

# The role of definitions in coordinating on perceptual meanings

Staffan Larsson

Centre for Linguistic Theory and Studies in Probability (CLASP)

Dept. of Philosophy, Linguistics and Theory of Science

University of Gothenburg

PotsDial 2021

25th Workshop on the Semantics and Pragmatics of Dialogue

# Outline

## Introduction

## Background

- TTR fundamentals

- Probabilistic TTR fundamentals

- Bayesian nets and the Naive Bayes classifier

- A ProbTTR Naive Bayes classifier

- Semantic classification

- Semantic learning

## Word Meaning Negotiation and semantic updates

## Learning perceptual meanings from definitions

- Learning a new meaning from example

- Learning a new meaning from definition

- Learning new evidence values

- Learning new evidence variables

## Definitions vs. examples

## Conclusions & future work

# Outline

## Introduction

## Background

- TTR fundamentals

- Probabilistic TTR fundamentals

- Bayesian nets and the Naive Bayes classifier

- A ProbTTR Naive Bayes classifier

- Semantic classification

- Semantic learning

## Word Meaning Negotiation and semantic updates

## Learning perceptual meanings from definitions

- Learning a new meaning from example

- Learning a new meaning from definition

- Learning new evidence values

- Learning new evidence variables

## Definitions vs. examples

## Conclusions & future work

# 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), Larsson (2015), Cooper *et al.* (2015b), Larsson (2020), Larsson and Cooper (2021)
- ▶ 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)


# Learning from demonstrations and definitions I

- ▶ Human first language learners typically learn from demonstrations, where a word becomes associated with a perceptual stimuli.
- ▶ This kind of semantic learning can be modeled as training a perceptual classifier on new perceptually available examples (Larsson, 2015), (Larsson, 2020).
- ▶ However, it also seems clear that at least adult humans can learn tentative new meanings, including perceptual meanings, from verbal descriptions.


# Learning from demonstrations and definitions II

- ▶ This paper explores how learning of perceptual meaning can be initialised by partial definitions given in interaction, provided that the words used in the definition themselves have known perceptual meanings.
- ▶ In brief, the idea is that definitions provide hints on a structure (here, a Naive Bayes classifier) connecting the defined concept with the concepts used in the definition.
- ▶ The defined concept is an unobserved variable (for a classifier, the class variable), and the concepts used in the definition are evidence variables.

# Classification




# Classification is subjective?







# Coordination process




# Classification is coordinated



# Classification is coordinated



# Coordination can be creative



## Previous work

- ▶ An account of probabilistic inference and classification in ProbTTR is introduced in Larsson and Cooper (2021), where it is also demonstrated how probabilistic classification of perceptual evidence can be combined with probabilistic reasoning.
- ▶ Building on Larsson and Cooper (2021), Larsson *et al.* (2021) propose a probabilistic account of semantic learning from interaction formulated in terms of a Probabilistic Type Theory with Records (ProbTTR) (Cooper *et al.*, 2015a).
- ▶ Starting from a probabilistic type theoretic formulations of naive Bayes classifiers, the account of semantic learning is illustrated with a simple language game (the fruit recognition game).
- ▶ Here, we will connect these strands of work in an attempt to provide a formal account of the role of definitions in semantic coordination, and in particular for perceptual meanings.

# Outline

## Introduction

## Background

- TTR fundamentals

- Probabilistic TTR fundamentals

- Bayesian nets and the Naive Bayes classifier

- A ProbTTR Naive Bayes classifier

- Semantic classification

- Semantic learning

## Word Meaning Negotiation and semantic updates

## Learning perceptual meanings from definitions

- Learning a new meaning from example

- Learning a new meaning from definition

- Learning new evidence values

- Learning new evidence variables

## Definitions vs. examples

## Conclusions & future work

# Type Theory with Records

- ▶ We want to use a framework which also encompasses accounts of many problems traditionally studied in formal semantics<sup>1</sup>.
- ▶ 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*.

---

<sup>1</sup>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).

# Outline

## Introduction

## Background

### TTR fundamentals

#### Probabilistic TTR fundamentals

#### Bayesian nets and the Naive Bayes classifier

#### A ProbTTR Naive Bayes classifier

#### Semantic classification

#### Semantic learning

## Word Meaning Negotiation and semantic updates

## Learning perceptual meanings from definitions

#### Learning a new meaning from example

#### Learning a new meaning from definition

#### Learning new evidence values

#### Learning new evidence variables

## Definitions vs. examples

## Conclusions & future work



# 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  $T_{\text{true}}$  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  $a, b : Ind$ .
- ▶ 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

- $a_1 : T_1,$
- $a_2 : T_2(a_1),$
- $\dots,$
- $a_n : T_n(a_1, a_2, \dots, a_{n-1}),$
- where  $T(a_1, \dots, a_n)$  represents a type  $T$  which depends on the objects  $a_1, \dots, a_n,$

- ...the record to the left is of the record type to the right.

$$\left[ \begin{array}{lcl} \ell_1 & = & a_1 \\ \ell_2 & = & a_2 \\ \dots & & \\ \ell_n & = & a_n \\ \dots & & \end{array} \right] : \left[ \begin{array}{lcl} \ell_1 & : & T_1 \\ \ell_2 & : & T_2(\ell_1) \\ \dots & & \\ \ell_n & : & T_n(\ell_1, \ell_2, \dots, \ell_{n-1}) \end{array} \right]$$

- $\ell_1, \dots, \ell_n$  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:

$$\left[ \begin{array}{lcl} \text{ref} & = & \text{obj}_{123} \\ \text{c}_{\text{man}} & = & \text{prf1} \\ \text{c}_{\text{run}} & = & \text{prf2} \end{array} \right] : \left[ \begin{array}{lcl} \text{ref} & : & \text{Ind} \\ \text{c}_{\text{man}} & : & \text{man}(\text{ref}) \\ \text{c}_{\text{run}} & : & \text{run}(\text{ref}) \end{array} \right]$$

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

# Outline

## Introduction

## Background

- TTR fundamentals

- Probabilistic TTR fundamentals**

- Bayesian nets and the Naive Bayes classifier

- A ProbTTR Naive Bayes classifier

- Semantic classification

- Semantic learning

## Word Meaning Negotiation and semantic updates

## Learning perceptual meanings from definitions

- Learning a new meaning from example

- Learning a new meaning from definition

- Learning new evidence values

- Learning new evidence variables

## Definitions vs. examples

## Conclusions & future work

# Probabilistic TTR fundamentals

- ▶ 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 Austinian propositions

- ▶ 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

$$\left[ \begin{array}{ll} \text{sit} & : \text{Sit} \\ \text{sit-type} & : \text{Type} \\ \text{prob} & : [0,1] \end{array} \right]$$

(where  $[0, 1]$  represents the type of real numbers between 0 and 1).

