

# A Tutorial on the Expectation-Maximization Algorithm Including Maximum-Likelihood Estimation and EM Training of Probabilistic Context-Free Grammars

**Detlef Prescher**

Language and Inference Technology Lab  
Institute for Logic, Language and Computation  
University of Amsterdam  
Nieuwe Achtergracht 166  
1018 WV Amsterdam  
The Netherlands  
`prescher@science.uva.nl`

October 2, 2003

Comments are welcome!

## 1 Introduction

The paper gives a brief review of the expectation-maximization algorithm (Dempster, Laird, and Rubin 1977) in the comprehensible framework of discrete mathematics. In Section 2, two prominent estimation methods, the relative-frequency estimation and the maximum-likelihood estimation are presented. Section 3 is dedicated to the expectation-maximization algorithm and a simpler variant, the generalized expectation-maximization algorithm. In Section 4, two loaded dice are rolled. A more interesting example is presented in Section 5: The estimation of probabilistic context-free grammars. Enjoy!

## 2 Estimation Methods

A statistics problem is a problem in which a **corpus**<sup>1</sup> that has been generated in accordance with some unknown **probability distribution** must be analyzed and some type of inference about the unknown distribution must be made. In other words, in a statistics problem there is a choice between two or more probability distributions which might have generated the corpus. In practice, there are often an infinite number of different possible distributions – statisticians bundle these into one single **probability model** – which might have generated the corpus. By analyzing the corpus, an attempt is made to learn about the unknown distribution. So, on the basis of the corpus, an **estimation method** selects one **instance** of the probability model, thereby aiming at finding the original distribution. In this section, two common estimation methods, the relative-frequency and the maximum-likelihood estimation, are presented.

---

<sup>1</sup>Statisticians use the term *sample* but computational linguists prefer the term *corpus*

## Corpora

**Definition 1** Let  $\mathcal{X}$  be a countable set. A real-valued function  $f: \mathcal{X} \rightarrow \mathcal{R}$  is called a **corpus**, if  $f$ 's values are non-negative numbers

$$f(x) \geq 0 \quad \text{for all } x \in \mathcal{X}$$

Each  $x \in \mathcal{X}$  is called a **type**, and each value of  $f$  is called a **type frequency**. The **corpus size**<sup>2</sup> is defined as

$$|f| = \sum_{x \in \mathcal{X}} f(x)$$

Finally, a corpus is called **non-empty** and **finite** if

$$0 < |f| < \infty$$

In this definition, type frequencies are defined as *non-negative real numbers*. The reason for not taking *natural numbers* is that some statistical estimation methods define type frequencies as *weighted occurrence frequencies* (which are not natural but non-negative real numbers). Later on, in the context of the EM algorithm, this point will become clear. The following definition shows that Definition 1 covers the standard notion of the term *corpus* (used in Computational Linguistics) and of the term *sample* (used in Statistics).

**Definition 2** Let  $x_1, \dots, x_n$  be a finite sequence of type instances from  $\mathcal{X}$ . Each  $x_i$  of this sequence is called a **token**. The **occurrence frequency** of a type  $x$  in the sequence is defined as the following count

$$f(x) = |\{i \mid x_i = x\}|$$

Obviously,  $f$  is a corpus in the sense of Definition 1, and it has the following properties: The type  $x$  does not occur in the sequence if  $f(x) = 0$ ; In any other case there are  $f(x)$  tokens in the sequence which are identical to  $x$ . Moreover, the corpus size  $|f|$  is identical to  $n$ , the number of tokens in the sequence.

## Relative-Frequency Estimation

Let us first present the notion of probability that we use throughout this paper.

**Definition 3** Let  $\mathcal{X}$  be a countable set of types. A real-valued function  $p: \mathcal{X} \rightarrow \mathcal{R}$  is called a **probability distribution on  $\mathcal{X}$** , if  $p$  has two properties: First,  $p$ 's values are non-negative numbers

$$p(x) \geq 0 \quad \text{for all } x \in \mathcal{X}$$

and second,  $p$ 's values sum to 1

$$\sum_{x \in \mathcal{X}} p(x) = 1$$

---

<sup>2</sup>Note that the corpus size  $|f|$  is well-defined: The order of summation is not relevant for the value of the (possible infinite) series  $\sum_{x \in \mathcal{X}} f(x)$ , since the types are countable and the type frequencies are non-negative numbers

Readers familiar to probability theory will certainly note that we use the term *probability distribution* in a sloppy way (on page 611, Duda et al. (2001) introduce the term *probability mass function* instead). Standardly (see e.g. DeGroot (1989)), probability distributions allocate a probability value  $p(A)$  to subsets  $A \subseteq \mathcal{X}$ , so-called **events** of an **event space**  $\mathcal{X}$ , such that three specific axioms are satisfied:

**Axiom 1**  $p(A) \geq 0$  for any event  $A$ .

**Axiom 2**  $p(\mathcal{X}) = 1$ .

**Axiom 3**  $p(\bigcup_{i=1}^{\infty} A_i) = \sum_{i=1}^{\infty} p(A_i)$  for any infinite sequence of disjoint events  $A_1, A_2, A_3, \dots$

Now, however, note that the probability distributions introduced in Definition 3 induce rather naturally the following probabilities for events  $A \subseteq \mathcal{X}$

$$p(A) := \sum_{x \in A} p(x)$$

Using the properties of  $p(x)$ , we can easily show that the probabilities  $p(A)$  satisfy the three axioms of probability theory. So, Definition 3 is justified and thus, for the rest of the paper, we are allowed to put axiomatic probability theory out of our minds.

**Definition 4** Let  $f$  be a non-empty and finite corpus. The probability distribution

$$\tilde{p}: \mathcal{X} \rightarrow [0, 1] \quad \text{where} \quad \tilde{p}(x) = \frac{f(x)}{|f|}$$

is called the **relative-frequency estimate** on  $f$ .

The relative-frequency estimation is the most comprehensible estimation method and has some nice properties which will be discussed in the context of the more general maximum-likelihood estimation. For now, however, note that it is very easy to check that  $\tilde{p}$ 's values sum to one:  $\sum_{x \in \mathcal{X}} \tilde{p}(x) = \sum_{x \in \mathcal{X}} |f|^{-1} \cdot f(x) = |f|^{-1} \cdot \sum_{x \in \mathcal{X}} f(x) = |f|^{-1} \cdot |f| = 1$ .

