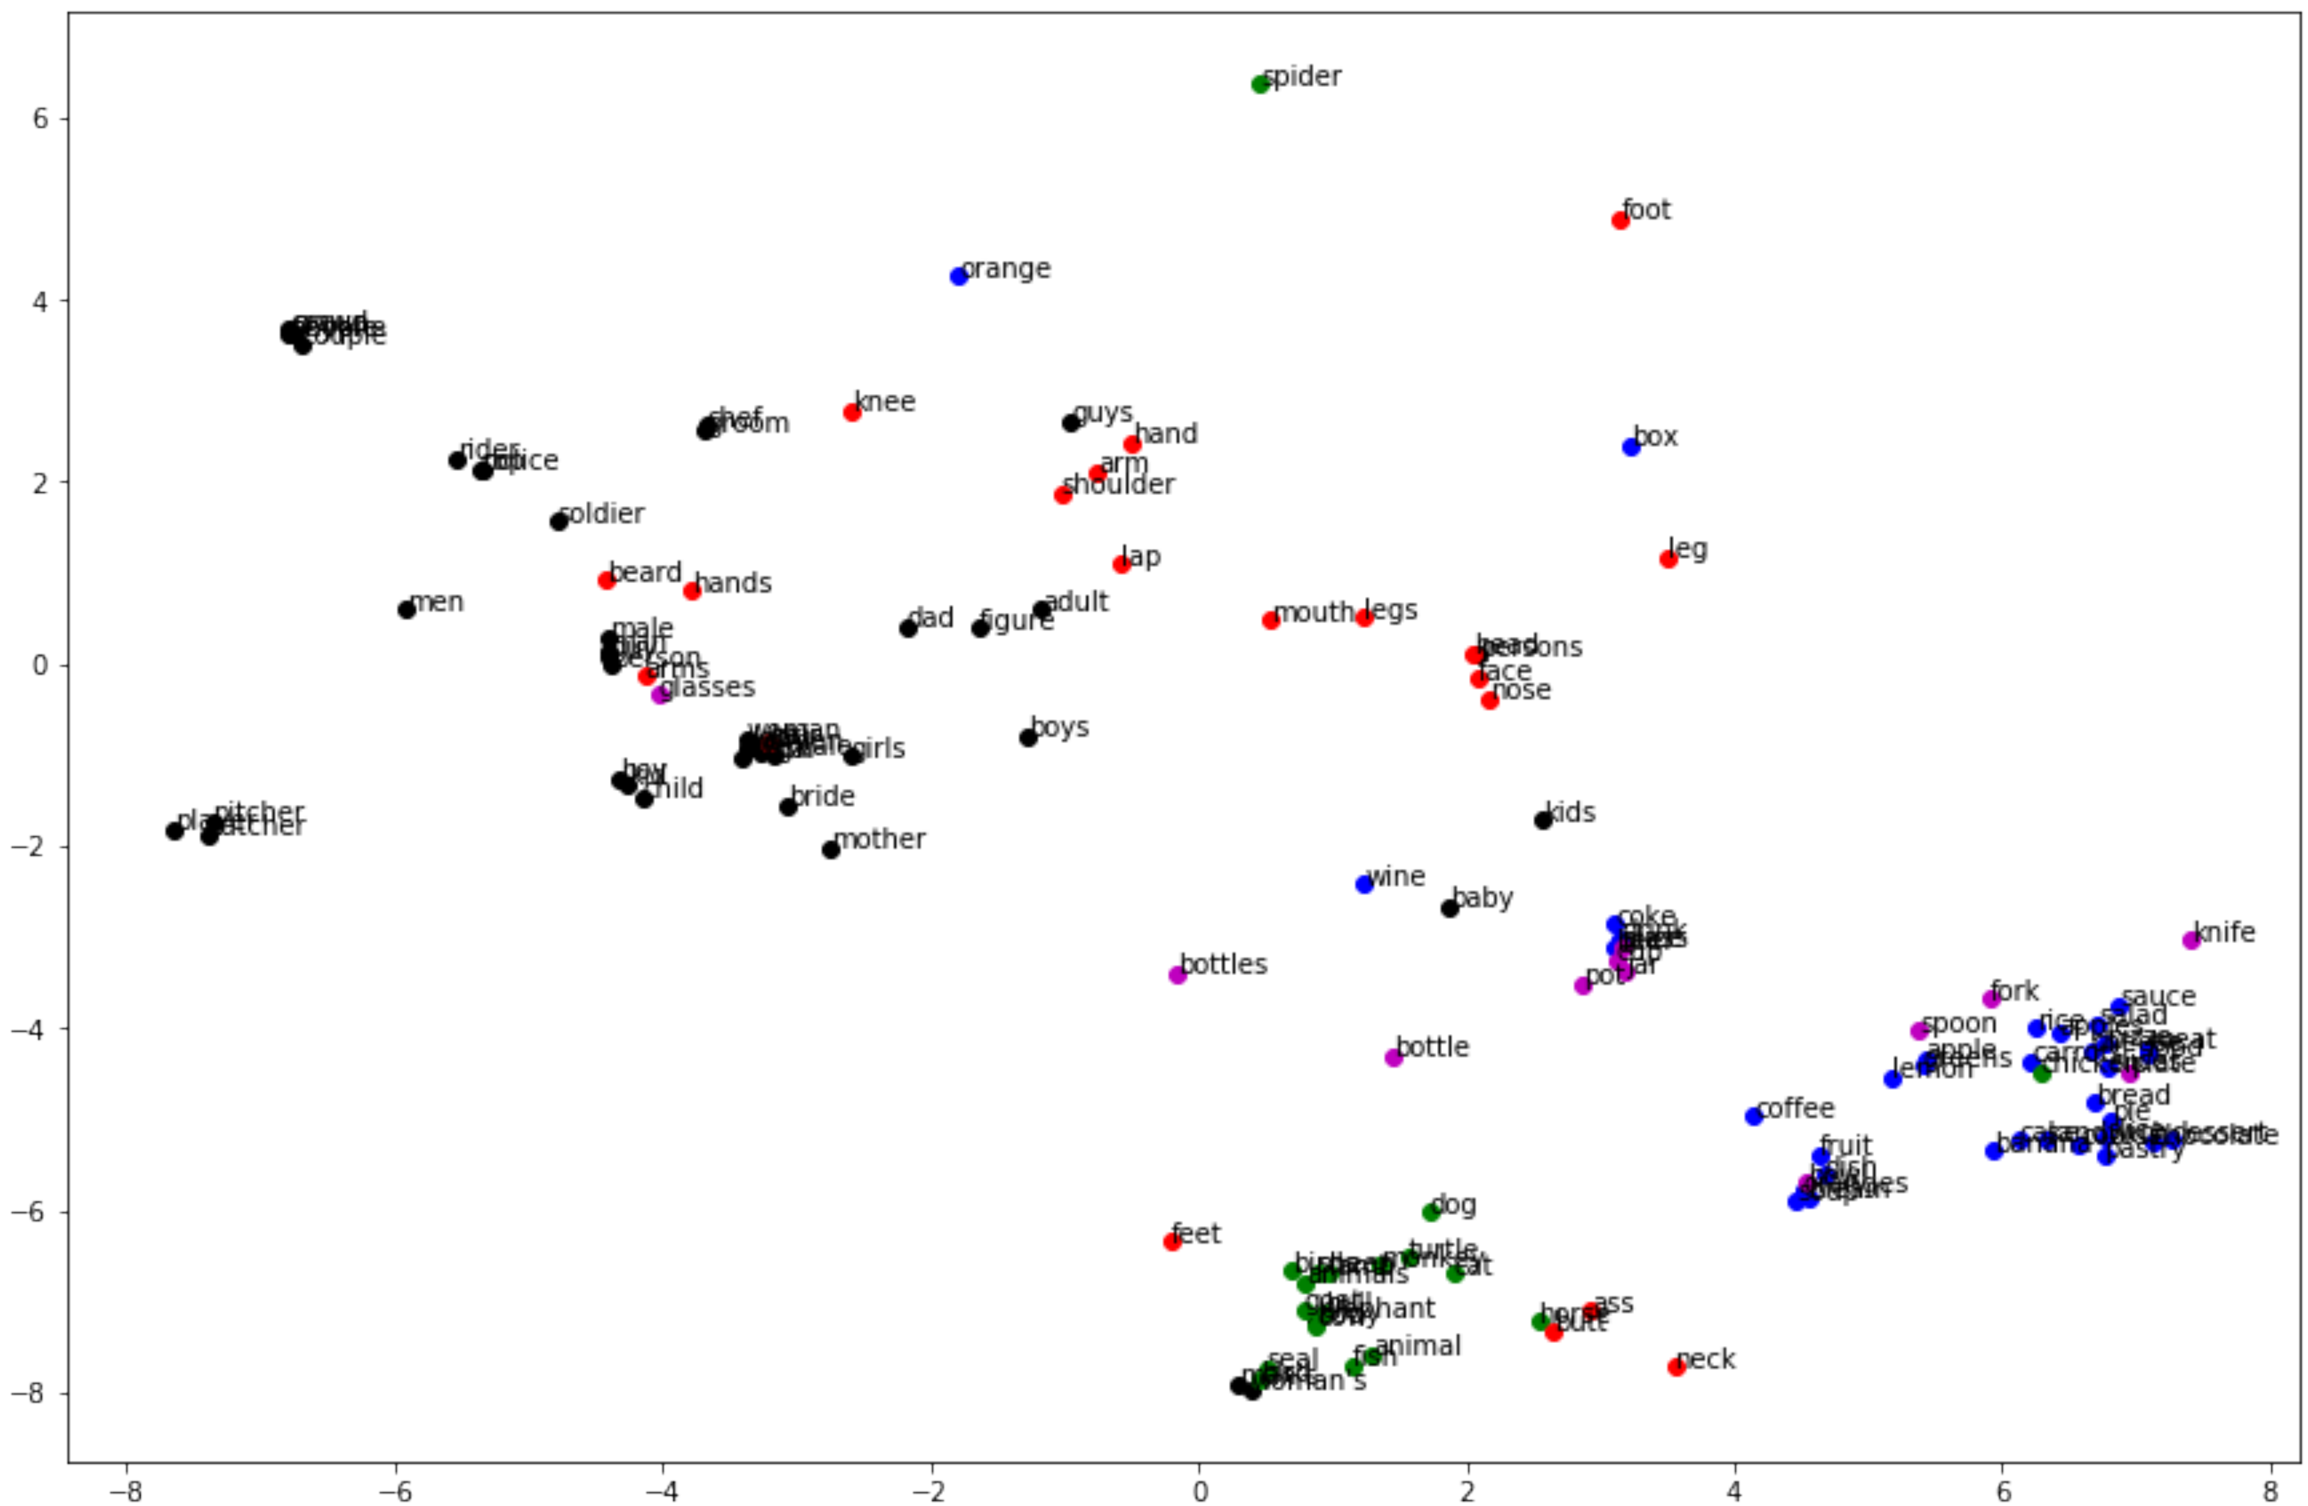

6
4
2
0
−2
−4
−6
−8
−8
−6
−4
−2
0
2
4
6
8
spider
foot
orange
box
couple
knee
guys
hand
arm
shoulder
chef
groom
rider
police
soldier
lap
leg
beard
hands
men
dad
figure
adult
mouth
legs
persons
face
nose
male
person
arms
glasses
boys
girls
boy
child
bride
mother
pitcher
catcher
kids
wine
baby
coke
bottles
pot
jar
knife
fork
sauce
spoon
salad
rice
apple
lemon
meat
chicken
coffee
bread
pie
pastry
chocolate
fruit
dish
dog
feet
turtle
animals
cat
elephant
ass
butt
neck
seal
fish
animal
woman's

mom	ext	adult female girl woman looking
	vis	child girl female blond eating
	sit	child little kid leg striped
	w2v	kids kid baby girl boy
wall	ext	corner edge brick side stone
	vis	brick wood picture open part
	sit	curtain lamp bed desk ceiling
	w2v	fence brick ceiling roof stone
road	ext	pavement sidewalk ground dirt path