- ▶ An object  $\varphi$  of the above type corresponds to, or encodes, the judgement

$$p(\varphi.\text{sit} : \varphi.\text{sit-type}) = \varphi.\text{prob}$$



# Conditional probabilities in ProbTTR

- ▶ We use  $p(T_1||T_2)$  to represent the probability that any situation  $s$  is of type  $T_1$ , given that  $s$  is of type  $T_2$ .
- ▶ Note that  $p(T_1||T_2)$ , is different from  $p(T_1|T_2)$ , the probability of there being something of type  $T_1$  given that there is something of type  $T_2$ .

# Random variables in TTR I

- ▶ Larsson and Cooper (2021) introduce a type theoretic counterpart of a random variable in Bayesian inference.
- ▶ To represent a single (discrete) random variable with a range of possible (mutually exclusive) values, ProbTTR uses a *variable type*  $V$  whose range is a set of *value types*  $\mathfrak{R}(V) = \{A_1, \dots, A_n\}$  such that the following conditions hold.
  - $A_j \sqsubseteq \mathbb{V}$  for  $1 \leq j \leq n$
  - $A_j \perp A_i$  for all  $i, j$  such that  $1 \leq i \neq j \leq n$
  - for any  $s$ ,  $p(s : \mathbb{V}) \in \{0, 1.0\}$  and  $p(s : \mathbb{V}) = \sum_{A \in \mathfrak{R}(\mathbb{V})} p(s : A)$

# Representing probability distributions

- For a situation  $s$ , a probability distribution over the  $m$  value types  $A_j \in \mathfrak{R}(\mathbb{A}), 1 \leq j \leq m$  belonging to a variable type  $\mathbb{A}$  can be written (as above) as a set of probabilistic Austinian propositions, e.g.

$$\left\{ \begin{array}{lcl} \text{sit} & = & s \\ \text{sit-type} & = & A_j \\ \text{prob} & = & p(s : A_j) \end{array} \right\} \mid A_j \in \mathfrak{R}(\mathbb{A}) \}$$

# Outline

## Introduction

## Background

- TTR fundamentals

- Probabilistic TTR fundamentals

- Bayesian nets and the Naive Bayes classifier**

- A ProbTTR Naive Bayes classifier

- Semantic classification

- Semantic learning

## Word Meaning Negotiation and semantic updates

## Learning perceptual meanings from definitions

- Learning a new meaning from example

- Learning a new meaning from definition

- Learning new evidence values

- Learning new evidence variables

## Definitions vs. examples


## Conclusions & future work

# 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.

# 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  $E_1, \dots, E_n$  (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:

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

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:

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

# Outline

## Introduction

## Background

- TTR fundamentals

- Probabilistic TTR fundamentals

- Bayesian nets and the Naive Bayes classifier

- A ProbTTR Naive Bayes classifier**

- Semantic classification

- Semantic learning

## Word Meaning Negotiation and semantic updates

## Learning perceptual meanings from definitions

- Learning a new meaning from example

- Learning a new meaning from definition

- Learning new evidence values

- Learning new evidence variables

## Definitions vs. examples

## Conclusions & future work



# A ProbTTR Naive Bayes classifier I

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

- We can encode the variables and values associated with a Naive Bayes classifier  $\kappa$  as a TTR record:

$$v(\kappa) = \left[ \begin{array}{lcl} \text{cvar} & = & \mathbb{C}^\kappa \\ \text{cvals} & = & \mathfrak{R}(\mathbb{C}^\kappa) \\ \text{evars} & = & \{\mathbb{E}_1^\kappa, \dots, \mathbb{E}_n^\kappa\} \\ \text{evals} & = & \left\{ \begin{bmatrix} \text{lbl}(\mathbb{E}_1^\kappa) & = & \mathfrak{R}(\mathbb{E}_1^\kappa) \\ \dots \\ \text{lbl}(\mathbb{E}_n^\kappa) & = & \mathfrak{R}(\mathbb{E}_n^\kappa) \end{bmatrix} \right\} \end{array} \right]$$

# A ProbTTR Naive Bayes classifier III

- ▶ To classify a situation  $s$  using a classifier  $\kappa$ , the evidence is acquired by observing and classifying  $s$  with respect to the evidence types.
- ▶ Larsson and Cooper (2021) define a ProbTTR Bayes classifier  $\kappa$  as a function
  - ▶ from a situation  $s$  (of the meet type of the evidence variable types  $\mathbb{E}_1^\kappa, \dots, \mathbb{E}_n^\kappa$ )
  - ▶ to a set of probabilistic Austinian propositions that define a probability distribution over the values of the class variable type  $\mathbb{C}^\kappa$ ,
  - ▶ given probability distributions over the values of each evidence variable type  $\mathbb{E}_1^\kappa, \dots, \mathbb{E}_n^\kappa$ .

# A ProbTTR Naive Bayes classifier IV

- A ProbTTR Naïve Bayes classifier is a function  $\kappa$  of the type

$$(\mathbb{E}_1^\kappa \wedge \dots \wedge \mathbb{E}_n^\kappa) \rightarrow \text{Set} \left( \begin{array}{ll} \text{sit} & : \text{Sit} \\ \text{sit-type} & : \text{Type} \\ \text{prob} & : [0,1] \end{array} \right)$$

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

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

where

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

# Outline

## Introduction

## Background

- TTR fundamentals

- Probabilistic TTR fundamentals

- Bayesian nets and the Naive Bayes classifier

- A ProbTTR Naive Bayes classifier

- Semantic classification**

- Semantic learning

## Word Meaning Negotiation and semantic updates

## Learning perceptual meanings from definitions

- Learning a new meaning from example

- Learning a new meaning from definition

- Learning new evidence values

- Learning new evidence variables

## Definitions vs. examples

## Conclusions & future work

# The fruit recognition game I

- ▶ Larsson and Cooper (2021) illustrate semantic classification using a Naive Bayes classifier in ProbTTR using the *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.

# The fruit recognition game II

- ▶ We will use shorthand for the types corresponding to an object being an apple vs. a pear:

- ▶  $Apple = \begin{bmatrix} x & : & Ind \\ c_{apple} & : & apple(x) \end{bmatrix}$

- ▶  $Pear = \begin{bmatrix} x & : & Ind \\ c_{pear} & : & pear(x) \end{bmatrix}$

# The fruit recognition game 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).

- $Ashape = \begin{bmatrix} x & : & Ind \\ c & : & ashape(x) \end{bmatrix}$

- $Pshape = \begin{bmatrix} x & : & Ind \\ c & : & pshape(x) \end{bmatrix}$


- $Green = \begin{bmatrix} x & : & Ind \\ c & : & green(x) \end{bmatrix}$

- $Red = \begin{bmatrix} x & : & Ind \\ c & : & red(x) \end{bmatrix}$



# The fruit recognition game 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  $\mathfrak{R}(Shape) = \{Ashape, Pshape\}$ .



