# Bayesian Classification, Inference in a Probabilistic Type Theory with Records


**Staffan Larsson**
Centre for Linguistic Theory and Studies in Probability (CLASP)
Dept. of Philosophy, Linguistics and Theory of Science
University of Gothenburg
sl@ling.gu.se

**Robin Cooper**
Centre for Linguistic Theory and Studies in Probability (CLASP)
Dept. of Philosophy, Linguistics and Theory of Science
University of Gothenburg
cooper@ling.gu.se



## Abstract

We propose a probabilistic account of semantic inference and classification formulated in terms of probabilistic type theory with records, building on Cooper et al. (2014, 2015). We suggest probabilistic type theoretic formulations of Naive Bayes Classifiers and Bayesian Networks. A central element of these constructions is a type-theoretic version of a random variable. We illustrate this account with a simple language game combining probabilistic classification of perceptual input with probabilistic (semantic) inference.


## 1 Introduction

A probabilistic type theory was presented in Cooper et al. (2014) and Cooper et al. (2015), which extends Cooper's Type Theory with Records (TTR, Cooper, 2012; Cooper and Ginzburg, 2015; Cooper, in prep). Non-probabilistic TTR (in common with other type theories) works with judgements of the form $a : T$ ("$a$ is of type $T$") and assumes that such judgements are categorical. In probabilistic TTR (probTTR) we associate probabilities with judgements: $p(a : T)$ ("the probability that $a$ is of type $T$").

TTR has been used previously for natural language semantics (see, for example, Cooper, 2005 and Cooper, 2012), and to analyze semantic coordination and learning (for example, Larsson and Cooper, 2009; Cooper and Larsson, 2009). It has also been applied to the analysis of interaction in dialogue (for example, Ginzburg, 2012 and Breitholtz, 2020), and in modelling robotic states and spatial cognition (for example, Dobnik et al., 2013).

Two main considerations motivated recasting TTR in probabilistic terms. First, a probabilistic type theory offers a natural framework for capturing the gradience of semantic judgements. This allows it to serve as the basis for an account of vagueness in interpretation, as shown by Fernández and Larsson (2014). Second, such a theory lends itself to developing a model of semantic learning that can be straightforwardly integrated into more general probabilistic explanations of learning and inference.

Furthermore, we believe this will provide the foundation for a unified probabilistic account of natural language semantics that accounts for reasoning (logical as well as non-logical/enthymematic as in Breitholtz, 2020), learning (semantic and factual) and interaction, and that integrates low-level, subsymbolic real-valued perceptual information and high-level symbolic information (Larsson, 2015).

In this paper we suggest a way of incorporating a probabilistic inference and classification into ProbTTR. We do this because we believe that vagueness, learning, inference and classification are central rather than peripheral notions in semantics, and that probabilistic reasoning is central to all of them. Also, in contrast to an approach where e.g. classifiers are implemented outside the semantic theory, we want the reasoning underlying an agent's behaviour to be as transparent as possible to the agent itself (and thereby potentially also to its interlocutors).

To incorporate a probabilistic inference and classification into ProbTTR, we will need to introduce a ProbTTR version of a random variable, not discussed in Cooper et al. (2015). We will also show how probabilistic classification of perceptual evidence can be combined with probabilistic reasoning.

We first provide a brief overview of TTR and Probabilistic TTR. Section 4 provides some background on probabilistic inference and classification. Section 5 introduces conditional probabilities and defines a type theoretic version of a random variable. We use these variables to characterise a Naive Bayes classifier in Section 6. We illustrate Naive




*Proceedings of the 1st and 2nd Workshops on Natural Logic Meets Machine Learning (NALOMA)*, pages 51–59
June 16, 2021. ©2021 Association for Computational Linguistics

Bayes classification with an example of semantic classification. In Section 7 we show how probabilistic perception and reasoning can be combined in ProbTTR. We then introduce a ProbTTR characterisation of Bayesian Networks, and briefly discuss semantic learning. In Section 10 we present our conclusions and discuss directions for future work.

## 2 TTR: A brief introduction

We give a brief sketch of those aspects of TTR which we will use in this paper. For more detailed accounts see Cooper and Ginzburg (2015); Cooper (in prep).