	vis	street pavement sidewalk ground where
	sit	van sidewalk grass rider car
	w2v	street bridge pavement hill bus
statue	ext	tower thing far short pillar
	vis	tall tallest tower fountain background
	sit	palm stairs tree bush bushes
	w2v	fountain pillar tower waterfall kneeling
bottom	ext	lower ground corner front closest
	vis	lower corner pic click are area
	sit	donut shelf doughnut cupcake top
	w2v	top leftmost side rightmost end
chef	ext	groom shirt hoodie man kid
	vis	cutting old bald groom man
	sit	female vest cutting bald object
	w2v	dish pizza hotdog food skier
chair	ext	seat table sitting couch wooden
	vis	seat bench wooden empty desk
	sit	table sofa empty pool wooden
	w2v	head sofa seat board couch

Evaluating Derived Concept Relations

Hypernymy

- Linked 589 terms from vocabulary (s+r+rp) to WordNet synset
- Identified 516 pairs of (term A, term B), where B is in closure of hyponym relation of A
- Rule: if A “likes” real superset of what B “likes”, A is hypern. of B.
  0.18 f-score (on denotational vectors)
- Entropy (Kiela *et al.* 2015): if A & B related, and entropy(A) > entropy(B), then hyper(A, B)
    - visual averages: 0.21 f-score
    - denotational vectors: 0.15 f-score
- False positives: “scarf” is a type of “woman”, “shirt” is a type of “man”, etc.
  false false positives: “cowboy” is a type of “dude”…

Evaluating Derived Concept Relations

Similarity / Relatedness / Compatibility

Model	MEN	SemSim	VisSim	Compatibility
w2v_ref	0.669	**0.687**	**0.580**	**0.251**
w2v_den	**0.765**	0.651	0.570	0.164
w2v_sit	0.586	0.515	0.409	0.166
baronimod	0.785	0.704	0.594	0.241
vis_av	0.523	0.526	0.486	0.287
wac_int	-0.373	-0.339	-0.294	-0.076
wac_den	-0.593	-0.615	-0.536	**-0.288**
wac_resp	**0.634**	**0.656**	**0.574**	0.276