## Maximum-Likelihood Estimation

**Definition 5** A non-empty set  $\mathcal{M}$  of probability distributions on a set  $\mathcal{X}$  of types is called a **probability model on  $\mathcal{X}$** , whereas the elements of  $\mathcal{M}$  are called **instances of the model  $\mathcal{M}$** . The **unrestricted probability model** is the set  $\mathcal{M}(\mathcal{X})$  of all probability distributions on the set of types

$$\mathcal{M}(\mathcal{X}) = \left\{ p: \mathcal{X} \rightarrow [0, 1] \mid \sum_{x \in \mathcal{X}} p(x) = 1 \right\}$$

A probability model  $\mathcal{M}$  is called **restricted** in all other cases

$$\mathcal{M} \subseteq \mathcal{M}(\mathcal{X}) \quad \text{and} \quad \mathcal{M} \neq \mathcal{M}(\mathcal{X})$$

In practice, most probability models are restricted since their instances are often defined on a set  $\mathcal{X}$  comprising multi-dimensional types such that certain parts of the types are statistically independent (see examples 3 and 4).

**Definition 6** Let  $f$  be a non-empty and finite corpus on a countable set  $\mathcal{X}$  of types. Let  $\mathcal{M}$  be a probability model on  $\mathcal{X}$ . The **probability of the corpus** allocated by an instance  $p$  of the model  $\mathcal{M}$  is defined as

$$L(f; p) = \prod_{x \in \mathcal{X}, f(x) \neq 0} p(x)^{f(x)}$$

An instance  $\hat{p}$  of the model  $\mathcal{M}$  is called a **maximum-likelihood estimate of  $\mathcal{M}$  on  $f$** , if and only if the corpus  $f$  is allocated a maximum probability by  $\hat{p}$

$$L(f; \hat{p}) = \max_{p \in \mathcal{M}} L(f; p)$$

By looking at this definition, we recognize that maximum-likelihood estimates are the solutions of a quite complex optimization problem. So, some nasty questions about maximum-likelihood estimation arise:

**Existence** Is there for any probability model and any corpus a maximum-likelihood estimate of the model on the corpus?

**Uniqueness** Is there for any probability model and any corpus a unique maximum-likelihood estimate of the model on the corpus?

**Computability** For which probability models and corpora can maximum-likelihood estimates be efficiently computed?

For some probability models  $\mathcal{M}$ , the following theorem gives a positive answer.

**Theorem 1** Let  $f$  be a non-empty and finite corpus on a countable set  $\mathcal{X}$  of types. Then:

- (i) The relative-frequency estimate  $\tilde{p}$  is a unique maximum-likelihood estimate of the unrestricted probability model  $\mathcal{M}(\mathcal{X})$  on  $f$ .
- (ii) The relative-frequency estimate  $\tilde{p}$  is a unique maximum-likelihood estimate of a (restricted or unrestricted) probability model  $\mathcal{M}$  on  $f$ , if and only if  $\tilde{p}$  is an instance of the model  $\mathcal{M}$ .

On a first glance, proposition (ii) seems to be more general than proposition (i), since proposition (i) is about one single probability model, the unrestricted model, whereas proposition (ii) gives some insight about the relation of the relative-frequency estimate to a maximum-likelihood estimate of arbitrary restricted probability models (which are the most common probability models in practice). Both propositions, however, are equivalent. Moreover, proposition (i) is equivalent to the famous information inequality of information theory, for which various proofs have been given in the literature.

## The Information Inequality of Information Theory

**Definition 7** The **relative entropy**  $D(p \parallel q)$  of the probability distribution  $p$  with respect to the probability distribution  $q$  is defined by

$$D(p \parallel q) = \sum_{x \in \mathcal{X}} p(x) \log \frac{p(x)}{q(x)}$$

(Based on continuity arguments, we use the convention that  $0 \log \frac{0}{q} = 0$  and  $p \log \frac{p}{0} = \infty$  and  $0 \log \frac{0}{0} = 0$ . The logarithm is calculated with respect to the base 2.)

Connecting maximum-likelihood estimation with the concept of relative entropy, the following theorem gives the important insight that the relative-entropy of the relative-frequency estimate is minimal with respect to a maximum-likelihood estimate.

**Theorem 2** *Let  $\tilde{p}$  be the relative-frequency estimate on a non-empty and finite corpus  $f$ , and let  $\mathcal{M}$  be a probability model on the set  $\mathcal{X}$  of types. Then: An instance  $\hat{p}$  of the model  $\mathcal{M}$  is a maximum-likelihood estimate of  $\mathcal{M}$  on  $f$ , if and only if the relative-entropy of  $\tilde{p}$  is minimal with respect to  $\hat{p}$*

$$D(\tilde{p} \parallel \hat{p}) = \min_{p \in \mathcal{M}} D(\tilde{p} \parallel p)$$

**Proof** First, the relative entropy  $D(\tilde{p} \parallel p)$  is simply the difference of two further entropy values, the so-called **cross-entropy**  $H(\tilde{p}; p) = -\sum_{x \in X} \tilde{p}(x) \log p(x)$  and the **entropy**  $H(\tilde{p}) = -\sum_{x \in X} \tilde{p}(x) \log \tilde{p}(x)$  of the relative-frequency estimate

$$D(\tilde{p} \parallel p) = H(\tilde{p}; p) - H(\tilde{p})$$

(Based on continuity arguments and in full agreement with the convention used in Definition 7, we use here that  $\tilde{p} \log 0 = -\infty$  and  $0 \log 0 = 0$ .) It follows that minimizing the relative entropy is equivalent to minimizing the cross-entropy (as a function of the instances  $p$  of the given probability model  $\mathcal{M}$ ). The cross-entropy, however, is proportional to the negative log-probability of the corpus  $f$

$$H(\tilde{p}; p) = -\frac{1}{|f|} \log L(f; p)$$

So, finally, minimizing the relative entropy  $D(\tilde{p} \parallel p)$  is equivalent to maximizing the corpus probability  $L(f; p)$ .<sup>3</sup>

Together with Theorem 2, the following theorem, the so-called **information inequality of information theory**, proves Theorem 1. The information inequality states simply that the relative entropy is a non-negative number – which is zero, if and only if the two probability distributions are equal.