$s : T$ represents a judgement that $s$ is of type $T$. A second kind of judgement (often written $T$true in Martin-Löf type theory) is the judgement that there is something of type $T$ ($T$ is non-empty).Types may be either *basic* or *complex* (in the sense that they are structured objects which have types or other objects introduced in the theory as components). One basic type that we will use is *Ind*, the type of individuals; another is *Real*, the type of real numbers. Among the complex types are *ptypes* which 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 situations, states or events in the world which instantiate the type. Thus $s$ : man($a$) can be glossed as "$s$ is a situation which shows (or proves) that $a$ is a man".

Another kind of complex type is *record types*. In TTR *records* are modelled as finite sets of fields. Each field is an ordered pair, $\langle \ell, o \rangle$, where $\ell$ is a *label* (drawn from a countably infinite stock of labels) and $o$ is an object which is a witness of some type. No two fields of a record can contain the same label. Importantly, $o$ can itself be a record. A *record type* is like a record except that the fields are of the form $\langle \ell, T \rangle$ where $\ell$ is a label as before and $T$ is a type. The basic intuition is that a record, $r$ is a witness for a record type, $T$, just in case for each field, $\langle \ell_i, T_i \rangle$, in $T$ there is a field, $\langle \ell_i, o_i \rangle$, in $r$ where $o_i : T_i$. (Note that this allows for the record to have additional fields with labels not included in the fields of the record type.) The types within fields in record types may *depend* on objects which can be found in the record which is being tested as a witness for the record type. We use a graphical display to represent both records and record types where each line represents a field. Example (1) represents the type of records which can be used to model situations where a man runs.

(1) $$\left[\begin{array}{lcl} \text{ref} & : & \textit{Ind} \\ \text{c}_{\text{man}} & : & \text{man(ref)} \\ \text{c}_{\text{run}} & : & \text{run(ref)} \end{array}\right]$$

A record of this type would be of the form

(2) $$\left[\begin{array}{lcl} \text{ref} & = & a \\ \text{c}_{\text{man}} & = & s \\ \text{c}_{\text{run}} & = & e \\ \ldots & & \end{array}\right]$$

where $a$ : *Ind*, $s$ : man($a$) and $e$ : run($a$).

We will introduce further details of TTR as we need them in subsequent sections.

## 3 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. The first is a judgement that a situation, $s$, is of type, $T$, with some probability. $p(s : T)$ is the probability that $s$ is a witness for $T$. The second is a judgement that there is some witness of type $T$. $p(T)$ is the probability that there is some witness for $T$. This introduces a distinction that is not normally made explicit in the notation used in probability theory.

It is useful to have type theoretic objects corresponding to judgements that a situation is of a type. Following terminology first introduced in Barwise (1989, Chap. 11), we call these *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 in (3).

(3) $$\left[\begin{array}{lcl} \text{sit} & : & \textit{Sit} \\ \text{sit-type} & : & \textit{Type} \\ \text{prob} & : & [0,1] \end{array}\right]$$

(where $[0, 1]$ represents the type of real numbers between 0 and 1). A probabilistic Austinian proposition $\varphi$ of this type corresponds to the judgement that $\varphi$.sit is of type $\varphi$.sit-type with probability $\varphi$.prob. That is,

