Bayesian Classification and Inference in a Probabilistic Type Theory with Records

Staffan Larsson and Robin Cooper
Centre for Linguistic Theory and Studies in Probability (CLASP)
Dept. of Philosophy, Linguistics and Theory of Science
University of Gothenburg

NALOMA Workshop, June 16, 2021

Outline

Outline

Introduction

- Questions
  - How is linguistic meaning related to perception?
  - How do we learn and agree on the meanings of our words?
- We are developing a formal *judgement-based semantics* where notions such as perception, classification, judgement, learning and dialogue coordination play a central role
  - See e.g. Cooper (2005), Cooper and Larsson (2009), Larsson (2011), Dobnik *et al.* (2011), Cooper (2012a), Dobnik and Cooper (2013), Cooper *et al.* (2015b)
- Key ideas:
  - modelling perceptual meanings as classifiers of real-valued perceptual data
  - modelling how agents learn and coordinate on meanings through interaction with other agents (semantic coordination)

Classification

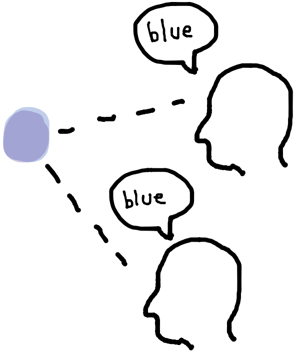

Classification is subjective?

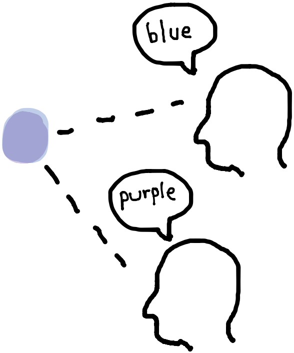

Coordination process

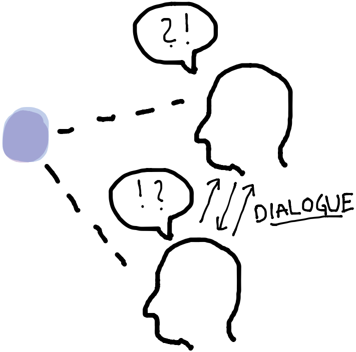

Classification is coordinated

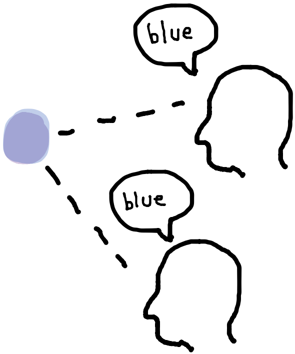

Classification is coordinated

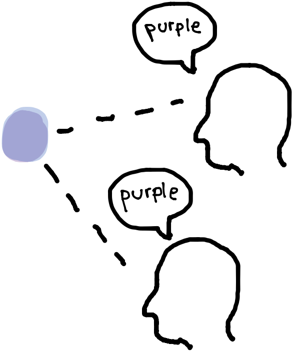

Coordination can be creative

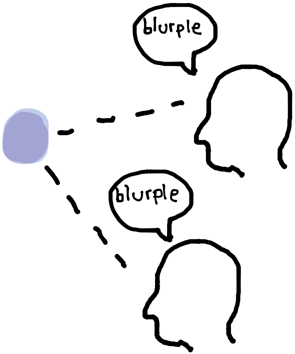

Type Theory with Records

- We want to use a framework which also encompasses accounts of many problems traditionally studied in formal semantics[1].
- We will be using Type Theory with Records, or TTR
  - see Cooper (2012a),Cooper (2012b),Cooper and Ginzburg (2015) and Cooper (in prep)
- TTR starts from the idea that information and meaning is founded on our ability to perceive and classify the world.
- Based on the notion of *judgements* of entities and situations being of certain *types*.

[1]Semantic phenomena which have been described using TTR include modelling of intensionality and mental attitudes (Cooper, 2005), dynamic generalised quantifiers (Cooper, 2004), co-predication and dot types in lexical innovation, frame semantics for temporal reasoning, reasoning in hypothetical contexts (Cooper, 2011), enthymematic reasoning (Breitholtz and Cooper, 2011), clarification requests (Cooper, 2010), negation (Cooper and Ginzburg, 2011), and information states in dialogue (Cooper, 1998; Ginzburg, 2012).

Probabilistic TTR

- A probabilistic type theory with records was introduced in Cooper *et al.* (2014) and Cooper *et al.* (2015a)
- In probabilistic TTR (probTTR) we associate probabilities with judgements: $p(a : T)$ ("the probability that a is of type T").

Why probabilistic TTR?

A single framework for modelling

- *vagueness* and *gradience* of semantic judgements (Fernández and Larsson, 2014).
- probabilistic *reasoning*, including logical and enthymematic inference (Breitholtz, 2020).
- *grounding language in perception* and the real world (Larsson, 2015; Larsson, 2020)
- semantic and factual *learning* and *coordination*
- *interaction in dialogue*, including interactive learning and reasoning

Why ProbTTR in NALOMA?