**Theorem 3** *Let  $p$  and  $q$  be two probability distributions. Then*

$$D(p \parallel q) \geq 0$$

*with equality if and only if  $p(x) = q(x)$  for all  $x \in \mathcal{X}$ .*

**Proof** See, e.g., Cover and Thomas (1991), page 26.

---

<sup>3</sup>For completeness, note that the **perplexity** of a corpus  $f$  allocated by a model instance  $p$  is defined as  $\text{perp}(f; p) = 2^{H(\tilde{p}; p)}$ . This yields  $\text{perp}(f; p) = |f| \sqrt{\frac{1}{L(f; p)}}$  and  $L(f; p) = \left(\frac{1}{\text{perp}(f; p)}\right)^{|f|}$  as well as the common interpretation that **the perplexity value measures the complexity of the given corpus from the model instance's view**: the perplexity is equal to the size of an imaginary word list from which the corpus can be generated by the model instance – assuming that all words on this list are equally probable. Moreover, the equations state that minimizing the corpus perplexity  $\text{perp}(f; p)$  is equivalent to maximizing the corpus probability  $L(f; p)$ .

## \*Maximum-Entropy Estimation

Readers only interested in the expectation-maximization algorithm are encouraged to omit this section. For completeness, however, note that the relative entropy is **asymmetric**. That means, in general

$$D(p||q) \neq D(q||p)$$

It is easy to check that the triangle inequality is not valid too. So, the relative entropy  $D(.||.)$  is not a “true” distance function. On the other hand,  $D(.||.)$  has some of the properties of a distance function. In particular, it is always non-negative and it is zero if and only if  $p = q$  (see Theorem 3). So far, however, we aimed at minimizing the relative entropy with respect to its **second** argument, filling the first argument slot of  $D(.||.)$  with the relative-frequency estimate  $\tilde{p}$ . Obviously, these observations raise the question, whether it is also possible to derive other “good” estimates by minimizing the relative entropy with respect to its **first** argument. So, in terms of Theorem 2, it might be interesting to ask for model instances  $p^* \in \mathcal{M}$  with

$$D(p^*||\tilde{p}) = \min_{p \in \mathcal{M}} D(p||\tilde{p})$$

For at least two reasons, however, this initial approach of relative-entropy estimation is too simplistic. First, it is tailored to probability models that lack any generalization power. Second, it does not provide deeper insight when estimating constrained probability models. Here are the details:

- A closer look at Definition 7 reveals that the relative entropy  $D(p||\tilde{p})$  is finite for those model instances  $p \in \mathcal{M}$  only that fulfill

$$\tilde{p}(x) = 0 \Rightarrow p(x) = 0$$

So, the initial approach would lead to model instances that are completely unable to generalize, since they are not allowed to allocate positive probabilities to at least some of the types not seen in the training corpus.

- Theorem 2 guarantees that the relative-frequency estimate  $\tilde{p}$  is a solution to the initial approach of relative-entropy estimation, whenever  $\tilde{p} \in \mathcal{M}$ . Now, Definition 8 introduces the constrained probability models  $M_{constr}$ , and indeed, it is easy to check that  $\tilde{p}$  is always an instance of these models. In other words, estimating constrained probability models by the approach above does not result in interesting model instances.

Clearly, all the mentioned drawbacks are due to the fact that the relative-entropy minimization is performed with respect to the relative-frequency estimate. As a resource, we switch simply to a more convenient reference distribution, thereby generalizing formally the initial problem setting. So, as the final request, we ask for model instances  $p^* \in \mathcal{M}$  with

$$D(p^*||p_0) = \min_{p \in \mathcal{M}} D(p||p_0)$$

In this setting, the **reference distribution**  $p_0 \in \mathcal{M}(\mathcal{X})$  is a given instance of the unrestricted probability model, and from what we have seen so far,  $p_0$  should allocate all types of interest a positive probability, and moreover,  $p_0$  should not be itself an instance of the probability

model  $\mathcal{M}$ . Indeed, this request will lead us to the interesting maximum-entropy estimates. Note first, that

$$D(p||p_0) = H(p;p_0) - H(p)$$

So, *minimizing*  $D(p||p_0)$  as a function of the model instances  $p$  is equivalent to *minimizing* the cross entropy  $H(p;p_0)$  and *simultaneously maximizing* the model entropy  $H(p)$ . Now, simultaneous optimization is a hard task in general, and this gives reason to focus firstly on maximizing the entropy  $H(p)$  in isolation. The following definition presents maximum-entropy estimation in terms of the well-known maximum-entropy principle (Jaynes 1957). Sloppily formulated, the maximum-entropy principle recommends to maximize the entropy  $H(p)$  as a function of the instances  $p$  of certain “constrained” probability models.

**Definition 8** Let  $f_1, \dots, f_d$  be a finite number of real-valued functions on a set  $\mathcal{X}$  of types, the so-called **feature functions**<sup>4</sup>. Let  $\tilde{p}$  be the relative-frequency estimate on a non-empty and finite corpus  $f$  on  $\mathcal{X}$ . Then, the **probability model constrained by the expected values of  $f_1 \dots f_d$  on  $f$**  is defined as

$$\mathcal{M}_{constr} = \left\{ p \in \mathcal{M}(\mathcal{X}) \mid E_p f_i = E_{\tilde{p}} f_i \text{ for } i = 1, \dots, d \right\}$$

Here, each  $E_p f_i$  is the **model instance’s expectation** of  $f_i$

$$E_p f_i = \sum_{x \in \mathcal{X}} p(x) f_i(x)$$

constrained to match  $E_{\tilde{p}} f_i$ , the **observed expectation** of  $f_i$

$$E_{\tilde{p}} f_i = \sum_{x \in \mathcal{X}} \tilde{p}(x) f_i(x)$$

Furthermore, a model instance  $p^* \in \mathcal{M}_{constr}$  is called a **maximum-entropy estimate** of  $\mathcal{M}_{constr}$  if and only if

$$H(p^*) = \max_{p \in \mathcal{M}_{constr}} H(p)$$

It is well-known that the maximum-entropy estimates have some nice properties. For example, as Definition 9 and Theorem 4 show, they can be identified to be the unique maximum-likelihood estimates of the so-called exponential models (which are also known as log-linear models).