(4) p($\varphi$.sit:$\varphi$.sit-type)= $\varphi$.prob

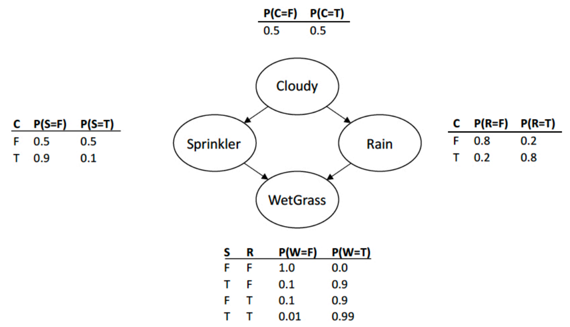


Figure 1: Example Bayesian Network

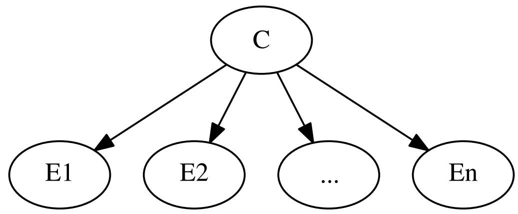


Figure 2: Naive Bayes classifier

## 4 Probabilistic Inference and Classification

A Bayesian Network is a Directed Acyclic Graph (DAG)[1]. The nodes of the DAG are random variables, each of whose values is the probability of one of the set of possible states that the variable denotes. Its directed edges express dependency relations among the variables. When the values of all the variables are specified, the graph describes a complete joint probability distribution (JPD) for its random variables (Pearl, 1990; Halpern, 2003).

Russell and Norvig (1995) give the example Bayesian Network in Figure 1. The only directly observable evidence is whether it is cloudy or not, and the queried variable is whether the grass is wet or not. We do not know if it is raining, or whether the sprinkler is on. Both of these factors depend on whether it is cloudy, and both affect the grass being wet.

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

(5) $p(W{=}\text{T}){=}\sum_{s,r,l} p(W{=}\text{T}, S{=}s, R{=}r, C{=}c)$

Here, $s$, $r$ and $l$ can be either T(rue) or F(alse).

The Bayesian network in Figure 1 allows us to simplify the computation of this JPD by encoding independence relations between variables, so that:

(6) $p(W, S, R, C){=}p(W|S, R)p(S|C)p(R|C)p(C)$

and hence

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

[1]This section briefly explains Bayesian nets and Naive Bayes classifiers, and introduces examples that will be used later. Readers familiar with this material can safely skip ahead to Section 5.

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. Figure 2 illustrates the relation between evidence variables and a class variable in a Naive Bayes classifier.

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

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

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:

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

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

Of course, if the assumption regarding the independence of the evidence variables does not hold, this estimation may be incorrect; this is the price to pay for the relative simplicity of the Naive Bayes classifier.

## 5 Type theoretic probabilistic inference and classification

We now turn to an account of probabilistic classification in ProbTTR. We first show how probabilistic inference can be modelled in ProbTTR. We then provide a Naive Bayes classifier with a detailed example. Finally, we generalise this account to Bayesian Networks.

### 5.1 Conditional probabilities in ProbTTR

We use $p(T_1||T_2)$ to represent the estimated[2] conditional probability that any situation, $s$, is of type

[2]Estimating $p(T_1||T_2)$ is part of the learning theory.

$T_1$ 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 $T_2$ , and $p(T_1|T_2)$, 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 $s : T_2$.

### 5.2 Random variables in TTR

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* $V$ whose range is a set of *value types* $\mathfrak{R}(V) = \{A_1, \ldots, A_n\}$ such that the following conditions hold.

(10) a. $A_j \sqsubseteq V$ 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 $s$, $p(s : V) \in \{0, 1.0\}$ and $p(s : V) = \sum_{T \in \mathfrak{R}(V)} p(s : T)$

(10 a) says that all value types for a variable type $V$ are subtypes of $V$. (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 $V = A_i \vee \ldots \vee A_n$. ($T_1 \vee T_2$ is the *join type* of $T_1$ and $T_2$. $a : T_1 \vee T_2$ just in case either $a : T_1$ or $a : T_2$). (10 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 $V$. (10 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 $V$ sum to 0; otherwise these probabilities sum to 1.0.

(10) encodes a conceptual difference between the probability that something has a property (such as colour, $p(s$:Colour$)$), and the probability that it has a certain value of a variable (e.g. $p(s$: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).

We can now formulate the example in Figure 1 in ProbTTR. We assume four binary variable types *Grass*, *Sprinkler*, *Raining* and *Cloudy* with corresponding variable value types as given in (11).

(11) $\mathfrak{R}$(*Grass*)=$\{$*GrassWet*, *GrassDry*$\}$
$\mathfrak{R}$(*Sprinkler*)=$\{$*SprinklerOn*, *SprinklerOff*$\}$
$\mathfrak{R}$(*Raining*)=$\{$*IsRaining*, *IsNotRaining*$\}$
$\mathfrak{R}$(*Cloudy*)=$\{$*ItIsCloudy*, *ItIsNotCloudy*$\}$

We specify that *Grass*=*GrassWet*$\vee$*GrassDry*, and similarly for the other variable types. This will ensure that *GrassWet*$\sqsubseteq$*Grass*, and similar subtyping constraints hold. Assuming that the variable types and variable value types are related as in (11) also entails that *GrassIsWet*$\perp$*GrassIsDry*, and similarly for the other variable value type pairs.