(Baroni *et al.* 2014) CBOW, 400dim

(Bruni *et al.* 2012) 372 out of 3,000

(Silberer & Lapata 2014) 721 out of 7,577

(Kruszewski & Baroni 2015) 1,859 out of 17,973

Predicting Incompatible Modifiers

```
[('left man', 190),
 ('right man', 159),
 ('man right', 153),
 ('the man', 129),
 ('man left', 111),
 ('man standing', 100),
 ('man sitting', 74),
 ('old man', 64),
 ('bald man', 62),
 ('closest man', 61),
 ('middle man', 52),
 ('black man', 45),
 ('standing man', 39),
 ('older man', 33),
 ('tallest man', 28),
 ('man eating', 26),
 ('taller man', 24),
 ('tall man', 22),
 ('blue man', 21),
 ('man glasses', 18)]
```

man

left ⟷ right

young ⟷ old

old ⟷ shirtless

shirt

plaid ⟷ green

red ⟷ gray

blue ⟷ yellow

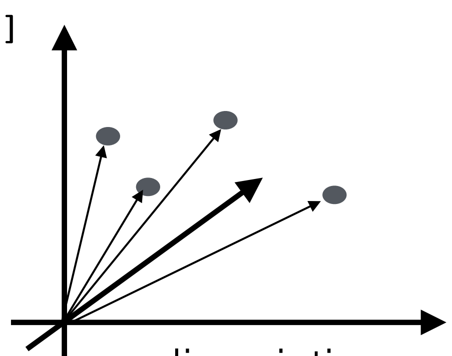

linguistic context

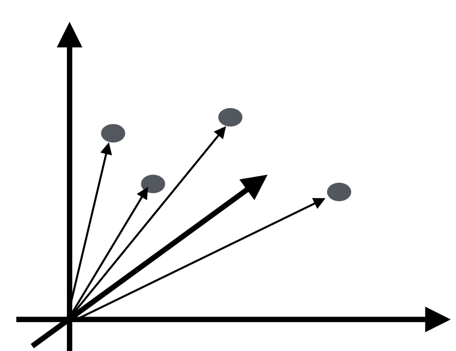

referential context

elephant

closest ⟷ back

big ⟷ baby

adult ⟷ smaller

Evaluating Derived Concept Relations

Similarity / Relatedness / Compatibility

Model	MEN	SemSim	VisSim	Compatibility
w2v_ref	0.669	**0.687**	**0.580**	**0.251**
w2v_den	**0.765**	0.651	0.570	0.164
w2v_sit	0.586	0.515	0.409	0.166
baronimod	0.785	0.704	0.594	0.241
vis_av	0.523	0.526	0.486	0.287
wac_int	-0.373	-0.339	-0.294	-0.076
wac_den	-0.593	-0.615	-0.536	**-0.288**
wac_resp	**0.634**	**0.656**	**0.574**	0.276

(Baroni *et al.* 2014) CBOW, 400dim

(Bruni *et al.* 2012) 372 out of 3,000

(Silberer & Lapata 2014) 721 out of 7,577

(Kruszewski & Baroni 2015) 1,859 out of 17,973

Overview

- Motivation: Knowledge from Testimony
- The Lexicon: Referential & Inferential Knowledge
- **Referential Knowledge: *Likeness***
  - Acquisition from Referential Interaction
  - Application in Reference Resolution
  - Application in Reference Generation
- **Inferential Knowledge**
  - … from Referential Knowledge / Referential Interaction ✔
  - … from Definitions
- **Towards Justifying Concepts**

Overview

- Motivation: Knowledge from Testimony
- The Lexicon: Referential & Inferential Knowledge
- **Referential Knowledge: *Likeness***
  - Acquisition from Referential Interaction
  - Application in Reference Resolution
  - Application in Reference Generation
- **Inferential Knowledge**
  - … from Referential Knowledge / Referential Interaction ✔
  - … from Definitions
- **Towards Justifying Concepts**

Overview

- Motivation: Knowledge from Testimony
- The Lexicon: Referential & Inferential Knowledge
- **Referential Knowledge: *Likeness***
  - Acquisition from Referential Interaction
  - Application in Reference Resolution
  - Application in Reference Generation
- **Inferential Knowledge**
  - … from Referential Knowledge / Referential Interaction
  - … from Definitions
- **Towards Justifying Concepts**