# The fruit recognition game V

- TTR encoding of variables and values associated with the Naive Bayes fruit classifier FruitC:

$$v(\text{FruitC}) = \left[ \begin{array}{lcl} \text{cvar} & = & \textit{Fruit} \\ \text{cval} & = & \{ \textit{Apple}, \textit{Pear} \} \\ \text{evars} & = & \{ \textit{Col}, \textit{Shp} \} \\ \text{evals} & = & \left\{ \begin{array}{lcl} \text{lbl}(\textit{Col}) & = & \{ \textit{Red}, \textit{Green} \} \\ \text{lbl}(\textit{Shp}) & = & \{ \textit{AShape}, \textit{PShape} \} \end{array} \right\} \end{array} \right]$$

(The function *lbl* takes a type *T* and returns a label unique to *T*)

# Classification in the fruit recognition game

- For a situation  $s$  the classifier  $\text{FruitC}(s)$  returns a probability distribution over the variable types of  $\text{Fruit}$ .

$$\text{FruitC}(s) = \left\{ \left[ \begin{array}{ll} \text{sit} & = s \\ \text{sit-type} & = F \\ \text{prob} & = p_{\mathcal{J}}^{\text{FruitC}}(s : F) \end{array} \right] \mid F \in \mathcal{R}(\text{Fruit}) \right\}$$

# Semantic classification using conditional probabilities I

- ▶ We follow Larsson and Cooper (2021) in showing how semantic classification (i.e., estimating a probability distribution over class value types) works
- ▶ In general, for class value types  $C_j \in \mathfrak{R}(\mathbb{C}^\kappa)$ , we have

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

# Semantic classification using conditional probabilities II

- ▶ Correspondingly, in the fruit recognition game, for each  $F \in \mathfrak{R}(Fruit)$  we have

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

# Semantic classification using conditional probabilities III

- ▶ Larsson (2015) shows how perceptual classification can be modelled in TTR
- ▶ Larsson (2020) reformulates and extends this formalisation to probabilistic classification.
- ▶ Larsson and Cooper (2021) suggests regarding the non-conditional probabilities (e.g.  $p(s : L)$  and  $p(s : S)$  above) 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.

# Semantic classification using conditional probabilities IV

- ▶ Such a classifier can be implemented in a number of different ways, e.g. as a neural network, as long as it outputs a probability distribution.
- ▶ The training of perceptual classifiers are outside the scope of this paper, but see Larsson (2013) and Fernández and Larsson (2014).

# Outline

## Introduction

## Background

- TTR fundamentals

- Probabilistic TTR fundamentals

- Bayesian nets and the Naive Bayes classifier

- A ProbTTR Naive Bayes classifier

- Semantic classification

## Semantic learning

## Word Meaning Negotiation and semantic updates

## Learning perceptual meanings from definitions

- Learning a new meaning from example

- Learning a new meaning from definition

- Learning new evidence values

- Learning new evidence variables

## Definitions vs. examples

## Conclusions & future work



# Frequentist semantic learning from examples I

- ▶ For the model of semantic classification that uses conditional probabilities, a central question is of course how to estimate conditional probabilities, of the form

$$p(C||E_1 \wedge \dots \wedge E_n) \text{ (where } C \in \mathfrak{R}(\mathbb{C}), E_i \in \mathfrak{R}(\mathbb{E}_i), 1 \leq i \leq n).$$

- ▶ Using Bayes rule and marginalising over the class value types, we get for a Naive Bayes classifier:

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

# Frequentist semantic learning from examples II

- ▶ For all combinations of evidence value types  $E_1, \dots, E_n$  and class value types  $C$ , we need
  - (a) the conditional probability of the evidence value types given the class value type,  $p(E_i || C)$ , and
  - (b) the prior of the class value type,  $p(C')$ .

# Computing conditional probabilities and priors I

- ▶ We compute likelihoods and probabilities as ratio of the frequencies of occurrences, summed over all judgements in the history:

$$p(E_i||C) = \frac{\sum_{j \in \mathcal{J}, j.\text{sit}=s} p(s : C)p(s : E_i)}{\sum_{j \in \mathcal{J}, j.\text{sit}=s} p(s : C)}$$

- ▶ This can be seen as a frequentist interpretation of epistemic probability.
- ▶ For the full account and motivation, see Larsson *et al.* (2021).

## Computing conditional probabilities and priors II

- ▶ In addition to conditional probabilities, we require the prior probabilities of the class value types  $C \in \mathfrak{R}(\mathbb{C})$ .
- ▶ We use  $p_{\mathfrak{J}}(T)$  to denote the prior probability that an arbitrary situation is of type  $T$  given  $\mathfrak{J}$ .

$$p_{\mathfrak{J}}(T) = \frac{\sum_{j \in \mathfrak{J}_T} j.\text{prob}}{P(\mathfrak{J})} \text{ if } P(\mathfrak{J}) > 0, \text{ otherwise } 0$$

- ▶ where  $P(\mathfrak{J})$  is the cardinality of situations in  $\mathfrak{J}$ , i.e. the total number of situations in  $\mathfrak{J}$ .

$$P(\mathfrak{J}) = |\{s | \exists j \in \mathfrak{J}, j.\text{sit} = s\}|$$

# Recording probabilities I

We can encode the relevant conditional probabilities and priors as a TTR record  $\pi(\kappa)$  containing conditional probabilities and priors for a classifier  $\kappa$ , where

- ▶ for  $1 \leq u \leq v$ ,  $E_u \in \mathfrak{R}(\mathbb{E}_1^\kappa) \cup \dots \cup \mathfrak{R}(\mathbb{E}_n^\kappa)$ , and
- ▶ where for  $1 \leq u \leq w$ ,  $C_u \in \mathfrak{R}(\mathbb{C}^\kappa)$

$$\pi(\kappa) = \left[ \begin{array}{l} \text{condps} \\ \text{priors} \end{array} = \left[ \begin{array}{l} \text{lbl}(C_1) = \left[ \begin{array}{l} \text{lbl}(E_1) = p(E_1||C_1) \\ \dots \\ \text{lbl}(E_v) = p(E_v||C_1) \end{array} \right] \\ \dots \\ \text{lbl}(C_w) = \left[ \begin{array}{l} \text{lbl}(E_1) = p(E_1||C_w) \\ \dots \\ \text{lbl}(E_v) = p(E_v||C_w) \end{array} \right] \\ \text{lbl}(C_1) = p_{\mathfrak{J}}(C_1) \\ \dots \\ \text{lbl}(C_w) = p_{\mathfrak{J}}(C_w) \end{array} \right] \right]$$