### 5.3 A ProbTTR Naive Bayes classifier

Corresponding to the evidence, class variables, and their values, we associate with a ProbTTR Naive Bayes classifier $\kappa$

(12) a. a collection of $m$ evidence variable types $\mathbb{E}^\kappa_1, \ldots, \mathbb{E}^\kappa_n$,

b. associated sets of evidence value types $\mathfrak{R}(\mathbb{E}^\kappa_1), \ldots, \mathfrak{R}(\mathbb{E}^\kappa_n)$,

c. a class variable type $\mathbb{C}^\kappa$, and

d. an associated set of class value types $\mathfrak{R}(\mathbb{C}^\kappa)$.

To classify a situation $s$ using a classifier $\kappa$, 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.

We define a ProbTTR Bayes classifier $\kappa$ as a function from a situation $s$ (of the meet type of the evidence variable types $\mathbb{E}^\kappa_1, \ldots, \mathbb{E}^\kappa_n$) 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 $\mathbb{E}^\kappa_1, \ldots, \mathbb{E}^\kappa_n$. Formally, a ProbTTR Naïve Bayes classifier is a function $\kappa$ of the type

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

such that if[3] $s : \mathbb{E}_1^\kappa \wedge \ldots \wedge \mathbb{E}_n^\kappa$, then

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

where

(15) $p^\kappa(s : C) =$

$$\sum_{\substack{E_1 \in \mathfrak{R}(E_1^\kappa)\\ \ldots\\ E_n \in \mathfrak{R}(E_n^\kappa)}} p^\kappa(C||E_1 \wedge \ldots \wedge E_n) p(s : E_1) \ldots p(s : E_n)$$

($T_1 \wedge T_2$ is the *meet type* of $T_1$ and $T_2$. $a : T_1 \wedge T_2$ just in case $a : T_1$ and $a : T_2$.)

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.

Note that the probabilities associated with the evidence are probabilities that the situation $s$ (the situation being classified) is of the various evidence value types. We do not assume that the evidence variables are known, only that we have a probability distribution over judgements of $s$ being of the associated evidence value types. We also do not use the priors of the evidence value types here, as that would give us the marginal probability of *any* situation being of the class value type $C$, rather than the situation $s$ being classified. Our ProbTTR notation allows us to make this distinction clear.

As above in (9), 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:

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

## 6 Semantic Classification: Example

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 an *Apple 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. (This paper describes only the classification step, leaving the learning step for future work.)

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

(17) a. $\textit{Apple} = \left[\begin{array}{lll}\text{x} & : & \textit{Ind}\\ \text{c}_{\text{apple}} & : & \text{apple(x)}\end{array}\right]$

b. $\textit{Pear} = \left[\begin{array}{lll}\text{x} & : & \textit{Ind}\\ \text{c}_{\text{pear}} & : & \text{pear(x)}\end{array}\right]$

We take it that the probability of judgements that something is of type *Ind* is always 1.0, and that

(18) $$p(s : \left[\begin{array}{lll}\text{x} & : & \textit{Ind}\\ \text{c} & : & T\text{(x)}\end{array}\right]) = p(s.c : T(s.x))$$

so that e.g. if

(19) $$s = \left[\begin{array}{lll}\text{x} & = & \text{a}\\ \text{c} & = & \text{prf}\end{array}\right],$$

then

(20) $p(s : \textit{Apple}) = p(\text{prf} : \text{apple(a)})$

.

Furthermore, we will assume that the objects in the Apple Recognition Game have one of two shapes (a-shape or p-shape) and one of two colours (green or red). We define shorthands for the record types involved.