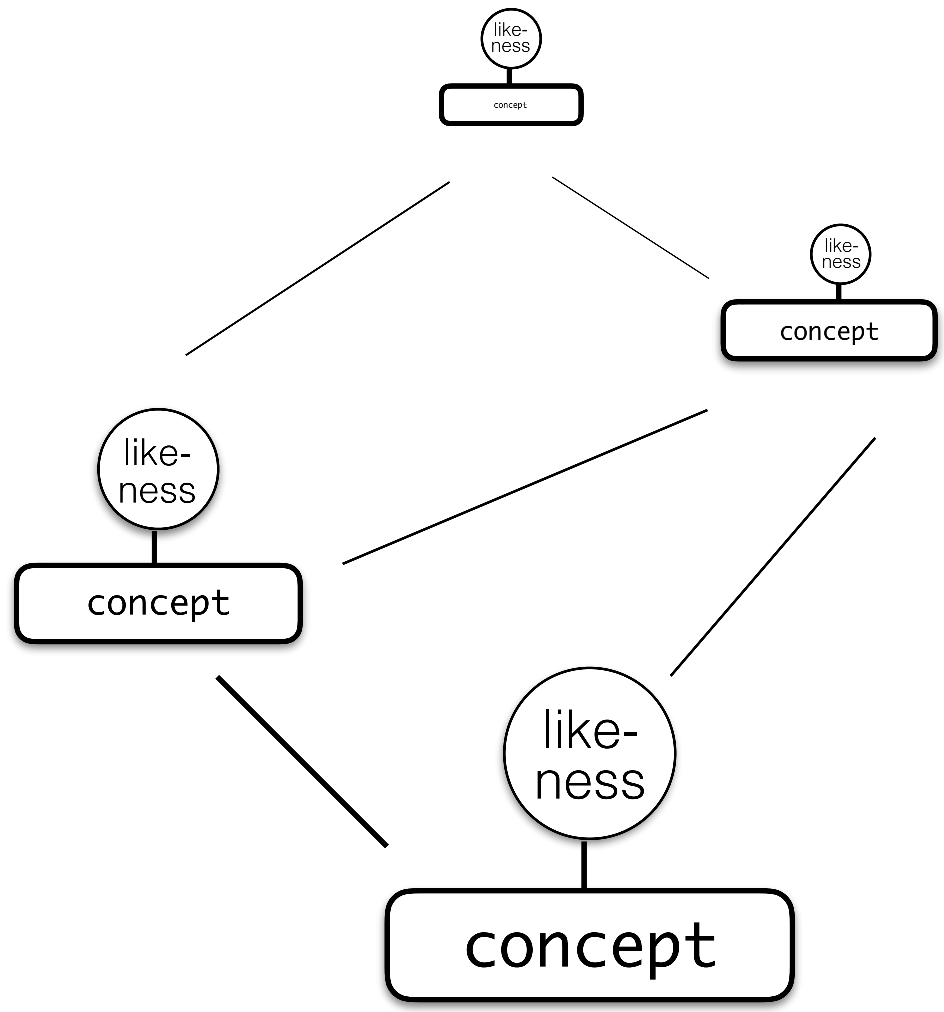

Learning from explicit definition

Recipe:

- Definition links definiendum to other concepts
- If those have likeness representation, do direct attribute prediction (Lampert *et al.* 2009)