# Recording probabilities II

Parts of record containing conditional probabilities and priors for the fruit classifier:

$$\pi(\text{FruitC}) = \begin{bmatrix} \text{condps} \\ \text{priors} \end{bmatrix} = \begin{bmatrix} \text{lbl}(\text{Apple}) \\ \text{lbl}(\text{Pear}) \\ \text{lbl}(\text{Apple}) \\ \text{lbl}(\text{Pear}) \end{bmatrix} = \begin{bmatrix} \begin{bmatrix} \text{lbl}(\text{Red}) & = & 0.63 \\ \text{lbl}(\text{Green}) & = & 0.37 \\ \text{lbl}(\text{AShape}) & = & 0.97 \\ \text{lbl}(\text{PShape}) & = & 0.03 \end{bmatrix} \\ \dots \\ 0.64 \\ 0.26 \end{bmatrix}$$

## Recording probabilities III

Accordingly, we replace

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

with

$$\hat{p}^{\kappa}(C||E_1 \wedge \dots \wedge E_n) = \frac{\mathfrak{p}_{\mathfrak{J}}^{\kappa}(C)p^{\kappa}(E_1||C) \dots p^{\kappa}(E_n||C)}{\sum_{C' \in \mathfrak{R}(\mathbb{C}^{\kappa})} \mathfrak{p}_{\mathfrak{J}}^{\kappa}(C')p^{\kappa}(E_1||C') \dots p^{\kappa}(E_n||C')}$$

where

- ▶  $p^{\kappa}(E||C) = \pi(\kappa).\text{condps}.\text{lbl}(C).\text{lbl}(E)$
- ▶  $\mathfrak{p}_{\mathfrak{J}}^{\kappa}(C) = \pi(\kappa).\text{priors}.\text{lbl}(C)$

## Recording probabilities IV

- ▶ What this buys us is the possibility of updating classifiers by manipulating records encoding them.
- ▶ We will exploit this in formulating semantic updates resulting from word meaning negotiations.



# Outline

## Introduction

## Background

- TTR fundamentals

- Probabilistic TTR fundamentals

- Bayesian nets and the Naive Bayes classifier

- A ProbTTR Naive Bayes classifier

- Semantic classification

- Semantic learning

## Word Meaning Negotiation and semantic updates

## Learning perceptual meanings from definitions

- Learning a new meaning from example

- Learning a new meaning from definition

- Learning new evidence values

- Learning new evidence variables

## Definitions vs. examples

## Conclusions & future work

# Word Meaning Negotiation and semantic updates I

- ▶ In Myrendal (2015) and Myrendal (2019), a taxonomy for dialogue acts involved in WMNs of so-called *trigger words*  $T$  in online discussion forum communication is presented.
- ▶ Two central dialogue acts are:
  - ▶ Explicification: Provides an explicit (partial or complete) definition of  $T$ . We will here refer to this as simply *definition*.
  - ▶ Exemplification: Providing examples of what the trigger word can mean, or usually means.

# Word Meaning Negotiation and semantic updates II

- ▶ To describe the effects of these dialogue acts (once they are grounded), Larsson and Myrendal (2017) propose an abstract formalism for conceptual updates
- ▶ We assume that a definition  $D$  of a word (or expression)  $T$  has been provided, or an example situation  $E$ .
- ▶  $D$  or  $E$  is then used for updating the meaning in question.
  - ▶  $\delta(T, D)$ :  $T$  updated with  $D$  as a partial definition of  $T$
  - ▶  $\epsilon(T, E)$ :  $T$  updated with  $E$  as an example of a situation described by  $T$

# Word Meaning Negotiation and semantic updates III

- ▶ The abstract meaning update functions serve as a sort of API between dialogue acts and their consequent meaning updates.
- ▶ We can see the learning from examples described above as part of the specification of  $\epsilon(T, E)$ .

# Word Meaning Negotiation and semantic updates IV

- ▶ While we leave the exact formulation for future work, updating with an example  $E$  in the frequentist learning paradigm amounts to
  1. adding example  $E$  to  $\mathfrak{J}$
  2. recomputing the conditional probabilities and priors based on the updated  $\mathfrak{J}$
  3. updating the probabilities and priors in the classifier record.

# Word Meaning Negotiation and semantic updates V

- We collect all information about our fruit classifier in FruitR

$$\text{FruitR} = v(\text{FruitC}) \cup \pi(\text{FruitC}) =$$

$$\left[ \begin{array}{ll} \text{cvar} & = \text{Fruit} \\ \text{cval} & = \{\text{Apple}, \text{Pear}\} \\ \text{evars} & = \{\text{Col}, \text{Shp}\} \\ \text{evals} & = \left\{ \left[ \begin{array}{ll} \text{lbl}(\text{Col}) & = \{\text{Red}, \text{Green}\} \\ \text{lbl}(\text{Shp}) & = \{\text{AShape}, \text{PShape}\} \end{array} \right] \right\} \\ \text{condps} & = \left[ \begin{array}{ll} \text{lbl}(\text{Apple}) & = \left[ \begin{array}{ll} \text{lbl}(\text{Red}) & = 0.63 \\ \text{lbl}(\text{Green}) & = 0.37 \\ \text{lbl}(\text{AShape}) & = 0.97 \\ \text{lbl}(\text{PShape}) & = 0.03 \end{array} \right] \\ \text{lbl}(\text{Pear}) & = \dots \end{array} \right] \\ \text{priors} & = \left[ \begin{array}{ll} \text{lbl}(\text{Apple}) & = 0.64 \\ \text{lbl}(\text{Pear}) & = 0.26 \end{array} \right] \end{array} \right]$$

- This will allow us to update any aspect of the classifier by manipulating this record

# Word Meaning Negotiation and semantic updates VI

- ▶ In Larsson *et al.* (2021), we show how conditional probabilities and priors are recomputed after adding a new example  $E$  to  $\mathfrak{J}$
- ▶ We can store the result of such computations in a record  $P'$

$$P' = \left[ \begin{array}{l} \text{condps} \\ \text{priors} \end{array} = \left[ \begin{array}{l} \text{lbl}(\text{Apple}) \\ \text{lbl}(\text{Pear}) \\ \text{lbl}(\text{Apple}) \\ \text{lbl}(\text{Pear}) \end{array} = \left[ \begin{array}{l} \text{lbl}(\text{Red}) = 0.63 \\ \text{lbl}(\text{Green}) = 0.37 \\ \text{lbl}(\text{AShape}) = 0.97 \\ \text{lbl}(\text{PShape}) = 0.03 \\ \dots \\ 0.64 \\ 0.26 \end{array} \right] \right] \right]$$

# Word Meaning Negotiation and semantic updates VII

- ▶ Now, step (3) could be formalised thus:

$$\text{FruitR}' = \text{FruitR}[\text{P}']$$

- ▶ Simplifying somewhat, if  $r_1$  and  $r_2$  are records, then  $r_1[r_2]$  is the union of  $r_1$  and  $r_2$  except that if a label  $\ell$  occurs in both  $r_1$  and  $r_2$ , the value of  $\ell$  in  $r_1[r_2]$  will be  $r_2.\ell$ .