- Probabilistic TTR provides a hybrid framework for natural language semantics, in the sense that it combines
  - Explainable probabilistic classification, inference and learning
  - "Black box" (e.g. neural net) classification, inference and learning.
- This enables us to model, for example, learning both from perceptual experience and from linguistic interaction.
- Although we focus here on semantic classification, probabilistic TTR is applicable to inference in general.

Outline

Outline

TTR fundamentals I

- $a : T$ is a judgment that a is of type T
- $: T$ is a judgement that there is something of type T
  - T is non-empty; often written Ttrue in Martin-Löf type theory

TTR fundamentals II

- Types may be either *basic* or *complex*
- Some basic types in TTR:
  - *Ind*, the type of an individual
  - *Real*, the type of real numbers
  - *[0,1]*, the type of real numbers between 0 and 1 (such as probabilities)

TTR fundamentals III

- Complex types are structured objects which have types or other objects introduced in the theory as components
- *ptypes* are constructed from a predicate and arguments of appropriate types as specified for the predicate.
- Examples are 'man(a)', 'see(a,b)' where $a, b : \mathit{Ind}$.
- The objects or *witnesses* of ptypes can be thought of as proofs in the form of situations, states or events in the world which instantiate the type.

Records and record types

- If
  - $a_1 : T_1$,
  - $a_2 : T_2(a_1)$,
  - $\ldots$,
  - $a_n : T_n(a_1, a_2, \ldots, a_{n-1})$,
  - where $T(a_1, \ldots, a_n)$ represents a type T which depends on the objects $a_1, \ldots, a_n$,
- ...the record to the left is of the record type to the right.

$$\left[\begin{array}{lll} \ell_1 & = & a_1 \\ \ell_2 & = & a_2 \\ \ldots & & \\ \ell_n & = & a_n \\ \ldots & & \end{array}\right] : \left[\begin{array}{lll} \ell_1 & : & T_1 \\ \ell_2 & : & T_2(l_1) \\ \ldots & & \\ \ell_n & : & T_n(\ell_1, l_2, \ldots, l_{n-1}) \end{array}\right]$$

- $\ell_1, \ldots \ell_n$ are *labels* which can be used elsewhere to refer to the values associated with them.

Records and record types

- A sample record and record type:

$$\begin{bmatrix} \text{ref} & = & \text{obj}_{123} \\ \text{c}_{\text{man}} & = & \text{prf1} \\ \text{c}_{\text{run}} & = & \text{prf2} \end{bmatrix} : \begin{bmatrix} \text{ref} & : & \textit{Ind} \\ \text{c}_{\text{man}} & : & \text{man(ref)} \\ \text{c}_{\text{run}} & : & \text{run(ref)} \end{bmatrix}$$

- The record on the left is of the record type on the right provided
  - $\text{obj}_{123} : \textit{Ind}$
  - $\text{prf1} : \text{man}(\text{obj}_{123})$
  - $\text{prf2} : \text{run}(\text{obj}_{123})$
- We will introduce further details of TTR as we need them in subsequent sections.

Outline

Probabilistic TTR fundamentals I

- The core of ProbTTR is the notion of probabilistic judgement.
- There are two kinds of judgement corresponding to the two kinds of judgement in non-probabilistic TTR:
  1. $p(s : T)$ – the probability that a situation, s, is of type, T
  2. $p(T)$ – the probability that there is some witness of type T.
- This introduces a distinction that is not normally made explicit in the notation used in probability theory.

Probabilistic TTR fundamentals II

- A *probabilistic Austinian proposition* is an object (a record) that corresponds to, or encodes, a probabilistic judgement.
- Probabilistic Austinian propositions are records of the type

$$\begin{bmatrix} \text{sit} & : & \textit{Sit} \\ \text{sit-type} & : & \textit{Type} \\ \text{prob} & : & [0,1] \end{bmatrix}$$

  (where $[0, 1]$ represents the type of real numbers between 0 and 1).
- An object φ of the above type corresponds to, or encodes, the judgement p(φ.sit:φ.sit-type)$=$ φ.prob

Outline

Bayesian inference I

- Bayesian Networks provide graphical models for probabilistic learning and inference (Pearl, 1990, Halpern, 2003).
- A Bayesian Network is a Directed Acyclic Graph (DAG).
- The nodes of the DAG are random variables
- Its directed edges express dependency relations among the variables.
- The graph describes a complete joint probability distribution (JPD) for its random variables.

Bayesian inference II

P(C=F)	P(C=T)
0.5	0.5

C	P(S=F)	P(S=T)
F	0.5	0.5
T	0.9	0.1

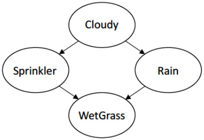

C	P(R=F)	P(R=T)
F	0.8	0.2
T	0.2	0.8

S	R	P(W=F)	P(W=T)
F	F	1.0	0.0
T	F	0.1	0.9
F	T	0.1	0.9
T	T	0.01	0.99

- Russell and Norvig (1995) give the example Bayesian Network above.
- The only directly observable evidence is whether it is cloudy or not, and the queried variable is whether the grass is wet or not.
- Whether it is raining and whether the sprinkler is on is not known, but both of these factors depend on whether it is cloudy, and both affect whether the grass is wet.