$$p(z|x) \propto \prod_{m=1}^{M} \left(\frac{p(a_m|x)}{p(a_m)} \right)^{a_m^z}$$

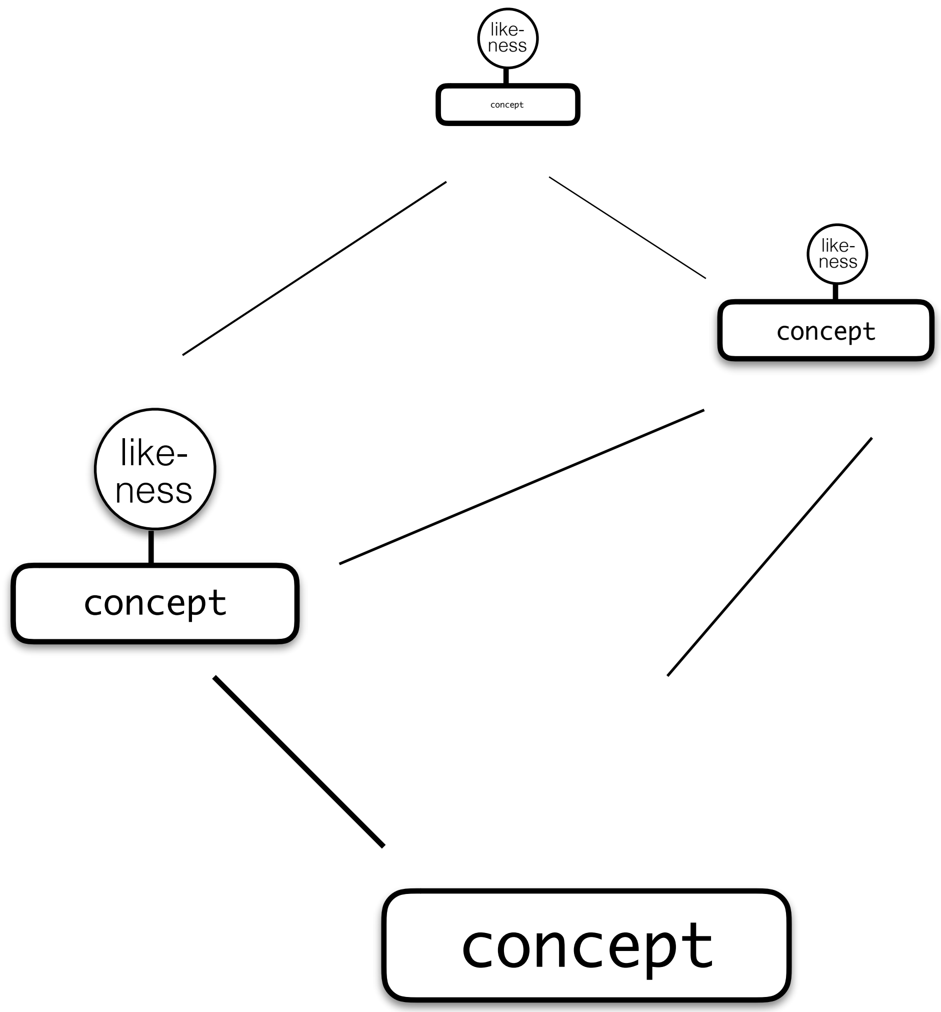

Learning from explicit definition

Recipe:

- Definition links definiendum to other concepts
- If those have likeness representation, do direct attribute prediction (Lampert *et al.* 2009)
- E.g., replace “wampimuk” with “small mouse mammal”

Zero-Shot Learning with Feature Norms

behavior	eats, walks, climbs, swims, runs
diet	drinks_water, eats_anything
shape_size	is_tall, is_large
anatomy	has_mouth, has_head, has_nose, has_tail, has_claws, has_jaws, has_neck, has_snout, has_feet, has_tongue
color_patterns	is_black, is_brown, is_white

botany	has_skin, has_seeds, has_stem, has_leaves, has_pulp
color_patterns	purple, white, green, has_green_top
shape_size	is_oval, is_long
texture_material	is_shiny

behavior	rolls
parts	has_step_through_frame, has_fork, has_2_wheels, has_chain, has_pedals has_gears, has_handlebar, has_bell, has_breaks has_seat, has_spokes
texture_material	made_of_metal
color_patterns	different_colors, is_black, is_red, is_grey, is_silver

(Silberer, Ferrari & Lapata, 2013), using feature norms of (McRae *et al.* 2005)

114 out of 509 concepts in vocab
instances for 340 of 637 attributes

Acc. on 20 test classes:
43.2%

D
E
F
G
H
I
J
K
L
M
N
O
P
Q
R
S
T
U
V
W
X
Y
Z

docks *noun*
1 a place where ships load and unload cargo.
dock *verb*

2 the place in a courtroom where the person on trial stands or sits.

doctor
doctors *noun*
a person who is trained to treat sick or injured people.

dodge
dodges dodging dodged *verb*
to avoid being hit by something by moving out of the way very quickly.
She ***dodged*** *the ball coming toward her.*

dolphin
dolphins *noun*
a fish-eating sea mammal. Dolphins breathe air, so they must swim to the surface often. They are friendly animals and are known for their intelligence. Dolphins are a type of small whale.

■ say **doll**-fin

dog
dogs *noun*
a mammal that is often kept as a pet. Dogs mainly eat meat and can be trained to perform certain tasks, such as herding sheep. Dogs are related to wolves and foxes (see **pet** on page 148).

collie dog

donation
donations *noun*
a gift, usually of money, that is made to a charity or another organization.
He made a large ***donation****.*

donkey
donkeys *noun*
a member of the horse family that has long ears and a soft, furry coat. Donkeys eat grass and in some countries are used for carrying people and goods.

door
doors *noun*
a piece of wood, glass, or metal that opens and shuts to provide a way into a room, cupboard, building, or vehicle.

dot
dots *noun*
a very small, round spot.
Ladybugs have ***dots*** *on them.*

double
adjective
twice as much.

A ***double*** *six.*
■ say **dub**-ul

doubtful
adjective
not sure, or unlikely.
He was ***doubtful*** *about his chances of winning.*
■ say **dout**-ful
doubt *verb*

dough
noun
a mixture of flour and either milk or water that is used to make bread or cakes.
■ say **doh**

doughnut
doughnuts *noun*
a sweet, round cake made from dough, which is fried in fat and covered in sugar.

■ say **doh**-nut

dove
doves *noun*
a bird that is a member of the pigeon family. Doves are often used as a symbol of peace.