- ▶ For example,  $\left[ \begin{array}{lcl} a & = & x \\ b & = & y \end{array} \right] \left[ \begin{array}{lcl} b & = & z \\ c & = & w \end{array} \right] = \left[ \begin{array}{lcl} a & = & x \\ b & = & z \\ c & = & w \end{array} \right]$

- ▶ See Cooper (in prep) for details.



# Outline

## Introduction

## Background

- TTR fundamentals

- Probabilistic TTR fundamentals

- Bayesian nets and the Naive Bayes classifier

- A ProbTTR Naive Bayes classifier

- Semantic classification

- Semantic learning

## Word Meaning Negotiation and semantic updates

## Learning perceptual meanings from definitions

- Learning a new meaning from example

- Learning a new meaning from definition

- Learning new evidence values

- Learning new evidence variables

## Definitions vs. examples

## Conclusions & future work

# Learning perceptual meanings from definitions I

- ▶ The work reviewed above showed how probabilistic classifiers can be trained from examples presented in interaction.
- ▶ However, this cannot be the whole story.
- ▶ Indeed, in terms of the dialogue acts for semantic coordination presented in Larsson and Myrendal (2017), we have only covered exemplification.
- ▶ What about partial definitions (explicifications)?
- ▶ What effect do they have on agent's takes on meanings, and how is learning from definitions related to learning from examples?

## Learning perceptual meanings from definitions II

- ▶ From the perspective of agents learning how to classify situations probabilistically, one might ask how agents learn the structure of the Bayes net (or as a special case, Naive Bayes classifier) used to classify situations.
- ▶ We propose to connect these two questions, by exploring the idea that the dependency structure of Bayesian classifiers can be derived (learned) from definitions, and that one effect of a definition can be to update the structure of a Bayesian classifier.
- ▶ (We are not claiming that this is the only way agents can learn such structures.)
- ▶ In the fruit recognition game,  $B$  learns how to take shape (a-shape or p-shape) and colour (red or green) into account when classifying apples and pears, by adjusting conditional probabilities and priors.

# Outline

## Introduction

## Background

- TTR fundamentals

- Probabilistic TTR fundamentals

- Bayesian nets and the Naive Bayes classifier

- A ProbTTR Naive Bayes classifier

- Semantic classification

- Semantic learning

## Word Meaning Negotiation and semantic updates

## Learning perceptual meanings from definitions

- Learning a new meaning from example

- Learning a new meaning from definition

- Learning new evidence values

- Learning new evidence variables

## Definitions vs. examples

## Conclusions & future work

## Learning a new meaning from example I

- ▶ Before going into learning new meanings from definitions, it might be helpful to show how learning new meanings from examples (demonstrations) could be accounted for.
- ▶ In Larsson and Cooper (2009), it is shown how *ontological* meaning (e.g. that pandas are (not) bears) can be learned from interaction, and how such learning can be modelled in TTR.
- ▶ We can imagine a version of the fruit recognition game where new fruits (i.e., new value types for the fruit variable type) are introduced by demonstration:

A: What fruit is this?

B: A pear.

A: Wrong, it's a Wax Jambu.

B: Okay.



## Learning a new meaning from example II

- ▶ In this example,  $B$  can learn both that Wax Jambus are fruits, and what they look like based on being provided with an example Wax Jambu that they can observe.
- ▶ From the context,  $B$  can figure out that Wax Jambus are fruits.
- ▶ In the general case such an inference can be based on a variety of factors, including the ongoing activity and linguistic evidence.
- ▶ In terms of a probabilistic classifier, learning this amounts to adding a new value (type) to the Fruit variable (type).
- ▶ This update can be formalised thus:

$$\text{FruitR}' = \text{FruitR}[\text{cvals} = F.\text{cvals} \cup \{C\}]$$

- ▶ Furthermore,  $B$  could add the new example to  $\mathfrak{J}$  and update the conditional probabilities, as detailed above.

# Learning a new meaning from example III

$$\left[ \begin{array}{lcl}
 \text{cvar} & = & \text{Fruit} \\
 \text{cval} & = & \{\text{Apple}, \text{Pear}\} \\
 \text{evars} & = & \{\text{Col}, \text{Shp}\} \\
 \text{evals} & = & \left\{ \left[ \begin{array}{lcl} \text{lbl}(\text{Col}) & = & \{\text{Red}, \text{Green}\} \\ \text{lbl}(\text{Shp}) & = & \{\text{AShape}, \text{PShape}\} \end{array} \right] \right\} \\
 \text{condps} & = & \left[ \begin{array}{lcl} \text{lbl}(\text{Apple}) & = & \left[ \begin{array}{lcl} \text{lbl}(\text{Red}) & = & 0.63 \\ \text{lbl}(\text{Green}) & = & 0.37 \\ \text{lbl}(\text{AShape}) & = & 0.97 \\ \text{lbl}(\text{PShape}) & = & 0.03 \end{array} \right] \\ \text{lbl}(\text{Pear}) & = & \dots \end{array} \right] \\
 \text{priors} & = & \left[ \begin{array}{lcl} \text{lbl}(\text{Apple}) & = & 0.64 \\ \text{lbl}(\text{Pear}) & = & 0.26 \end{array} \right]
 \end{array} \right]$$