Bayesian inference III

- From this Bayesian Network we can compute the marginal probability of the grass being wet ($W = \mathrm{T}$):

$$p(W = \mathrm{T}) = \sum_{s,r,l \in \{T,F\}} p(W = \mathrm{T}, S = s, R = r, C = c)$$

- The Bayesian network allows us to simplify the computation of this JPD by encoding independence relations between variables, so that:

$$p(W, S, R, C) = p(W|S, R)p(S|C)p(R|C)p(C)$$

- and hence $p(W = \mathrm{T}) =$

$$\sum_{s,r,l \in \{T,F\}} p(W = \mathrm{T}|S = s, R = r)p(S = s|C = c)p(R = r|C = c)p(C = c)$$

Naive Bayes classifier I

- A standard Naive Bayes model is a Bayesian network with a single class variable C that influences a set of evidence variables $E_1, \ldots, E_n$ (the evidence), which do not depend on each other.

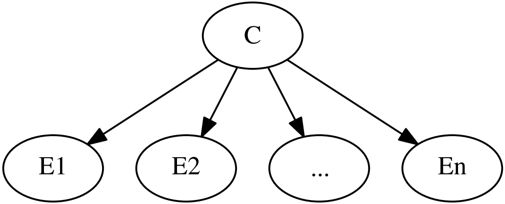

Naive Bayes classifier II

- A Naive Bayes classifier computes the marginal probability of a class, given the evidence:

$$p(c) = \sum_{e_1,\ldots,e_n} p(c \mid e_1, \ldots, e_n) p(e_1) \ldots p(e_n)$$

where c is the value of C, e_i is the value of E_i $(1 \leq i \leq n)$ and the conditional probability of the class given the evidence is estimated thus:

$$\hat{p}(c \mid e_1, \ldots, e_n) = \frac{p(c)p(e_1 \mid c) \ldots p(e_n \mid c)}{\sum_{C=c'} p(c')p(e_1 \mid c') \ldots p(e_n \mid c')}$$

Outline

Conditional probabilities in ProbTTR

- We use $p(T_1 || T_2)$ to represent the estimated conditional probability that any situation, s, is of type T_1 given that it is of type T_2.
- This contrasts with two other probability judgements in probTTR:
  - $p(s_1 : T_1 | s_2 : T_2)$, the probability that a particular situation, s_1, is of type T_1 given that s_2 is of type T_2
  - $p(T_1 | T_2)$, the probability that there is a situation of type T_1 given that there is a situation of type T_2.
- In addition there are "mixed" probabilities such as $p(T_1 | s : T_2)$, the probability that there is a situation of type T_1 given that $s : T_2$.

Outline

Random variables in TTR I

- To do probabilistic inference in ProbTTR, we need a type theoretic counterpart of a random variable in probabilistic inference.
- Assume a single (discrete) random variable with a range of possible (mutually exclusive) values.
- We introduce a *variable type* $\mathbb{V}$ whose range is a set of *value types* $\mathfrak{R}(\mathbb{V}) = \{A_1, \ldots, A_n\}$ such that the following conditions hold.
  a. $A_j \sqsubseteq \mathbb{V}$ for $1 \leq j \leq n$
  b. $A_j \perp A_i$ for all i, j such that $1 \leq i \neq j \leq n$
  c. for any s, $p(s : \mathbb{V}) \in \{0, 1.0\}$ and $p(s : \mathbb{V}) = \sum_{A \in \mathfrak{R}(\mathbb{V})} p(s : A)$

Random variables in TTR II

a. $A_j \sqsubseteq \mathbb{V}$ for $1 \leq j \leq n$

- (a) says that all value types for a variable type $\mathbb{V}$ are subtypes of $\mathbb{V}$.
  - (A type T_1 is a subtype of type T_2, $T_1 \sqsubseteq T_2$, just in case $a : T_1$ implies $a : T_2$ no matter what we assign to the basic types.)
- A simple way of achieving this is to let $\mathbb{V} = A_i \vee \ldots \vee A_n$.
  - ($T_1 \vee T_2$ is the *join type* of T_1 and T_2. $a : T_1 \vee T_2$ just in case either $a : T_1$ or $a : T_2$).

Random variables in TTR III

b. $A_j \perp A_i$ for all i, j such that $1 \leq i \neq j \leq n$

c. for any s, $p(s : \mathbb{V}) \in \{0, 1.0\}$ and $p(s : \mathbb{V}) = \sum_{A \in \mathfrak{R}(\mathbb{V})} p(s : A)$

- (b) says that all value types for a given variable type V are mutually exclusive, i.e. there are no objects that are of two value types for V.
- (c) says that the probability of a situation s being of a variable type V is either 0 or 1.0.
  - If it is 0 (i.e., the variable has no value for the situation), then the probabilities that s is of each of the value types for $\mathbb{V}$ sum to 0;
  - otherwise these probabilities sum to 1.0.

Random variables in TTR IV

- (c) encodes a conceptual difference between the probability that something has a property (such as colour, $p(s{:}\text{Colour})$), and the probability that it has a certain value of a variable (e.g. $p(s{:}\text{Green})$).
- If the probability distribution over different values (colours) sums to 1.0, then the probability that the object in question has a colour is 1.0.
- The probability that an object has colour is either 0 or 1.0.
- We assume that certain ontological/conceptual type judgements of the form “physical objects have colour” are categorical (which in a probabilistic framework means they have probability 0 or 1.0).