Overview

- Motivation: Knowledge from Testimony
- The Lexicon: Referential & Inferential Knowledge
- **Referential Knowledge: *Likeness***
    - Acquisition from Referential Interaction
    - Application in Reference Resolution
    - Application in Reference Generation
- **Inferential Knowledge**
    - … from Referential Knowledge / Referential Interaction
    - … from Definitions ✓
- **Towards Justifying Concepts**

Overview

- Motivation: Knowledge from Testimony
- The Lexicon: Referential & Inferential Knowledge
- **Referential Knowledge: *Likeness***
  - Acquisition from Referential Interaction
  - Application in Reference Resolution
  - Application in Reference Generation
- **Inferential Knowledge**
  - … from Referential Knowledge / Referential Interaction
  - … from Definitions
- **Towards Justifying Concepts**

Justifying Concepts

- Show that you are a competent language user by
  - Applying terms (mostly) correctly
  - Being able to defend your decisions / being able to make your concepts explicit
  - Learning from being corrected

User: Draw a new card

System: *Ok.*

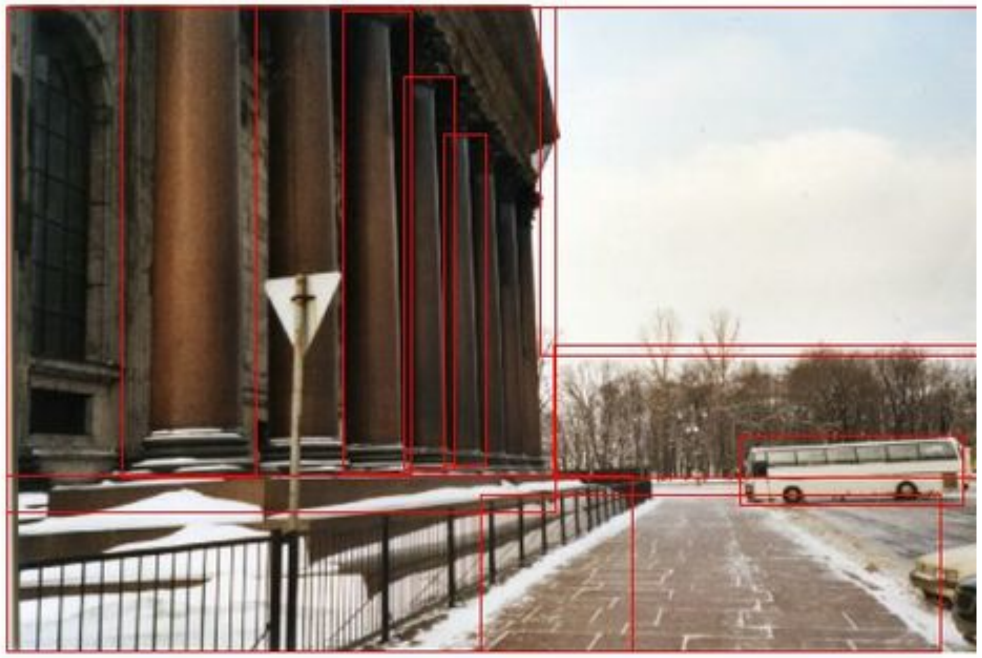

User: get a new collage men

System: *Ok.*

U: Find the man with the book

S: *Let me see. Did you mean this?*

U: yes

S: *Thanks!*

U: explain

S: *Okidoki.*

S: *0 is best for "man"*

S: *5 is best for "book"*

S: *Overal best: 5*

S: *Rank of region 5 for "man": 4*

S: *Rank of region 5 for "book": 1*

Justifying Concepts

- U: *Find the Rottweiler.*
- S: Picture 4.
- U: *Explain*.
- S: I have never seen a Rottweiler, but
  I know that it is a type of dog.
  4 is the only dog.

Justifying Concepts

- Show that you are a competent language user by
  - Applying terms (mostly) correctly
  - Being able to defend your decisions / **being able to make your concepts explicit**
  - Learning from being corrected
- Do people actually appreciate this? Does it work?

Learning

Application

correction:
“No, *this* is an X”
“No, this is similar, but Xs don’t have A”

demonstration:
“*This* is an X.”

expl. definition:
“An X is …”

impl. definition:
“bla bla bla X bla bla”