# Learning a new meaning from example IV

$$\left[ \begin{array}{lcl} \text{cvar} & = & \textit{Fruit} \\ \text{cval} & = & \{ \textit{Apple}, \textit{Pear}, \mathbf{WaxJambu} \} \\ \text{evars} & = & \{ \textit{Col}, \textit{Shp} \} \\ \text{evals} & = & \left\{ \left[ \begin{array}{lcl} \text{lbl}(\textit{Col}) & = & \{ \textit{Red}, \textit{Green} \} \\ \text{lbl}(\textit{Shp}) & = & \{ \textit{AShape}, \textit{PShape} \} \end{array} \right] \right\} \\ \text{condps} & = & \left[ \begin{array}{lcl} \text{lbl}(\textit{Apple}) & = & \left[ \begin{array}{lcl} \text{lbl}(\textit{Red}) & = & 0.63 \\ \text{lbl}(\textit{Green}) & = & 0.37 \\ \text{lbl}(\textit{AShape}) & = & 0.97 \\ \text{lbl}(\textit{PShape}) & = & 0.03 \end{array} \right] \\ \text{lbl}(\textit{Pear}) & = & \dots \\ \text{lbl}(\mathbf{WaxJambu}) & = & \dots \end{array} \right] \\ \text{priors} & = & \left[ \begin{array}{lcl} \text{lbl}(\textit{Apple}) & = & 0.64 \\ \text{lbl}(\textit{Pear}) & = & 0.26 \\ \text{lbl}(\mathbf{WaxJambu}) & = & \dots \end{array} \right] \end{array} \right]$$



## Learning a new meaning from example $V$

- ▶ Next, we will see how  $B$  can instead learn from partial definitions,
- ▶ Such definitions do not provide perceptually available evidence...
- ▶ ...but do seem to offer help in guiding  $B$ 's learning of the structure of the classifier, as well as associated probabilities.

# Outline

## Introduction

## Background

- TTR fundamentals

- Probabilistic TTR fundamentals

- Bayesian nets and the Naive Bayes classifier

- A ProbTTR Naive Bayes classifier

- Semantic classification

- Semantic learning

## Word Meaning Negotiation and semantic updates

## Learning perceptual meanings from definitions

- Learning a new meaning from example

- Learning a new meaning from definition

- Learning new evidence values

- Learning new evidence variables

## Definitions vs. examples

## Conclusions & future work

# Learning a new meaning from definition I

- ▶ We can imagine another language game where  $A$  asks  $B$  to fetch different fruits in a fruit storage, where several types of fruits are available, some of them unknown to  $B$ :

A: Get me an apple please

B: (fetches apple) there you go

A: Thanks. Now get me a Wax Jambu!

B: A Wax Jambu?

A: They are pear-shaped and red.

- ▶ We can call this the *fruit fetching game*.

## Learning a new meaning from definition II

- ▶ Let's assume that our learning agent  $B$  from  $A$ 's second utterance learns that Wax Jambus are fruits.
- ▶ However,  $B$  has not been presented with an example fruit to use for training.
- ▶ In this sense  $B$  does not yet know what Wax Jambus look like.
- ▶ It seems plausible that  $B$  in this case might be able to use  $A$ 's definition of Wax Jambu to distinguish Wax Jambus from other fruits (even if this ability will not be as developed as it might later be after seeing several Wax Jambus).

## Learning a new meaning from definition III

- ▶ How, then, could we model the effects of A's definition, which (with pronoun resolved) can be paraphrased as "Wax Jambus are pear-shaped and red"?
- ▶ Firstly, by adding a value type *WaxJambu* to the fruit classifier:

$$\text{FruitR}' = \text{FruitR}[\text{cvals} = \text{FruitR.cvals} \cup \{\text{WaxJambu}\}]$$

# Learning a new meaning from definition IV

- ▶ Secondly, by recomputing probabilities, assigning high values to  $p(Pshape||WaxJambu)$  and  $p(Red||WaxJambu)$ , and lowering other probabilities accordingly.
- ▶ For simplicity, we assume here that the high values are 1 and that conditional probabilities for other values of the same variables are lowered to 0.

$$FruitR'' = FruitR' \left[ \text{conds} = \left[ \text{lbl}(WaxJambu) = \begin{bmatrix} \text{lbl}(Red) = 1.0 \\ \text{lbl}(Green) = 0.0 \\ \text{lbl}(PShape) = 1.0 \\ \text{lbl}(AShape) = 0.0 \end{bmatrix} \right] \right]$$

# Learning a new meaning from definition V

$$\left[ \begin{array}{lcl}
 \text{cvar} & = & \text{Fruit} \\
 \text{cval} & = & \{\text{Apple}, \text{Pear}\} \\
 \text{evars} & = & \{\text{Col}, \text{Shp}\} \\
 \text{evals} & = & \left\{ \left[ \begin{array}{lcl} \text{lbl}(\text{Col}) & = & \{\text{Red}, \text{Green}\} \\ \text{lbl}(\text{Shp}) & = & \{\text{AShape}, \text{PShape}\} \end{array} \right] \right\} \\
 \text{condps} & = & \left[ \begin{array}{lcl} \text{lbl}(\text{Apple}) & = & \left[ \begin{array}{lcl} \text{lbl}(\text{Red}) & = & 0.63 \\ \text{lbl}(\text{Green}) & = & 0.37 \\ \text{lbl}(\text{AShape}) & = & 0.97 \\ \text{lbl}(\text{PShape}) & = & 0.03 \end{array} \right] \\ \text{lbl}(\text{Pear}) & = & \dots \end{array} \right] \\
 \text{priors} & = & \left[ \begin{array}{lcl} \text{lbl}(\text{Apple}) & = & 0.64 \\ \text{lbl}(\text{Pear}) & = & 0.26 \end{array} \right]
 \end{array} \right]$$

# Learning a new meaning from definition VI

$$\left[ \begin{array}{lcl}
 \text{cvar} & = & \text{Fruit} \\
 \text{cval} & = & \{\text{Apple}, \text{Pear}, \mathbf{WaxJambu}\} \\
 \text{evars} & = & \{\text{Col}, \text{Shp}\} \\
 \text{evals} & = & \left\{ \begin{array}{lcl}
 \text{lbl}(\text{Col}) & = & \{\text{Red}, \text{Green}\} \\
 \text{lbl}(\text{Shp}) & = & \{\text{AShape}, \text{PShape}\}
 \end{array} \right\} \\
 \\
 \text{condps} & = & \left[ \begin{array}{lcl}
 \text{lbl}(\text{Apple}) & = & \begin{bmatrix} \text{lbl}(\text{Red}) = 0.63 \\ \text{lbl}(\text{Green}) = 0.37 \\ \text{lbl}(\text{AShape}) = 0.97 \\ \text{lbl}(\text{PShape}) = 0.03 \end{bmatrix} \\
 \text{lbl}(\text{Pear}) & = & \dots \\
 \text{lbl}(\mathbf{WaxJambu}) & = & \begin{bmatrix} \text{lbl}(\text{Red}) = 1.0 \\ \text{lbl}(\text{Green}) = 0.0 \\ \text{lbl}(\text{PShape}) = 1.0 \\ \text{lbl}(\text{AShape}) = 0.0 \end{bmatrix}
 \end{array} \right] \\
 \\
 \text{priors} & = & \left[ \begin{array}{lcl}
 \text{lbl}(\text{Apple}) & = & 0.64 \\
 \text{lbl}(\text{Pear}) & = & 0.26 \\
 \text{lbl}(\mathbf{WaxJambu}) & = & \dots
 \end{array} \right]
 \end{array} \right]$$



## Learning a new meaning from definition VII

- ▶ Hence, the ProbTTR implementation of the  $\delta^+(T, D)$  function should be such that  $\delta^+(\llbracket \text{Wax Jambu} \rrbracket, \llbracket \text{pear-shaped and red} \rrbracket)$  results in these updates.
- ▶ Equipped with the updated mental fruit classifier,  $B$  now goes off to fetch a Wax Jambu in a storage room, despite never having seen one.
- ▶ One way of finding the right type of fruit in the storage is to simply going through the fruits in storage one by one and classify them, until one is classified as the sought type (here, Wax Jambu)<sup>2</sup>.