Random variables in TTR V

- Sprinkler example:
- Four binary variable types *Grass*, *Sprinkler*, *Raining* and *Cloudy* with corresponding variable value types:
  $\mathfrak{R}(Grass)=\{GrassWet, GrassDry\}$
  $\mathfrak{R}(Sprinkler)=\{SprinklerOn, SprinklerOff\}$
  $\mathfrak{R}(Raining)=\{IsRaining, IsNotRaining\}$
  $\mathfrak{R}(Cloudy)=\{ItIsCloudy, ItIsNotCloudy\}$
- We assume $Grass=GrassWet \vee GrassDry$, and similarly for the other variable types.
- This ensures
  a. $GrassWet \sqsubseteq Grass$ etc.
  b. $GrassWet \perp GrassDry$ etc.
  c. $p(s : Grass) = p(s : GrassWet) + p(s : GrassDry)$

Outline

A ProbTTR Naive Bayes classifier I

- Corresponding to the evidence, class variables, and their values, we associate with a ProbTTR Naive Bayes classifier κ
  a. a collection of *evidence variable types* $\mathbb{E}_1^\kappa, \ldots, \mathbb{E}_n^\kappa$,
  b. associated sets of *evidence value types* $\mathfrak{R}(\mathbb{E}_1^\kappa), \ldots, \mathfrak{R}(\mathbb{E}_n^\kappa)$,
  c. a *class variable type* $\mathbb{C}^\kappa$, and
  d. an associated set of *class value types* $\mathfrak{R}(\mathbb{C}^\kappa)$.

A ProbTTR Naive Bayes classifier II

- To classify a situation s using a classifier κ, the evidence is acquired by observing and classifying s with respect to the evidence types.
- This can be done through another layer of probabilistic classification based on yet another set of evidence types.
- Type judgements can also be obtained directly from probabilistic or non-probabilistic classification of low-level sensory readings supplied by observation.

A ProbTTR Naive Bayes classifier III

- A ProbTTR Naïve Bayes classifier is a function κ of the type

$$(\mathbb{E}_1^\kappa \wedge \ldots \wedge \mathbb{E}_n^\kappa) \to \text{Set}(\begin{bmatrix} \text{sit} & : & \mathit{Sit} \\ \text{sit-type} & : & \mathit{Type} \\ \text{prob} & : & [0,1] \end{bmatrix})$$

such that if $s : \mathbb{E}_1^\kappa \wedge \ldots \wedge \mathbb{E}_n^\kappa$, then

$$\kappa(s) = \{\begin{bmatrix} \text{sit} & = & s \\ \text{sit-type} & = & C \\ \text{prob} & = & p^\kappa(s : C) \end{bmatrix} \mid C \in \mathfrak{R}(\mathbb{C}^\kappa)\}$$

where

$$p^\kappa(s : C) = \sum_{\substack{E_1 \in \mathfrak{R}(\mathbb{E}_1^\kappa) \\ \vdots \\ E_n \in \mathfrak{R}(\mathbb{E}_n^\kappa)}} p^\kappa(C || E_1 \wedge \ldots \wedge E_n) p(s : E_1) \ldots p(s : E_n)$$

A ProbTTR Naive Bayes classifier IV

- ($T_1 \wedge T_2$ is the *meet type* of T_1 and T_2. $a : T_1 \wedge T_2$ just in case $a : T_1$ and $a : T_2$.)
- When using κ, we are interested in the marginal probability $p^{\kappa}(s : C)$ of the situation s being of a class value type C in light of the evidence concerning s.
- As in the case of standard Bayesian Networks, we obtain the marginal probabilities of a class value type C by summing over all combinations of evidence value types.
- The classifier gives a probability distribution over the class value types, encoded as a set of probabilistic Austinian propositions.

A ProbTTR Naive Bayes classifier V

- As above, for the Naive Bayes classifier we estimate the conditional probability of the class given the evidence using the assumption that the evidence variable types are independent:

  $\hat{p}^{\kappa}(C||E_1 \wedge \ldots \wedge E_n) =$

  $$\frac{p(C)p(E_1||C)\ldots p(E_n||C)}{\sum_{C' \in \mathfrak{R}(\mathbb{C}^{\kappa})} p(C')p(E_1||C')\ldots p(E_n||C')}$$

- Estimating $p(C)$ and $p(E_i||C)$ is part of the learning theory

Outline

Semantic Classification: Example I

- We will now illustrate classification in ProbTTR using a Naive Bayes classifier for fruits.
- We can imagine this classification taking place in the setting of a *Fruit Recognition Game*.
- In this game a teacher shows a learning agent fruits (for simplicity, we assume there are only apples and pears in this instance of the game).
- The agent makes a guess, the teacher provides the correct answer, and the agent learns from these observations.
  - Here, we only describe the classification step.