**Definition 9** Let  $f_1, \dots, f_d$  be a finite number of feature functions on a set  $\mathcal{X}$  of types. The **exponential model of  $f_1, \dots, f_d$**  is defined by

$$\mathcal{M}_{exp} = \left\{ p \in \mathcal{M}(\mathcal{X}) \mid p(x) = \frac{1}{Z_\lambda} e^{\lambda_1 f_1(x) + \dots + \lambda_d f_d(x)} \text{ with } \lambda_1, \dots, \lambda_d, Z_\lambda \in \mathcal{R} \right\}$$

---

<sup>4</sup>Each of these feature functions can be thought of as being constructed by inspecting the set of types, thereby measuring a specific property of the types  $x \in \mathcal{X}$ . For example, if working in a formal-grammar framework, then it might be worthy to look (at least) at some feature functions  $f_r$  directly associated to the rules  $r$  of the given formal grammar. The “measure”  $f_r(x)$  of a specific rule  $r$  for the analyzes  $x \in \mathcal{X}$  of the grammar might be calculated, for example, in terms of the occurrence frequency of  $r$  in the sequence of those rules which are necessary to produce  $x$ . For instance, Chi (1999) studied this approach for the context-free grammar formalism. Note, however, that there is in general no recipe for constructing “good” feature functions: Often, it is really an intellectual challenge to find those feature functions that describe the given data as best as possible (or at least in a satisfying manner).

Here, the **normalizing constant**  $Z_\lambda$  (with  $\lambda$  as a short form for the sequence  $\lambda_1, \dots, \lambda_d$ ) guarantees that  $p \in \mathcal{M}(\mathcal{X})$ , and it is given by

$$Z_\lambda = \sum_{x \in \mathcal{X}} e^{\lambda_1 f_1(x) + \dots + \lambda_d f_d(x)}$$

**Theorem 4** *Let  $f$  be a non-empty and finite corpus, and  $f_1, \dots, f_d$  be a finite number of feature functions on a set  $\mathcal{X}$  of types. Then*

- (i) *The maximum-entropy estimates of  $\mathcal{M}_{constr}$  are instances of  $\mathcal{M}_{exp}$ , and the maximum-likelihood estimates of  $\mathcal{M}_{exp}$  on  $f$  are instances of  $\mathcal{M}_{constr}$ .*
- (ii) *If  $p^* \in \mathcal{M}_{constr} \cap \mathcal{M}_{exp}$ , then  $p^*$  is both a unique maximum-entropy estimate of  $\mathcal{M}_{constr}$  and a unique maximum-likelihood estimate of  $\mathcal{M}_{exp}$  on  $f$ .*

Part (i) of the theorem simply suggests the form of the maximum-entropy or maximum-likelihood estimates we are looking for. By combining both findings of (i), however, the search space is drastically reduced for both estimation methods: We simply have to look at the intersection of the involved probability models. In turn, exactly this fact makes the second part of the theorem so valuable. If there is a maximum-entropy **or** a maximum-likelihood estimate, then it is in the intersection of both models, and thus according to Part (ii), it is a unique estimate, and even more, it is both a maximum-entropy **and** a maximum-likelihood estimate.

**Proof** See e.g. Cover and Thomas (1991), pages 266-278. For an interesting alternate proof of (ii), see Ratnaparkhi (1997). Note, however, that the proof of Ratnaparkhi's Theorem 1 is incorrect, whenever the set  $\mathcal{X}$  of types is infinite. Although Ratnaparkhi's proof is very elegant, it relies on the existence of a uniform distribution on  $\mathcal{X}$  that simply does not exist in this special case. By contrast, Cover and Thomas prove Theorem 11.1.1 without using a uniform distribution on  $\mathcal{X}$ , and so, they achieve indeed the more general result.

Finally, we are coming back to our request of minimizing the relative entropy with respect to a given reference distribution  $p_0 \in \mathcal{M}(\mathcal{X})$ . For constrained probability models, the relevant results differ not much from the results described in Theorem 4. So, let

$$\mathcal{M}_{exp.ref} = \left\{ p \in \mathcal{M}(\mathcal{X}) \mid p(x) = \frac{1}{Z_\lambda} e^{\lambda_1 f_1(x) + \dots + \lambda_d f_d(x)} \cdot p_0(x) \text{ with } \lambda_1, \dots, \lambda_d, Z_\lambda \in \mathcal{R} \right\}$$

Then, along the lines of the proof of Theorem 4 it can be also proven that the following propositions are valid.

- (i) *The minimum relative-entropy estimates of  $\mathcal{M}_{constr}$  are instances of  $\mathcal{M}_{exp.ref}$ , and the maximum-likelihood estimates of  $\mathcal{M}_{exp.ref}$  on  $f$  are instances of  $\mathcal{M}_{constr}$ .*
- (ii) *If  $p^* \in \mathcal{M}_{constr} \cap \mathcal{M}_{exp.ref}$ , then  $p^*$  is both a unique minimum relative-entropy estimate of  $\mathcal{M}_{constr}$  and a unique maximum-likelihood estimate of  $\mathcal{M}_{exp.ref}$  on  $f$ .*

Being a bit sloppy, we summarize all results as follows

$$\text{maximum-likelihood estimation} \text{ of } \left\{ \begin{array}{l} \text{exponential models} \\ \text{exponential models} \\ \text{with reference distributions} \end{array} \right\} \iff \left\{ \begin{array}{l} \text{maximum-entropy} \\ \text{estimation} \\ \text{minimum relative-entropy} \\ \text{estimation} \end{array} \right\}$$

### 3 The Expectation-Maximization Algorithm

Both relative-frequency and maximum-likelihood estimation are common estimation methods with a two-fold input, a corpus and a probability model<sup>5</sup> such that the instances of the model might have generated the corpus. The output of both estimation methods is simply an instance of the probability model, ideally, the unknown distribution that generated the corpus. In contrast to this setting, in which we are almost *completely informed* (the only thing that is not known to us is the unknown distribution that generated the corpus), the expectation-maximization algorithm is designed to estimate an instance of the probability model for settings, in which we are *incompletely informed*.