---

<sup>2</sup>One can imagine a continuation of the game, where  $B$  shows the retrieved fruit to  $A$  and receives feedback on whether it was right kind of fruit or not, and trains on this example in the normal way.

# Outline

## Introduction

## Background

- TTR fundamentals

- Probabilistic TTR fundamentals

- Bayesian nets and the Naive Bayes classifier

- A ProbTTR Naive Bayes classifier

- Semantic classification

- Semantic learning

## Word Meaning Negotiation and semantic updates

## Learning perceptual meanings from definitions

- Learning a new meaning from example

- Learning a new meaning from definition

- Learning new evidence values**

- Learning new evidence variables

## Definitions vs. examples

## Conclusions & future work

# Learning new evidence values I

- ▶ Above,  $A$ 's definition only included evidence values that were already used in the fruit classifier.
- ▶ However,  $A$  may also introduce a unknown value previously unknown to  $B$ :

A: Get me a Mango please!

B: A Mango?

A: They have an oblong shape.

## Learning new evidence values II

- ▶ In this case,  $B$  needs to both add the new class value *Mango* and a new evidence value *Oblong* (for the variable *Shp*):

$$\text{FruitR}' = \text{FruitR} [ \text{cvals} = \text{FruitR.cvals} \cup \{ \text{Mango} \} ]$$

$$\text{FruitR}'' = \text{FruitR}' [ \text{evals} = \left[ \begin{array}{l} \text{lbl}(\text{Shp}) = \text{FruitR.evals.lbl}(\text{Shp}) \\ \cup \{ \text{Oblong} \} \end{array} \right] ]$$

- ▶ Finally, as before, the conditional probabilities are shifted to favour the evidence variable given in the definition:

$$\text{FruitR}''' = \text{FruitR}'' [ \text{condps} = \left[ \begin{array}{l} \text{lbl}(\text{Mango}) = \left[ \begin{array}{l} \text{lbl}(\text{PShape}) = 0.0 \\ \text{lbl}(\text{AShape}) = 0.0 \\ \text{lbl}(\text{Oblong}) = 1.0 \end{array} \right] \end{array} \right] ]$$

# Learning new evidence values III

$$\left[ \begin{array}{lcl}
 \text{cvar} & = & \text{Fruit} \\
 \text{cval} & = & \{\text{Apple}, \text{Pear}\} \\
 \text{evars} & = & \{\text{Col}, \text{Shp}\} \\
 \text{evals} & = & \left\{ \left[ \begin{array}{lcl} \text{lbl}(\text{Col}) & = & \{\text{Red}, \text{Green}\} \\ \text{lbl}(\text{Shp}) & = & \{\text{AShape}, \text{PShape}\} \end{array} \right] \right\} \\
 \text{condps} & = & \left[ \begin{array}{lcl} \text{lbl}(\text{Apple}) & = & \left[ \begin{array}{lcl} \text{lbl}(\text{Red}) & = & 0.63 \\ \text{lbl}(\text{Green}) & = & 0.37 \\ \text{lbl}(\text{AShape}) & = & 0.97 \\ \text{lbl}(\text{PShape}) & = & 0.03 \end{array} \right] \\ \text{lbl}(\text{Pear}) & = & \dots \end{array} \right] \\
 \text{priors} & = & \left[ \begin{array}{lcl} \text{lbl}(\text{Apple}) & = & 0.64 \\ \text{lbl}(\text{Pear}) & = & 0.26 \end{array} \right]
 \end{array} \right]$$

# Learning new evidence values IV

$$\left[ \begin{array}{lcl}
 \text{cvar} & = & \text{Fruit} \\
 \text{cval} & = & \{\text{Apple}, \text{Pear}, \mathbf{Mango}\} \\
 \text{evars} & = & \{\text{Col}, \text{Shp}\} \\
 \text{evals} & = & \left\{ \begin{array}{lcl}
 \text{lbl}(\text{Col}) & = & \{\text{Red}, \text{Green}\} \\
 \text{lbl}(\text{Shp}) & = & \{\text{AShape}, \text{PShape}, \mathbf{Oblong}\}
 \end{array} \right\} \\
 \text{condps} & = & \left[ \begin{array}{lcl}
 \text{lbl}(\text{Apple}) & = & \left[ \begin{array}{lcl}
 \text{lbl}(\text{Red}) & = & 0.63 \\
 \text{lbl}(\text{Green}) & = & 0.37 \\
 \text{lbl}(\text{AShape}) & = & 0.97 \\
 \text{lbl}(\text{PShape}) & = & 0.03
 \end{array} \right] \\
 \text{lbl}(\text{Pear}) & = & \dots \\
 \text{lbl}(\mathbf{Mango}) & = & \left[ \begin{array}{lcl}
 \text{lbl}(\text{PShape}) & = & 0.0 \\
 \text{lbl}(\text{AShape}) & = & 0.0 \\
 \text{lbl}(\mathbf{Oblong}) & = & 1.0
 \end{array} \right]
 \end{array} \right] \\
 \text{priors} & = & \left[ \begin{array}{lcl}
 \text{lbl}(\text{Apple}) & = & 0.64 \\
 \text{lbl}(\text{Pear}) & = & 0.26 \\
 \text{lbl}(\mathbf{Mango}) & = & \dots
 \end{array} \right]
 \end{array} \right]$$

# Learning new evidence values $V$

- ▶ We assume here that  $B$  was familiar with the shape value type *Oblong*, but had not previously considered it relevant to fruit classification
- ▶ A more complicated situation arises when a previously unknown value for a known variable is introduced, e.g. a new shape.
- ▶ In such cases, perceptually available examples may be necessary to train the updated classifier on.