Semantic Classification: Example II

- We will use shorthand for the types corresponding to an object being an apple vs. a pear:
  - $Apple = \begin{bmatrix} \text{x} & : & \mathit{Ind} \\ \text{c}_{\text{apple}} & : & \text{apple(x)} \end{bmatrix}$
  - $Pear = \begin{bmatrix} \text{x} & : & \mathit{Ind} \\ \text{c}_{\text{pear}} & : & \text{pear(x)} \end{bmatrix}$

Semantic Classification: Example III

- Objects in the Fruit Recognition Game have one of two shapes (a-shape or p-shape) and one of two colours (green or red).
  - $Ashape = \left[\begin{array}{lll} \text{x} & : & Ind \\ \text{c} & : & \text{ashape(x)} \end{array}\right]$
  - $Pshape = \left[\begin{array}{lll} \text{x} & : & Ind \\ \text{c} & : & \text{pshape(x)} \end{array}\right]$
  - $Green = \left[\begin{array}{lll} \text{x} & : & Ind \\ \text{c} & : & \text{green(x)} \end{array}\right]$
  - $Red = \left[\begin{array}{lll} \text{x} & : & Ind \\ \text{c} & : & \text{red(x)} \end{array}\right]$

Semantic Classification: Example IV

- The class variable type is *Fruit*, with value types $\mathfrak{R}(Fruit) = \{Apple, Pear\}$.
- The evidence variable types are
  - *Col*(our), with value types $\mathfrak{R}(Col) = \{Green, Red\}$
  - Shape, with value types $\mathfrak{R}(Shape) = \{Ashape, Pshape\}$.

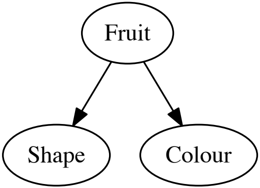

Classification in the Apple game

- For a situation s the classifier FruitC(s) returns a set of probabilistic Austinian propositions asserting that s instantiates a certain type of fruit.
- This set is a probability distribution over the variable types of *Fruit*.

$$\text{FruitC}(s) = \{ \begin{bmatrix} \text{sit} & = & s \\ \text{sit-type} & = & F \\ \text{prob} & = & p_{\mathcal{J}}^{\text{FruitC}}(s : F) \end{bmatrix} \mid F \in \mathcal{R}(\mathit{Fruit}) \}$$

- Probability of a fruit type judgement in the Fruit Recognition Game:

$$p^{\text{FruitC}}(s : F) = \sum_{\substack{L \in \mathcal{R}(\text{Col}) \\ S \in \mathcal{R}(\text{Shape})}} p(F || L \wedge S) p(s : L) p(s : S)$$

Classification in the Apple game, cont'd

- To determine the probability that a situation is of the apple type, we sum over the various evidence type values for apple.
- $p^{\mathrm{FruitC}}(s : Apple) =$

$$\sum_{\substack{L \in \mathcal{R}(\mathrm{Col}) \\ S \in \mathcal{R}(\mathrm{Shape})}} p(Apple||L \wedge S)p(s : L)p(s : S) =$$

$$p(Apple||Green \wedge Ashape)p(s : Green)p(s : Ashape)+$$
$$p(Apple||Green \wedge Pshape)p(s : Green)p(s : Pshape)+$$
$$p(Apple||Red \wedge Ashape)p(s : Red)p(s : Ashape)+$$
$$p(Apple||Red \wedge Pshape)p(s : Red)p(s : Pshape)$$

Conditional probabilities used by classifier

- Conditional probabilities for the fruit classifier are derived from previous judgements of the form $p(F||C \wedge S)$
- The example values in the matrix below illustrates a JPD for the apple classifier:

Apple/Pear	*Ashape*	*Pshape*
Green	0.93/0.07	0.63/0.37
Red	0.56/0.44	0.13/0.87

- For example, $p(Apple||Green \wedge Ashape) = 0.93$

Evidence used by the classifier

- The non-conditional probabilities are derived from the agents' take on the particular situation being classified; let's call it s_5.

	T=*Ashape*	T=*Pshape*	T=*Green*	T=*Red*
$p(s_5{:}T)$	0.90	0.10	0.80	0.20

- We can think of these probabilities as resulting from probabilistic classification of real-valued visual input, where a classifier assigns to each image a probability that the image shows a situation of the respective type.

Classification in the Apple game, cont'd

With these numbers in place, we can compute the probability that the fruit being classified is an apple:

$$
\begin{aligned}
&p^{\text{FruitC}}(\mathrm{s}_5\text{: } \mathit{Apple}) = \\
&\quad p(\mathit{Apple}||\mathit{Green} \wedge \mathit{Ashape})p(s : \mathit{Green})p(s : \mathit{Ashape}) + \\
&\quad p(\mathit{Apple}||\mathit{Green} \wedge \mathit{Pshape})p(s : \mathit{Green})p(s : \mathit{Pshape}) + \\
&\quad p(\mathit{Apple}||\mathit{Red} \wedge \mathit{Ashape})p(s : \mathit{Red})p(s : \mathit{Ashape}) + \\
&\quad p(\mathit{Apple}||\mathit{Red} \wedge \mathit{Pshape})p(s : \mathit{Red})p(s : \mathit{Pshape}) = \\
&\quad 0.93{*}0.80{*}0.90 + \\
&\quad 0.63{*}0.80{*}0.10 + \\
&\quad 0.56{*}0.20{*}0.90 + \\
&\quad 0.13{*}0.20{*}0.10 = \\
&\quad 0.67 + 0.05 + 0.10 + 0.00 = \\
&\quad 0.82
\end{aligned}
$$

Outline

Perceiving evidence I

- Where do the non-conditional probabilities of the evidence variables concerning the situation s being classified come from?
- We suggest regarding these probabilities as resulting from probabilistic classification of real-valued (non-symbolic) visual input, where a classifier assigns to each image a probability that the image shows a situation of the respective type.
- Such a classifier can be implemented in a number of different ways, e.g. as a deep neural network, as long as it outputs a probability distribution.
- Larsson (2015) shows how perceptual classification can be modelled in TTR, and Larsson (2020) reformulates and extends this formalisation to probabilistic classification.