(21) a. $\textit{Ashape} = \left[\begin{array}{lll}\text{x} & : & \textit{Ind}\\ \text{c} & : & \text{ashape(x)}\end{array}\right]$

b. $\textit{Pshape} = \left[\begin{array}{lll}\text{x} & : & \textit{Ind}\\ \text{c} & : & \text{pshape(x)}\end{array}\right]$

[3] Recall that that $\mathbb{E}_1^\kappa \ldots \mathbb{E}_n^\kappa$ are variable types and that for any variable type $V$ and situation $s$, $p(s : V) \in \{0, 1.0\}$. Therefore, any type judgement regarding a variable type, such as that involved in the classifier function, can be regarded as categorical.

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

d. $Red = \begin{bmatrix} \text{x} & : & \textit{Ind} \\ \text{c} & : & \text{red(x)} \end{bmatrix}$

The class variable type is *Fruit*, with value types $\Re(\textit{Fruit}) = \{\textit{Apple}, \textit{Pear}\}$. The evidence variable types are (i) *Col*(our), with value types $\Re(\textit{Col}) = \{\textit{Green}, \textit{Red}\}$, and (ii) Shape, with value types $\Re(\textit{Shape}) = \{\textit{Ashape}, \textit{Pshape}\}$. Figure 3 shows the evidence and class types of the Apple Recognition Game in a simple Bayesian Network.

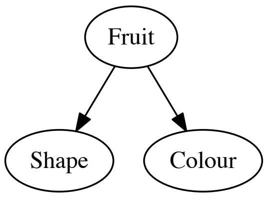


Figure 3: Bayesian Network for the Apple Recognition Game

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*.

(22) FruitC($s$) =

$$\{ \begin{bmatrix} \text{sit} = s \\ \text{sit-type} = F \\ \text{prob} = p^{\text{FruitC}}(s : F) \end{bmatrix} \mid F \in \Re(\textit{Fruit})\}$$

We compute the probability of a classification in the Apple Recognition Game with the equation in (23), which is a special case of (15).

(23) for each $F \in \Re(\textit{Fruit})$, $p^{\textit{FruitC}}(s : F) =$

$$\sum_{\substack{L \in \Re(\textit{Col}) \\ S \in \Re(\textit{Shape})}} p(F||L \wedge S)p(s : L)p(s : S)$$

Therefore, to determine the probability that a situation is of the apple type, we sum over the various evidence type values for apple.

(24) $p^{\text{FruitC}}(s{:}\textit{Apple}) =$
$p(\textit{Apple}||\textit{Green} \wedge \textit{Ashape})p(s{:}\textit{Green})p(s{:}\textit{Ashape}) +$
$p(\textit{Apple}||\textit{Green} \wedge \textit{Pshape})p(s{:}\textit{Green})p(s{:}\textit{Pshape}) +$
$p(\textit{Apple}||\textit{Red} \wedge \textit{Ashape})p(s{:}\textit{Red})p(s{:}\textit{Ashape}) +$
$p(\textit{Apple}||\textit{Red} \wedge \textit{Pshape})p(s{:}\textit{Red})p(s{:}\textit{Pshape})$

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 in (25) illustrate a JPD for the Bayesian Network in Figure 3.

(25)

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

For each square with *Apple*/*Pear* type values, the conditional probabilities of the fruit being an apple and of its being a pear sum to 1. These probabilities are estimated using (16). For example:

(26) $\hat{p}(\textit{Apple}||\textit{Green} \wedge \textit{Ashape}) =$

$$\frac{p(\textit{Apple})p(\textit{Green}||\textit{Apple})p(\textit{Ashape}||\textit{Apple})}{\sum_{F' \in \{\textit{Apple},\textit{Pear}\}} p(F')p(\textit{Green}||F')p(\textit{Ashape}||F')}$$

The non-conditional probabilities in (24) are derived from the agents' take on the particular situation being classified; let us call it $s_5$.

(27)

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

With these numbers in place, we can compute the probability that the fruit shown in $s_5$is an apple:

(28) $p^{\text{FruitC}}(s_5\text{: } \textit{Apple}) =$
$0.93 * 0.80 * 0.90 + 0.63 * 0.80 * 0.10 +$
$0.56 * 0.20 * 0.90 + 0.13 * 0.20 * 0.10 =$
$0.67 + 0.05 + 0.10 + 0.00 =$
$0.82$

In this section, we have shown how a Naive Bayes classifier, taking as input [1] judgements about how a situation $s$ is classified with respect to a set of evidence value types, [2] conditional probabilities of some situation being of an evidence value type given that it is of a class value type, can be cast in ProbTTR.