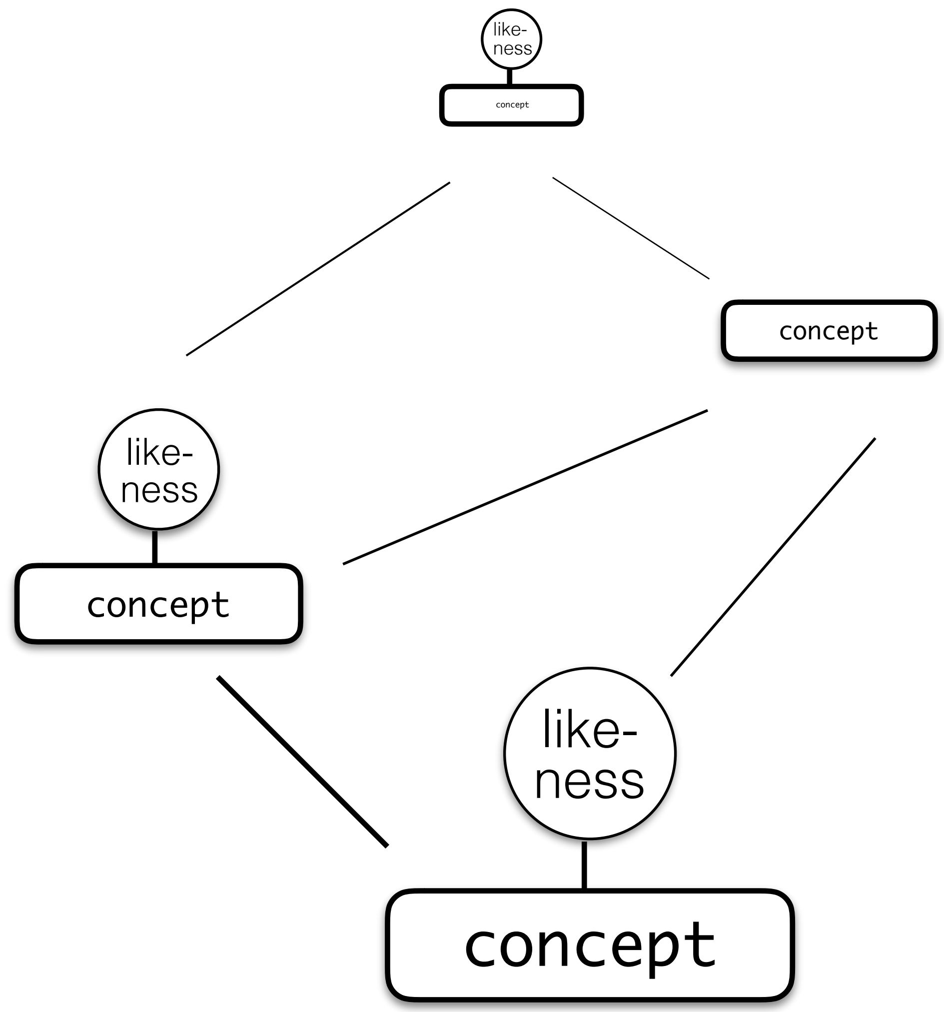

selection:
“*This* is an X.”

justification:
“I think it’s this, because an X is a type of Y, and this is a Y”
“Xs are …”

Overview

- Motivation: Knowledge from Testimony
- The Lexicon: Referential & Inferential Knowledge
- **Referential Knowledge: *Likeness***
    - Acquisition from Referential Interaction
    - Application in Reference Resolution
    - Application in Reference Generation
- **Inferential Knowledge**
    - … from Referential Knowledge / Referential Interaction
    - … from Definitions
- **Towards Justifying Concepts**

Loose Ends

- Integrate this into probabilistic logic.
- Use inferential knowledge to drive actual inferences…
- Discourse representations.
- Learn syntax / composition from referential interaction.

Current / Future Work

- Assembling a better tutor by structuring the training data (Z&S, EACL 2017, ACL 2017, forth)
- Improving generation with situational constraints

Dialogue Systems Group - Bielefeld University

- *Post-Docs*
  - Julian Hough (PhD QMU London)
  - **Sina Zarrieß** (Phd Stuttgart)
  - Iwan de Kok (PhD U Twente)



- *PhD Students*
  - Ting Han
  - Soledad López
  - Birte Carlmeyer *
  - Simon Betz *

(* co-supervised)




- *Alumni*
  - **Casey Kennington** (PhD; now Boise State Univ)
  - Spyros Kousidis (Post-Doc; now Carmeq GmbH)
  - Timo Baumann (PhD; now Univ. Hamburg)
  - Gabriel Skantze (Post-Doc; now KTH, Stockhlm)
  - Okko Buß (PhD; now Carmeq GmbH)
  - Michaela Atterer (Post-Doc)




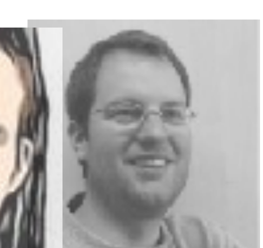
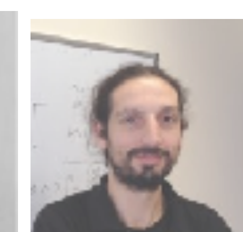

Thank you!