Perceiving evidence II

- A probabilistic perceptual classifier, corresponding to an evidence variable type $\mathbb{E}_i$ $(1 \leq i \leq n)$, provides a mapping
  - from perceptual input (of a type $\mathfrak{V}$, e.g. a digital image)
  - onto a probability distribution over evidence value types in $\mathfrak{R}(\mathbb{E}_i^\kappa)$,
- ...the latter encoded as a set of probabilistic Austinian propositions:
  - $\pi_{E_i^\kappa} : Sit_{\mathfrak{V}} \rightarrow \left\{ \begin{bmatrix} \text{sit} & : & Sit_{\mathfrak{V}} \\ \text{sit-type} & : & RecType_{E_i} \\ \text{prob} & : & [0,1] \end{bmatrix} \mid E_i \in \mathfrak{R}(\mathbb{E}_i^\kappa) \right\}$
- where $Sit_{\mathfrak{V}}$ is the type of situations where perception of some object (labelled x) yields visual information (labelled c) concerning x:
  - $Sit_{\mathfrak{V}} = \begin{bmatrix} \text{x} & : & Ind \\ \text{c} & : & \mathfrak{V} \end{bmatrix}$

Perceiving evidence III

- $RecType_R$ is the (singleton) type of record types that are identical to R, so that e.g.
  - $T : RecType_{Green}$ iff $T : RecType$ and $T = Green$
- In the Apple game, an agent would be equipped with visual classifiers corresponding to *Shape* and *Col*, where e.g.
  - $\pi_{Col} : \begin{bmatrix} \text{x} & : & Ind \\ \text{c} & : & \mathfrak{V} \end{bmatrix} \rightarrow \left\{ \begin{bmatrix} \text{sit} & : & Sit_{\mathfrak{V}} \\ \text{sit-type} & : & RecType_{Green} \\ \text{prob} & : & [0,1] \end{bmatrix}, \begin{bmatrix} \text{sit} & : & Sit_{\mathfrak{V}} \\ \text{sit-type} & : & RecType_{Red} \\ \text{prob} & : & [0,1] \end{bmatrix} \right\}$

Perceiving evidence IV

- If we e.g. assume $s_5 = \begin{bmatrix} \text{x} & = & \text{a}_{453} \\ \text{c} & = & \text{Img}_{9876} \end{bmatrix}$ where
  - $\text{a}_{453} : \mathit{Ind}$
  - $\text{Img}_{9876} : \mathfrak{V}$
- and we assume that
  - $\pi_{Col}(s_5) = \{ \begin{bmatrix} \text{sit} & = & s_5 \\ \text{sit-type} & = & \mathit{Green} \\ \text{prob} & = & 0.8 \end{bmatrix}, \begin{bmatrix} \text{sit} & = & s_5 \\ \text{sit-type} & = & \mathit{Red} \\ \text{prob} & = & 0.2 \end{bmatrix} \}$
- then
  - $p(s_5 : \mathit{Green}) = 0.8$
  - $p(s_5 : \mathit{Red}) = 0.2$
- This illustrates how ProbTTR allows combining probabilistic perceptual classification and probabilistic reasoning.

Outline

Bayesian networks in TTR I

- To extend the above to full Bayesian networks, we need to distinguish evidence variables from *unobserved variables*, and incorporate the latter into our classifier.
- A TTR Bayes net classifier is associated with
  - $\mathbb{E}_1^\kappa, \ldots, \mathbb{E}_n^\kappa$ is a collection of evidence variable types,
  - $\mathfrak{R}(\mathbb{E}_1^\kappa), \ldots, \mathfrak{R}(\mathbb{E}_n^\kappa)$ are sets of evidence value types,
  - $\mathbb{I}_1^\kappa, \ldots, \mathbb{I}_m^\kappa$ is a collection of unobserved variable types,
  - $\mathfrak{R}(\mathbb{I}_1^\kappa), \ldots, \mathfrak{R}(\mathbb{I}_m^\kappa)$ are sets of unobserved value types.