## 7 Perceiving evidence

We might at this point ask, where do the non-conditional probabilities concerning the situation $s$ being classified (exemplified in 27) 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 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. Adapting the notion of a probabilistic TTR classifier to the current setting, a probabilistic perceptual (here, visual) classifier corresponding to an evidence value type $E_i (1 \leq i \leq n)$ 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 $\mathfrak{R}(E_i^\kappa)$, encoded as a set of probabilistic Austinian propositions:

(29) $\pi_{E_i^\kappa}$:*Sit*$_\mathfrak{V} \rightarrow$
$$\{\begin{bmatrix}\text{sit} : \textit{Sit}_\mathfrak{V} \\ \text{sit-type} : \textit{RecType}_R \\ \text{prob} : [0,1]\end{bmatrix} \mid R \in \mathfrak{R}(E_i^\kappa)\}$$

where *Sit*$_\mathfrak{V}$ is the type of situations where perception of some object (labelled x) yields visual information (labelled c) concerning x:

(30) $$\textit{Sit}_\mathfrak{V} = \begin{bmatrix} \text{x} & : & \textit{Ind} \\ \text{c} & : & \mathfrak{V} \end{bmatrix}$$

and where *RecType*$_R$ is the (singleton) type of record types identical to $R$, so that e.g.

(31) $T$:*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.

(32) $$\pi_{Col} : \begin{bmatrix} \text{x} & : & \textit{Ind} \\ \text{c} & : & \mathfrak{V} \end{bmatrix} \rightarrow$$
$$\{\begin{bmatrix}\text{sit} : \textit{Sit}_\mathfrak{V} \\ \text{sit-type} : \textit{RecType}_{Green} \\ \text{prob} : [0,1]\end{bmatrix}, \begin{bmatrix}\text{sit} : \textit{Sit}_\mathfrak{V} \\ \text{sit-type} : \textit{RecType}_{Red} \\ \text{prob} : [0,1]\end{bmatrix}\}$$

If we take $s_5$ to be e.g.

(33) $$\begin{bmatrix} \text{x} & = & \text{a}_{453} \\ \text{c} & = & \text{Img}_{9876} \end{bmatrix}$$

where

(34) a. $\text{a}_{453}$:*Ind*

b. $\text{Img}_{9876}$:$\mathfrak{V}$

and we assume that

(35) $\pi_{Col}(s_5)$=
$$\{\begin{bmatrix}\text{sit} = s_5 \\ \text{sit-type} = \textit{Green} \\ \text{prob} = 0.8\end{bmatrix}, \begin{bmatrix}\text{sit} = s_5 \\ \text{sit-type} = \textit{Red} \\ \text{prob} = 0.2\end{bmatrix}\}$$

then (4) yields that

(36) a. p($s_5$:*Green*)=0.8

b. p($s_5$:*Red*)=0.2

which, incidentally, are the probabilities also shown in (27). This illustrates how ProbTTR allows combining probabilistic perceptual classification and probabilistic reasoning.

## 8 Bayesian networks in TTR

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.

Given this, a TTR Bayes net classifier is a function $\kappa$ of type

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

such that if $s : \mathbb{E}_1^\kappa \wedge \ldots \wedge \mathbb{E}_n^\kappa$ and $1 \leq j \leq m$, then

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

$$p^{\kappa}(s:I_j) = \sum_{\substack{I_1 \in \mathfrak{R}(\mathbb{I}_1^{\kappa}) \\ \vdots \\ I_{j-1} \in \mathfrak{R}(\mathbb{I}_{j-1}^{\kappa}) \\ I_{j+1} \in \mathfrak{R}(\mathbb{I}_{j+1}^{\kappa}) \\ \vdots \\ I_m \in \mathfrak{R}(\mathbb{I}_m^{\kappa}) \\ E_1 \in \mathfrak{R}(\mathbb{E}_1^{\kappa}) \\ \vdots \\ E_n \in \mathfrak{R}(\mathbb{E}_n^{\kappa})}} p(I_j || I_1 \wedge \ldots \wedge I_{j-1} \wedge I_{j+1} \wedge \ldots \wedge I_m \wedge E_1 \wedge \ldots \wedge E_n) p(s:E_1) \ldots p(s:E_n)$$

Figure 4: A TTR Bayes net classifier

where $p^{\kappa}(s:I_j)$ is as in Figure 4.