To be more specific, instead of a **complete-data corpus**, the input of the expectation-maximization algorithm is an **incomplete-data corpus** together with a so-called **symbolic analyzer**. A symbolic analyzer is a device assigning to each **incomplete-data type** a **set of analyzes**, each analysis being a **complete-data type**. As a result, the missing complete-data corpus can be partly compensated by the expectation-maximization algorithm: The application of the the symbolic analyzer to the incomplete-data corpus leads to an *ambiguous complete-data corpus*. The ambiguity arises as a consequence of the inherent **analytical ambiguity** of the symbolic analyzer: the analyzer can replace each token of the incomplete-data corpus by a set of complete-data types – the set of its analyzes – but clearly, the symbolic analyzer is not able to resolve the analytical ambiguity.

The expectation-maximization algorithm performs a sequence of runs over the resulting ambiguous complete-data corpus. Each of these runs consists of an **expectation step** followed by a **maximization step**. In the **E step**, the expectation-maximization algorithm pairs the symbolic analyzer with an instance of the probability model. This combination yields a **statistical analyzer** which is able to **resolve the analytical ambiguity** introduced by the symbolic analyzer. In the **M step**, the expectation-maximization algorithm calculates an ordinary maximum-likelihood estimate on the resolved complete-data corpus.

In general, however, a sequence of such runs is necessary. The reason is that we never know which instance of the given probability model leads to a good statistical analyzer, and thus, which instance of the probability model shall be used in the E-step. The expectation-maximization algorithm provides a simple but somehow surprising solution to this serious problem. At the beginning, a randomly generated **starting instance** of the given probability model is used for the first E-step. In further iterations, the estimate of the M-step is used for the next E-step.

---

<sup>5</sup>We associate the relative-frequency estimate with the unrestricted probability model

## Symbolic and Statistical Analyzers

**Definition 10** Let  $\mathcal{X}$  and  $\mathcal{Y}$  be non-empty and countable sets. A function

$$\mathcal{A}: \mathcal{Y} \rightarrow 2^{\mathcal{X}}$$

is called a **symbolic analyzer** if the (possibly empty) **sets of analyzes**  $\mathcal{A}(y) \subseteq \mathcal{X}$  are pair-wise disjoint, and the union of all sets of analyzes  $\mathcal{A}(y)$  is complete

$$\mathcal{X} = \sum_{y \in \mathcal{Y}} \mathcal{A}(y)$$

In this case,  $\mathcal{Y}$  is called the set of **incomplete-data types**, whereas  $\mathcal{X}$  is called the set of **complete-data types**. So, in other words, the analyzes  $\mathcal{A}(y)$  of the incomplete-data types  $y$  form a partition of the complete-data  $\mathcal{X}$ . Therefore, for each  $x \in \mathcal{X}$  exists a unique  $y \in \mathcal{Y}$ , the so-called **yield** of  $x$ , such that  $x$  is an analysis of  $y$

$$y = \text{yield}(x) \quad \text{if and only if} \quad x \in \mathcal{A}(y)$$

For example, if working in a formal-grammar framework, the *grammatical sentences* can be interpreted as the incomplete-data types, whereas the *grammatical analyzes of the sentences* are the complete-data types. So, in terms of Definition 10, a so-called *parser* – a device assigning a set of grammatical analyzes to a given sentence – is clearly a symbolic analyzer: The most important thing to check is that the parser does not assign a given grammatical analysis to two different sentences – which is pretty obvious, if the *sentence words* are part of the grammatical analyzes.

**Definition 11** A pair  $\langle \mathcal{A}, p \rangle$  consisting of a symbolic analyzer  $\mathcal{A}$  and a probability distribution  $p$  on the complete-data types  $\mathcal{X}$  is called a **statistical analyzer**. We use a statistical analyzer to **induce probabilities** for the incomplete-data types  $y \in \mathcal{Y}$

$$p(y) := \sum_{x \in \mathcal{A}(y)} p(x)$$

Even more important, we use a statistical analyzer to **resolve the analytical ambiguity** of an incomplete-data type  $y \in \mathcal{Y}$  by looking at the **conditional probabilities of the analyzes**  $x \in \mathcal{A}(y)$

$$p(x|y) := \frac{p(x)}{p(y)} \quad \text{where } y = \text{yield}(x)$$

It is easy to check that the statistical analyzer induces a proper probability distribution on the set  $\mathcal{Y}$  of incomplete-data types

$$\sum_{y \in \mathcal{Y}} p(y) = \sum_{y \in \mathcal{Y}} \sum_{x \in \mathcal{A}(y)} p(x) = \sum_{x \in \mathcal{X}} p(x) = 1$$

Moreover, the statistical analyzer induces also proper conditional probability distributions on the sets of analyzes  $\mathcal{A}(y)$

$$\sum_{x \in \mathcal{A}(y)} p(x|y) = \sum_{x \in \mathcal{A}(y)} \frac{p(x)}{p(y)} = \frac{\sum_{x \in \mathcal{A}(y)} p(x)}{p(y)} = \frac{p(y)}{p(y)} = 1$$

Of course, by defining  $p(x|y) = 0$  for  $y \neq \text{yield}(x)$ ,  $p(\cdot|y)$  is even a probability distribution on the full set  $\mathcal{X}$  of analyzes.

# Input, Procedure, and Output of the EM Algorithm

**Definition 12** *The input of the expectation-maximization (EM) algorithm is*

- (i) a **symbolic analyzer**, i.e., a function  $\mathcal{A}$  which assigns a **set of analyzes**  $\mathcal{A}(y) \subseteq \mathcal{X}$  to each **incomplete-data type**  $y \in \mathcal{Y}$ , such that all sets of analyzes form a partition of the set  $\mathcal{X}$  of **complete-data types**

$$\mathcal{X} = \sum_{y \in \mathcal{Y}} \mathcal{A}(y)$$

- (ii) a non-empty and finite **incomplete-data corpus**, i.e., a frequency distribution  $f$  on the set of incomplete-data types

$$f: \mathcal{Y} \rightarrow \mathcal{R} \quad \text{such that} \quad f(y) \geq 0 \text{ for all } y \in \mathcal{Y} \quad \text{and} \quad 0 < |f| < \infty$$

- (iii) a **complete-data model**  $\mathcal{M} \subseteq \mathcal{M}(\mathcal{X})$ , i.e., each instance  $p \in \mathcal{M}$  is a probability distribution on the set of complete-data types

$$p: \mathcal{X} \rightarrow [0, 1] \quad \text{and} \quad \sum_{x \in \mathcal{X}} p(x) = 1$$

- (\*) **implicit input**: an **incomplete-data model**  $\mathcal{M} \subseteq \mathcal{M}(\mathcal{Y})$  induced by the symbolic analyzer and the complete-data model. To see this, recall Definition 11. Together with a given instance of the complete-data model, the symbolic analyzer constitutes a statistical analyzer which, in turn, induces the following instance of the incomplete-data model

$$p: \mathcal{Y} \rightarrow [0, 1] \quad \text{and} \quad p(y) = \sum_{x \in \mathcal{A}(y)} p(x)$$

(Note: For both complete and incomplete data, the same notation symbols  $\mathcal{M}$  and  $p$  are used. The sloppy notation, however, is justified, because the incomplete-data model is a marginal of the complete-data model.)