Bayesian networks in TTR II

- We can use a TTR Bayes net as a classifier, i.e., a function κ of type

$$\mathbb{E}^\kappa_1 \wedge \ldots \wedge \mathbb{E}^\kappa_n \rightarrow \text{Set}(\begin{bmatrix} \text{sit} & : & \textit{Sit} \\ \text{sit-type} & : & \textit{Type} \\ \text{prob} & : & [0,1] \end{bmatrix})$$

such that if $s : \mathbb{E}^\kappa_1 \wedge \ldots \wedge \mathbb{E}^\kappa_n$ and $1 \leq j \leq m$, then

$$\kappa(s) = \{ \begin{bmatrix} \text{sit} & = & s \\ \text{sit-type} & = & I_j \\ \text{prob} & = & p^\kappa(s : I_j) \end{bmatrix} \mid I_j \in \mathfrak{R}(\mathbb{I}^\kappa_j) \}$$

where

$$p^\kappa(s : I_j) = \sum_{\substack{I_1 \in \mathfrak{R}(\mathbb{I}^\kappa_1) \\ \cdots \\ I_{j-1} \in \mathfrak{R}(\mathbb{I}^\kappa_{j-1}) \\ I_{j+1} \in \mathfrak{R}(\mathbb{I}^\kappa_{j+1}) \\ \cdots \\ I_m \in \mathfrak{R}(\mathbb{I}^\kappa_m) \\ E_1 \in \mathfrak{R}(\mathbb{E}^\kappa_1) \\ \cdots \\ E_n \in \mathfrak{R}(\mathbb{E}^\kappa_n)}} p(I_j || I_1 \wedge \ldots \wedge I_{j-1} \wedge I_{j+1} \wedge \ldots \wedge I_m \wedge E_1 \wedge \ldots \wedge E_n) p(s : E_1) \ldots p(s : E_n)$$

Bayesian networks in TTR III

- The dependencies encoded in a Bayes net will affect how the conditional probability

  $p(C||I_1 \wedge \ldots I_{j-1} \wedge I_{j+1} \wedge I_m \wedge E_1 \wedge \ldots \wedge E_n)$

  is computed.
- In the sprinkler example, we have three unobserved variable types *Grass*, *Sprinkler* and *Rain*, and one evidence variable type *Cloudy*.

Bayesian networks in TTR IV

- For $S \in \mathfrak{R}(\mathit{Sprinkler}), R \in \mathfrak{R}(\mathit{Rain}), L \in \mathfrak{R}(\mathit{Cloudy})$ and $G \in \mathfrak{R}(\mathit{Grass})$, the dependencies encoded in the Bayesian network above entail that $p(G||S \wedge R \wedge L) =$

$$p(G||S \wedge R)p(S||L)p(R||L)$$

and hence for $G \in \mathfrak{R}(\mathit{Grass})$,

$$p^{\kappa}(s : G) = \sum_{\substack{S \in \mathfrak{R}(\mathit{Sprinkler}) \\ R \in \mathfrak{R}(\mathit{Raining}) \\ L \in \mathfrak{R}(\mathit{Cloudy})}} p(G||S \wedge R)p(S||L)p(R||L)p(s : L)$$

Outline

Conclusion

- We have proposed a Bayesian account of semantic classification and inference formulated in terms of probabilistic type theory.
- Elements that were added to ProbTTR:
  - The notion of a random variable, seen as a type
  - Evidence, class and unobserved variable types and value types
  - Naive Bayes Classifiers and Bayesian Networks
- This gives us one of the building blocks for a probabilistic type theory that
  - combines explainable probabilistic reasoning and learning with black-box (e.g. neural net) perceptual classification,
  - allowing us to model (semantic and factual) learning from perceptual experience and linguistic interaction.

Future work

- Learning in the frequentist model (submitted)
- Classification and learning using linear transformation model (relating to Bernardy *et al.* (2018); work in progress)
- Investigate how probabilistic dependencies can be learned from interaction (submitted)
- Implementing probabilistic TTR (work in progress)
- Explore using ProbTTR as a framework for NLI tasks, especially those taking both visual and linguistic input, ideally making use of previous TTR and ProbTTR work on various phenomena in natural language semantics
- ...

Bernardy, Jean-Philippe; Blanck, Rasmus; Chatzikyriakidis, Stergios; and Lappin, Shalom 2018.
A compositional bayesian semantics for natural language.
In *Proceedings of the first international workshop on language cognition and computational models.*
1–10.

Breitholtz, Ellen and Cooper, Robin 2011.
Enthymemes as rhetorical resources.
In Artstein, Ron; Core, Mark; DeVault, David; Georgila, Kallirroi; Kaiser, Elsi; and Stent, Amanda, editors 2011, *Proceedings of the 15th Workshop on the Semantics and Pragmatics of Dialogue (SemDial 2011)*, Los Angeles (USA). Institute for Creative Technologies.
149–157.

Breitholtz, Ellen 2020.
Enthymemes and Topoi in Dialogue: The Use of Common Sense Reasoning in Conversation.
Brill, Leiden, The Netherlands.

Cooper, Robin and Ginzburg, Jonathan 2011.
Negation in dialogue.
In Artstein, Ron; Core, Mark; DeVault, David; Georgila, Kallirroi; Kaiser, Elsi; and Stent, Amanda, editors 2011, *SemDial 2011 (Los Angelogue): Proceedings of the 15th Workshop on the Semantics and Pragmatics of Dialogue*.
130–139.

Cooper, Robin and Ginzburg, Jonathan 2015.
Type theory with records for natural language semantics.
In Lappin, Shalom and Fox, Chris, editors 2015, *The Handbook of Contemporary Semantic Theory, Second Edition*. Wiley-Blackwell, Oxford and Malden.
375–407.