# Outline

## Introduction

## Background

- TTR fundamentals

- Probabilistic TTR fundamentals

- Bayesian nets and the Naive Bayes classifier

- A ProbTTR Naive Bayes classifier

- Semantic classification

- Semantic learning

## Word Meaning Negotiation and semantic updates

## Learning perceptual meanings from definitions

- Learning a new meaning from example

- Learning a new meaning from definition

- Learning new evidence values

- Learning new evidence variables

## Definitions vs. examples

## Conclusions & future work



# Learning new evidence variables I

- ▶ Finally, a definition may introduce a new evidence variable, along with a value:  
  
A: Get me a Kumqat!  
B: A Kumqat?  
A: They are small
  
- ▶ We assume that *B* is already has a *Size* classifier and knows that *Small* is a *Size* (along with, say, *MidSize* and *Large*).

# Learning new evidence variables II

- ▶ Given this, the resulting updates to  $B$ 's fruit classifier could be described thus:

$$\text{FruitR}' = \text{FruitR} [ \text{values} = \text{FruitR.cvals} \cup \{Kumqat\} ] ]$$

$$\text{FruitR}'' = \text{FruitR}' [ \text{evars} = \text{FruitR.evvars} \cup \{Size\} ] ]$$

$$\text{FruitR}''' = \text{FruitR}'' [ \text{evals} = [ \text{lbl}(Size) = \{Large, MidSize, Small\} ] ] ]$$

$$\text{FruitR}'''' = \text{FruitR}''' [ \text{condps} = \left[ \text{lbl}(Kumqat) = \begin{bmatrix} \text{lbl}(Small) = 1.0 \\ \text{lbl}(Large) = 0.0 \\ \text{lbl}(Midsize) = 0.0 \end{bmatrix} \right] ] ]$$

## Learning new evidence variables III

- ▶ This example also raises the question about partial definitions that only mention a value of one of the evidence variables.
- ▶ What should the conditional probabilities for a situation being of the value types for the other evidence variable types (not mentioned in the definition) given that the situation is of the new class value type?
- ▶ For now, we note that several options are available - assuming uniform distributions, or asking for more information ("What colour is a Kumqat? What shape?") and use the response to infer new conditional probabilities.

# Outline

## Introduction

## Background

- TTR fundamentals

- Probabilistic TTR fundamentals

- Bayesian nets and the Naive Bayes classifier

- A ProbTTR Naive Bayes classifier

- Semantic classification

- Semantic learning

## Word Meaning Negotiation and semantic updates

## Learning perceptual meanings from definitions

- Learning a new meaning from example

- Learning a new meaning from definition

- Learning new evidence values

- Learning new evidence variables

## Definitions vs. examples

## Conclusions & future work

# Definitions vs. examples I

- ▶ If we want to model how meanings are affected by both definitions and examples, we will need to say something about the trade-off between definitions and examples.
- ▶ For example, while a definition may be useful until examples have been observed, at some point the observed examples may override a definition.
- ▶ In the proposed account, definitions affect conditional probabilities only in the short run.

## Definitions vs. examples II

- ▶ Assuming conditional probabilities are recomputed when receiving new relevant observations, the probabilities resulting from proposed definitions (e.g. in the fruit fetching game) will be overwritten as soon as an observation of an instance of the defined concept has been made (an actual fruit of the defined type has been observed).
- ▶ This is perhaps not obviously wrong – it is at least theoretically possible that definitions are categorically superseded by observations – but a more flexible trade-off between definitions and examples (observations) would probably be desirable.

## Definitions vs. examples III

- ▶ There are ways of achieving this in the frequentist approach, e.g. by letting a definition lead to adding some relatively high number  $N$  of “fake” observations in line with the definition to  $\mathfrak{J}$ .
- ▶ By manipulation of  $N$ , the relative importance of definitions relative to observations can be regulated.
- ▶ If such approaches are deemed unsatisfying for theoretical or empirical reasons, it may be necessary to move to a different learning method.
- ▶ Future work thus includes working out alternative learning approaches that can better account for the trade-off between definitions and examples (see forthcoming talk at CLASP ReInAct workshop!)

# Outline

## Introduction

## Background

- TTR fundamentals

- Probabilistic TTR fundamentals

- Bayesian nets and the Naive Bayes classifier

- A ProbTTR Naive Bayes classifier

- Semantic classification

- Semantic learning

## Word Meaning Negotiation and semantic updates

## Learning perceptual meanings from definitions

- Learning a new meaning from example

- Learning a new meaning from definition

- Learning new evidence values

- Learning new evidence variables

## Definitions vs. examples

## Conclusions & future work



# Conclusion

- ▶ We have shown how (partial) definitions offered in word meaning negotiations can help learners structure probabilistic classifiers that are used to compute probabilistic semantic judgements.
- ▶ Technically, this was achieved by encoding a Naive Bayes classifier as a TTR record structure which can be updated by definitions.

# Future work

- ▶ Parsing natural language into an appropriate representation for updating classifiers
- ▶ Formulating a general update rule for carrying out such updates (of which several examples were given above)
- ▶ Generalising the account to Bayes nets (and other types of probabilistic classifiers)
- ▶ We also want to study actual definitions from human-human dialogues, rather than invented ones.

 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.

 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 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:](https://sites.google.com/site/typetheorywithrecords/drafts)

[//sites.google.com/site/typetheorywithrecords/drafts](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 and Cooper, Robin 2009.

Towards a formal view of corrective feedback.



In Alishahi, Afra; Poibeau, Thierry; and Villavicencio, Aline, editors 2009, *Proceedings of the Workshop on Cognitive Aspects of Computational Language Acquisition*. EACL. 1–9.

 Larsson, Staffan and Cooper, Robin 2021.

Bayesian classification and inference in a probabilistic type theory with records.

In *Proceedings of NALOMA 2021*.

 Larsson, Staffan and Myrendal, Jenny 2017.

Dialogue acts and updates for semantic coordination.

*SEMDIAL 2017 SaarDial* 59.

 Larsson, Staffan; Bernardy, Jean-Philippe; and Cooper, Robin 2021.

Semantic learning in a probabilistic type theory with records.

In *Proceedings of Workshop on Computing Semantics with Types, Frames and Related Structures 2021*.

 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 2013.

Formal semantics for perceptual classification.

*Journal of Logic and Computation*.



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.



Myrendal, Jenny 2015.

*Word Meaning Negotiation in Online Discussion Forum Communication.*

Ph.D. Dissertation, University of Gothenburg.



Myrendal, Jenny 2019.

Negotiating meanings online: Disagreements about word meaning in discussion forum communication.

*Discourse Studies* 21(3):317–339.



Pearl, J. 1990.

Bayesian decision methods.

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