- (iv) a (randomly generated) **starting instance**  $p_0$  of the complete-data model  $\mathcal{M}$ .  
(Note: If permitted by  $\mathcal{M}$ , then  $p_0$  should not assign to any  $x \in \mathcal{X}$  a probability of zero.)

**Definition 13** *The procedure of the EM algorithm is*

- (1) for each  $i = 1, 2, 3, \dots$  do  
 (2)  $q := p_{i-1}$   
 (3) **E-step**: compute the **complete-data corpus**  $f_q: \mathcal{X} \rightarrow \mathcal{R}$  expected by  $q$

$$f_q(x) := f(y) \cdot q(x|y) \quad \text{where } y = \text{yield}(x)$$

- (4) **M-step**: compute a maximum-likelihood estimate  $\hat{p}$  of  $\mathcal{M}$  on  $f_q$

$$L(f_q; \hat{p}) = \max_{p \in \mathcal{M}} L(f_q, p)$$

(Implicit pre-condition of the EM algorithm: it exists!)

- (5)  $p_i := \hat{p}$   
 (6) end // for each  $i$   
 (7) print  $p_0, p_1, p_2, p_3, \dots$

In line (3) of the EM procedure, a complete-data corpus  $f_q(x)$  has to be generated on the basis of the incomplete-data corpus  $f(y)$  and the conditional probabilities  $q(x|y)$  of the analyzes of  $y$  (conditional probabilities are introduced in Definition 11). In fact, this generation procedure is conceptually very easy: according to the conditional probabilities  $q(x|y)$ , the frequency  $f(y)$  has to be distributed among the complete-data types  $x \in \mathcal{A}(y)$ . Moreover, there exists a simple reversed procedure (summation of all frequencies  $f_q(x)$  with  $x \in \mathcal{A}(y)$ ) which guarantees that the original corpus  $f(y)$  can be recovered from the generated corpus  $f_q(x)$ . Finally, the size of both corpora is the same

$$|f_q| = |f|$$

In line (4) of the EM procedure, it is stated that a maximum-likelihood estimate  $\hat{p}$  of the complete-data model has to be computed on the complete-data corpus  $f_q$  expected by  $q$ . Recall for this purpose that the probability of  $f_q$  allocated by an instance  $p \in \mathcal{M}$  is defined as

$$L(f_q; p) = \prod_{x \in \mathcal{X}, f_q(x) \neq 0} p(x)^{f_q(x)}$$

In contrast, the probability of the incomplete-data corpus  $f$  allocated by an instance  $p$  of the incomplete-data model is much more complex. Using Definition 12.\*, we get an expression involving a product of sums

$$L(f; p) = \prod_{y \in \mathcal{Y}, f(y) \neq 0} \left( \sum_{x \in \mathcal{A}(y)} p(x) \right)^{f(y)}$$

Nevertheless, the following theorem reveals that the EM algorithm aims at finding an instance of the incomplete-data model which possibly maximizes the probability of the incomplete-data corpus.

**Theorem 5** *The output of the EM algorithm is: A sequence of instances of the complete-data model  $\mathcal{M}$ , the so-called **EM re-estimates**,*

$$p_0, p_1, p_2, p_3, \dots$$

*such that the sequence of probabilities allocated to the incomplete-data corpus is monotonic increasing*

$$L(f; p_0) \leq L(f; p_1) \leq L(f; p_2) \leq L(f; p_3) \leq \dots$$

It is common wisdom that the sequence of EM re-estimates will converge to a (local) maximum-likelihood estimate of the incomplete-data model on the incomplete-data corpus. As proven by Wu (1983), however, the EM algorithm will do this only in specific circumstances. Of course, it is guaranteed that the sequence of corpus probabilities (allocated by the EM re-estimates) must converge. However, we are more interested in the behavior of the EM re-estimates itself. Now, intuitively, the EM algorithm might get stuck in a saddle point or even a local minimum of the corpus-probability function, whereas the associated model instances are hopping uncontrolled around (for example, on a circle-like path in the “space” of all model instances).

**Proof** See theorems 6 and 7.

# The Generalized Expectation-Maximization Algorithm

The EM algorithm performs a sequence of maximum-likelihood estimations on complete data, resulting in good re-estimates on incomplete-data (“good” in the sense of Theorem 5). The following theorem, however, reveals that the EM algorithm might overdo it somehow, since there exist alternative M-steps which can be easier performed, and which result in re-estimates having the same property as the EM re-estimates.

**Definition 14** *A generalized expectation-maximization (GEM) algorithm has exactly the same input as the EM-algorithm, but an easier M-step is performed in its procedure:*

(4) **M-step (GEM):** compute an instance  $\hat{p}$  of the complete-data model  $\mathcal{M}$  such that

$$L(f_q; \hat{p}) \geq L(f_q; q)$$

**Theorem 6** *The output of a GEM algorithm is: A sequence of instances of the complete-data model  $\mathcal{M}$ , the so-called **GEM re-estimates**, such that the sequence of probabilities allocated to the incomplete-data corpus is monotonic increasing.*