Cooper, Robin and Larsson, Staffan 2009.
Compositional and ontological semantics in learning from corrective feedback and explicit definition.

In Edlund, Jens; Gustafson, Joakim; Hjalmarsson, Anna; and Skantze, Gabriel, editors 2009, *Proceedings of DiaHolmia, 2009 Workshop on the Semantics and Pragmatics of Dialogue*.

Cooper, Robin; Dobnik, Simon; Lappin, Shalom; and Larsson, Staffan 2014.
A probabilistic rich type theory for semantic interpretation.
In *Proceedings of the EACL 2014 Workshop on Type Theory and Natural Language Semantics (TTNLS)*. Gothenburg, Association of Computational Linguistics.
72–79.

Cooper, Robin; Dobnik, Simon; Lappin, Shalom; and Larsson, Staffan 2015a.
Probabilistic type theory and natural language semantics.
Linguistic Issues in Language Technology 10 1–43.

Cooper, Robin; Dobnik, Simon; Larsson, Staffan; and Lappin, Shalom 2015b.
Probabilistic type theory and natural language semantics.

LiLT (Linguistic Issues in Language Technology) 10.

Cooper, Robin 1998.
Information states, attitudes and dependent record types.
In *ITALLC98*.
85–106.

Cooper, Robin 2004.
Dynamic generalised quantifiers and hypothetical contexts.
In Svennerlind, Christer, editor 2004, *Ursus Philosophicus, a festschrift for Björn Haglund*. Department of Philosophy, University of Gothenburg, Gothenburg (Sweden).

Cooper, Robin 2005.
Records and record types in semantic theory.
Journal of Logic and Computation 15(2):99–112.

Cooper, Robin 2010.
Generalized quantifiers and clarification content.

In Łupkowski, Paweł and Purver, Matthew, editors 2010, *Aspects of Semantics and Pragmatics of Dialogue. SemDial 2010, 14th Workshop on the Semantics and Pragmatics of Dialogue*, Poznań. Polish Society for Cognitive Science.

Cooper, Robin 2011.
Copredication, quantification and frames.
In Pogodalla, Sylvain and Prost, Jean-Philippe, editors 2011, *LACL*, volume 6736 of *Lecture Notes in Computer Science*. Springer. 64–79.

Cooper, Robin 2012a.
Type theory and semantics in flux.
In Kempson, Ruth; Asher, Nicholas; and Fernando, Tim, editors 2012a, *Handbook of the Philosophy of Science*, volume 14: Philosophy of Linguistics. Elsevier BV.
General editors: Dov M. Gabbay, Paul Thagard and John Woods.

Cooper, Robin 2012b.
Type theory and semantics in flux.

In Kempson, Ruth; Asher, Nicholas; and Fernando, Tim, editors 2012b, *Handbook of the Philosophy of Science*, volume 14: Philosophy of Linguistics. Elsevier BV.
General editors: Dov M. Gabbay, Paul Thagard and John Woods.

Cooper, Robin prep.
From perception to communication: An analysis of meaning and action using a theory of types with records (TTR).
Draft available from https://sites.google.com/site/typetheorywithrecords/drafts.

Dobnik, Simon and Cooper, Robin 2013.
Spatial descriptions in type theory with records.
In *Proceedings of IWCS 2013 Workshop on Computational Models of Spatial Language Interpretation and Generation (CoSLI-3)*, Potsdam, Germany. Association for Computational Linguistics.
1–6.

Dobnik, Simon; Larsson, Staffan; and Cooper, Robin 2011.
Toward perceptually grounded formal semantics.

In *Proceedings of the Workshop on Integrating Language and Vision at NIPS 2011*, Sierra Nevada, Spain. Neural Information Processing Systems Foundation (NIPS).

Fernández, Raquel and Larsson, Staffan 2014.
Vagueness and learning: A type-theoretic approach.
In *Proceedings of the 3rd Joint Conference on Lexical and Computational Semantics (∗SEM 2014)*.

Ginzburg, Jonathan 2012.
The Interactive Stance: Meaning for Conversation.
Oxford University Press, Oxford.

Halpern, J. 2003.
Reasoning About Uncertainty.
MIT Press, Cambridge MA.

Larsson, Staffan 2011.
The TTR perceptron: Dynamic perceptual meanings and semantic coordination.

In *Proceedings of the 15th Workshop on the Semantics and Pragmatics of Dialogue (SemDial 2011)*, Los Angeles (USA).

Larsson, Staffan 2015.
Formal semantics for perceptual classification.
Journal of Logic and Computation 25(2):335–369.
Published online 2013-12-18.

Larsson, Staffan 2020.
Discrete and probabilistic classifier-based semantics.
In *Proceedings of the Probability and Meaning Conference (PaM 2020)*, Gothenburg. Association for Computational Linguistics.
62–68.

Pearl, J. 1990.
Bayesian decision methods.
In Shafer, G. and Pearl, J., editors 1990, *Readings in Uncertain Reasoning*. Morgan Kaufmann.
345–352.

Russell, Stuart and Norvig, Peter 1995.

Artificial Intelligence: A Modern Approach.
Prentice Hall Series in Artificial Intelligence. Englewood Cliffs, New Jersey.