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*. 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 in Figure 1 entail that

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

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

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

(40) $p^{\kappa}(s:G) =$

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

## 9 Semantic learning

A central question is, of course, how we get the conditional and prior probabilities used for classification. This is the role of the semantic learning component. For a ProbTTR classifier, the learning component needs to estimate the probabilities required for computing $p(C||E_1 \wedge \ldots \wedge E_n)$.

In Cooper et al. (2015) a solution is sketched, based on the idea that an agent makes judgements based on a finite string of probabilistic Austinian propositions, the *judgement history* $\mathfrak{J}$. When an agent $A$ encounters a new situation $s$ and wants to know if it is of type $T$ or not, $A$ uses probabilistic reasoning to determine $p(s:T)$ on the basis of $A$'s previous judgements $\mathfrak{J}$. We are currently working on casting a couple of learning theories in ProbTTR, and this will be reported in future work.

## 10 Conclusions

Cooper et al. (2014) and Cooper et al. (2015) presented a probabilistic formulation of a rich type theory with records, and used it as the foundation for a compositional semantics in which a probabilistic judgement that a situation is of a certain type plays a central role. The basic types and type judgements at the foundation of the type system correspond to perceptual judgements concerning objects and events in the world, rather than to entities in a model, and set theoretic constructions defined on them. This approach grounds meaning in observational judgements concerning the likelihood of situations holding in the world.

Here, we have proposed a Bayesian account of semantic classification and inference formulated in terms of probabilistic type theory. We have suggested probabilistic type theoretic formulations of Naive Bayes Classifiers and Bayesian Networks. A central element of these constructions is a ProbTTR version of a random variable.

Future work includes applying Bayesian inference and classification in ProbTTR to a variety of problems in natural language semantics, including vagueness (where some initial steps are taken in Fernández and Larsson (2014)), probabilistic reasoning in dialogue, and learning grounded meanings from interaction (along the lines of Larsson (2013)). We will also implement this integrated system in order to demonstrate its viability as a computational model of natural language learning, reasoning and interaction.


## Acknowledgments

This work was supported by Swedish Research Council grants 2014-39, Centre for Linguistic Theory and Studies in Probability (CLASP) at the University of Gothenburg, and 2016-01162, Incremental Reasoning in Dialogue.

## References


Jon Barwise. 1989. *The Situation in Logic*. CSLI Publications, Stanford.

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

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

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

Robin Cooper. in 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`.

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

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

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

Robin Cooper and Staffan Larsson. 2009. Compositional and ontological semantics in learning from corrective feedback and explicit definition. In *Proceedings of DiaHolmia: 2009 Workshop on the Semantics and Pragmatics of Dialogue*, pages 59–66. Department of Speech, Music and Hearing, KTH.

Simon Dobnik, Robin Cooper, and Staffan Larsson. 2013. Modelling language, action, and perception in Type Theory with Records. In Denys Duchier and Yannick Parmentier, editors, *Constraint Solving and Language Processing - 7th International Workshop on Constraint Solving and Language Processing, CSLP 2012, Orleans, France, September 13-14, 2012. Revised Selected Papers*, number 8114 in Publications on Logic, Language and Information (FoLLI). Springer, Berlin, Heidelberg.

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

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

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

Staffan Larsson. 2013. Formal semantics for perceptual classification. *Journal of Logic and Computation*.

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

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

Staffan Larsson and Robin Cooper. 2009. Towards a formal view of corrective feedback. In *Proceedings of the Workshop on Cognitive Aspects of Computational Language Acquisition*, pages 1–9. EACL.

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

Stuart Russell and Peter Norvig. 1995. *Artificial Intelligence: A Modern Approach*. Prentice Hall Series in Artificial Intelligence. Englewood Cliffs, New Jersey.