**Proof** Various proofs have been given in the literature. The first one was presented by Dempster et al. (1977). For other variants of the EM algorithm, the book of McLachlan and Krishnan (1997) is a good source. Here, we present something along the lines of the original proof. Clearly, a proof of the theorem requires somehow that we are able to express *the probability of the given incomplete-data corpus  $f$*  in terms of the *the probabilities of complete-data corpora  $f_q$*  which are involved in the M-steps of the GEM algorithm (where both types of corpora are allocated a probability by the same instance  $p$  of the model  $\mathcal{M}$ ). A certain entity, which we would like to call the **expected cross-entropy on the analyzes**, plays a major role for solving this task. To be specific, the expected cross-entropy on the analyzes is defined as the expectation of certain cross-entropy values  $H_{\mathcal{A}(y)}(q, p)$  which are calculated on the different sets  $\mathcal{A}(y)$  of analyzes. Then, of course, the “expectation” is calculated on the basis of the relative-frequency estimate  $\tilde{p}$  of the given incomplete-data corpus

$$H_{\mathcal{A}}(q; p) = \sum_{y \in \mathcal{Y}} \tilde{p}(y) \cdot H_{\mathcal{A}(y)}(q; p)$$

Now, for two instances  $q$  and  $p$  of the complete-data model, their conditional probabilities  $q(x|y)$  and  $p(x|y)$  form proper probability distributions on the set  $\mathcal{A}(y)$  of analyzes of  $y$  (see Definition 11). So, the cross-entropy  $H_{\mathcal{A}(y)}(q; p)$  on the set  $\mathcal{A}(y)$  is simply given by

$$H_{\mathcal{A}(y)}(q; p) = - \sum_{x \in \mathcal{A}(y)} q(x|y) \log p(x|y)$$

Recalling the central task of this proof, a bunch of relatively straight-forward calculations leads to the following interesting equation<sup>6</sup>

$$L(f; p) = \left(2^{H_{\mathcal{A}}(q; p)}\right)^{|f|} \cdot L(f_q; p)$$

---

<sup>6</sup>It is easier to show that

$$H(\tilde{p}; p) = H(\tilde{p}_q; p) - H_{\mathcal{A}}(q; p).$$

Here,  $\tilde{p}$  is the relative-frequency estimate on the incomplete-data corpus  $f$ , whereas  $\tilde{p}_q$  is the relative-frequency estimate on the complete-data corpus  $f_q$ . However, by defining an “average perplexity of the analyzes”,

Using this equation, we can state that

$$\frac{L(f; p)}{L(f; q)} = \left( 2^{H_{\mathcal{A}}(q; p) - H_{\mathcal{A}}(q, q)} \right)^{|f|} \cdot \frac{L(f_q; p)}{L(f_q; q)}$$

In what follows, we will show that, after each M-step of a GEM algorithm (i.e. for  $p$  being a GEM re-estimate  $\hat{p}$ ), both of the factors on the right-hand side of this equation are not less than one. First, an iterated application of the information inequality of information theory (see Theorem 3) yields

$$\begin{aligned} H_{\mathcal{A}}(q; p) - H_{\mathcal{A}}(q, q) &= \sum_{y \in \mathcal{Y}} \tilde{p}(y) \cdot \left( H_{\mathcal{A}(y)}(q; p) - H_{\mathcal{A}(y)}(q, q) \right) \\ &= \sum_{y \in \mathcal{Y}} \tilde{p}(y) \cdot D_{\mathcal{A}(y)}(q \| p) \\ &\geq 0 \end{aligned}$$

So, the first factor is never (i.e. for no model instance  $p$ ) less than one

$$\left( 2^{H_{\mathcal{A}}(q; p) - H_{\mathcal{A}}(q, q)} \right)^{|f|} \geq 1$$

Second, by definition of the M-step of a GEM algorithm, the second factor is also not less than one

$$\frac{L(f_q; \hat{p})}{L(f_q; q)} \geq 1$$

So, it follows

$$\frac{L(f; \hat{p})}{L(f; q)} \geq 1$$

yielding that the probability of the incomplete-data corpus allocated by the GEM re-estimate  $\hat{p}$  is not less than the probability of the incomplete-data corpus allocated by the model instance  $q$  (which is either the starting instance  $p_0$  of the GEM algorithm or the previously calculated GEM re-estimate)

$$L(f; \hat{p}) \geq L(f; q)$$

**Theorem 7** *An EM algorithm is a GEM algorithm.*

**Proof** In the M-step of an EM algorithm, a model instance  $\hat{p}$  is selected such that

$$L(f_q; \hat{p}) = \max_{p \in \mathcal{M}} L(f_q, p)$$

So, especially

$$L(f_q; \hat{p}) \geq L(f_q, q)$$

and the requirements of the M-step of a GEM algorithm are met.

---

  $\text{perp}_{\mathcal{A}}(q; p) := 2^{H_{\mathcal{A}}(q; p)}$  (see also Footnote 3), the true spirit of the equation can be revealed:

$$L(f_q; p) = L(f; p) \cdot \left( \frac{1}{\text{perp}_{\mathcal{A}}(q; p)} \right)^{|f|}$$

This equation states that the probability of a complete-data corpus (generated by a statistical analyzer) is the product of the probability of the given incomplete-data corpus and  $|f|$ -times the average probability of the different corpora of analyzes (as generated for each of the  $|f|$  tokens of the incomplete-data corpus).

## 4 Rolling Two Dice

**Example 1** We shall now consider an experiment in which two loaded dice are rolled, and we shall compute the relative-frequency estimate on a corpus of outcomes.

If we assume that the two dice are distinguishable, each outcome can be represented as a pair of numbers  $(x_1, x_2)$ , where  $x_1$  is the number that appears on the first die and  $x_2$  is the number that appears on the second die. So, for this experiment, an appropriate set  $\mathcal{X}$  of types comprises the following 36 outcomes:

$(x_1, x_2)$	$x_2 = 1$	$x_2 = 2$	$x_2 = 3$	$x_2 = 4$	$x_2 = 5$	$x_2 = 6$
$x_1 = 1$	(1, 1)	(1, 2)	(1, 3)	(1, 4)	(1, 5)	(1, 6)
$x_1 = 2$	(2, 1)	(2, 2)	(2, 3)	(2, 4)	(2, 5)	(2, 6)
$x_1 = 3$	(3, 1)	(3, 2)	(3, 3)	(3, 4)	(3, 5)	(3, 6)
$x_1 = 4$	(4, 1)	(4, 2)	(4, 3)	(4, 4)	(4, 5)	(4, 6)
$x_1 = 5$	(5, 1)	(5, 2)	(5, 3)	(5, 4)	(5, 5)	(5, 6)
$x_1 = 6$	(6, 1)	(6, 2)	(6, 3)	(6, 4)	(6, 5)	(6, 6)

If we throw the two dice a 100 000 times, then the following occurrence frequencies might arise

$f(x_1, x_2)$	$x_2 = 1$	$x_2 = 2$	$x_2 = 3$	$x_2 = 4$	$x_2 = 5$	$x_2 = 6$
$x_1 = 1$	3790	3773	1520	1498	2233	2298
$x_1 = 2$	3735	3794	1497	1462	2269	2184
$x_1 = 3$	4903	4956	1969	2035	2883	3010
$x_1 = 4$	2495	2519	1026	1049	1487	1451
$x_1 = 5$	3820	3735	1517	1498	2276	2191
$x_1 = 6$	6369	6290	2600	2510	3685	3673

The size of this corpus is  $|f| = 100\,000$ . So, the relative-frequency estimate  $\tilde{p}$  on  $f$  can be easily computed (see Definition 4)

$\tilde{p}(x_1, x_2)$	$x_2 = 1$	$x_2 = 2$	$x_2 = 3$	$x_2 = 4$	$x_2 = 5$	$x_2 = 6$
$x_1 = 1$	0.03790	0.03773	0.01520	0.01498	0.02233	0.02298
$x_1 = 2$	0.03735	0.03794	0.01497	0.01462	0.02269	0.02184
$x_1 = 3$	0.04903	0.04956	0.01969	0.02035	0.02883	0.03010
$x_1 = 4$	0.02495	0.02519	0.01026	0.01049	0.01487	0.01451
$x_1 = 5$	0.03820	0.03735	0.01517	0.01498	0.02276	0.02191
$x_1 = 6$	0.06369	0.06290	0.02600	0.02510	0.03685	0.03673

**Example 2** We shall consider again Experiment 1 in which two loaded dice are rolled, but we shall now compute the relative-frequency estimate on the corpus of outcomes of the first die, as well as on the corpus of outcomes of the second die.

If we look at the same corpus as in Example 1, then the corpus  $f_1$  of outcomes of the first die can be calculated as  $f_1(x_1) = \sum_{x_2} f(x_1, x_2)$ . An analog summation yields the corpus of outcomes of the second die,  $f_2(x_2) = \sum_{x_1} f(x_1, x_2)$ . Obviously, the sizes of all corpora are identical  $|f_1| = |f_2| = |f| = 100\,000$ . So, the relative-frequency estimates  $\tilde{p}_1$  on  $f_1$  and  $\tilde{p}_2$  on

$f_2$  are calculated as follows

$f_1(x_1)$	$x_1$	$\tilde{p}_1(x_1)$	$x_1$	$f_2(x_2)$	$x_2$	$\tilde{p}_2(x_2)$	$x_2$
15112	1	0.15112	1	25112	1	0.25112	1
14941	2	0.14941	2	25067	2	0.25067	2
19756	3	0.19756	3	10129	3	0.10129	3
10027	4	0.10027	4	10052	4	0.10052	4
15037	5	0.15037	5	14833	5	0.14833	5
25127	6	0.25127	6	14807	6	0.14807	6

**Example 3** We shall consider again Experiment 1 in which two loaded dice are rolled, but we shall now compute a maximum-likelihood estimate of the probability model which assumes that the numbers appearing on the first and second die are statistically independent.

First, recall the definition of statistical independence (see e.g. Duda et al. (2001), page 613).

**Definition 15** The variables  $x_1$  and  $x_2$  are said to be **statistically independent** given a joint probability distribution  $p$  on  $\mathcal{X}$  if and only if

$$p(x_1, x_2) = p_1(x_1) \cdot p_2(x_2)$$

where  $p_1$  and  $p_2$  are the **marginal distributions** for  $x_1$  and  $x_2$

$$p_1(x_1) = \sum_{x_2} p(x_1, x_2)$$

$$p_2(x_2) = \sum_{x_1} p(x_1, x_2)$$

So, let  $\mathcal{M}_{1/2}$  be the probability model which assumes that the numbers appearing on the first and second die are statistically independent

$$\mathcal{M}_{1/2} = \{p \in \mathcal{M}(\mathcal{X}) \mid x_1 \text{ and } x_2 \text{ are statistically independent given } p\}$$

In Example 1, we have calculated the relative-frequency estimator  $\tilde{p}$ . Theorem 1 states that  $\tilde{p}$  is the unique maximum-likelihood estimate of the unrestricted model  $\mathcal{M}(\mathcal{X})$ . Thus,  $\tilde{p}$  is also a candidate for a maximum-likelihood estimate of  $\mathcal{M}_{1/2}$ . Unfortunately, however,  $x_1$  and  $x_2$  are **not** statistically independent given  $\tilde{p}$  (see e.g.  $\tilde{p}(1, 1) = 0.03790$  and  $\tilde{p}_1(1) \cdot \tilde{p}_2(1) = 0.0379493$ ). This has two consequences for the experiment in which two (loaded) dice are rolled:

- the probability model, which assumes that the numbers appearing on the first and second die are statistically independent, is a **restricted model** (see Definition 5), and
- **the relative-frequency estimate is in general not a maximum-likelihood estimate of the standard probability model** assuming that the numbers appearing on the first and second die are statistically independent.

Therefore, we are now following Definition 6 to compute the maximum-likelihood estimate of  $\mathcal{M}_{1/2}$ . Using the independence property, the probability of the corpus  $f$  allocated by an instance  $p$  of the model  $\mathcal{M}_{1/2}$  can be calculated as

$$L(f; p) = \left( \prod_{x_1=1, \dots, 6} p_1(x_1)^{f_1(x_1)} \right) \cdot \left( \prod_{x_2=1, \dots, 6} p_2(x_2)^{f_2(x_2)} \right) = L(f_1; p_1) \cdot L(f_2; p_2)$$

Definition 6 states that the maximum-likelihood estimate  $\hat{p}$  of  $\mathcal{M}_{1/2}$  on  $f$  must maximize  $L(f;p)$ . A product, however, is maximized, if and only if its factors are simultaneously maximized. Theorem 1 states that the corpus probabilities  $L(f_i;p_i)$  are maximized by the relative-frequency estimators  $\tilde{p}_i$ . Therefore, the product of the relative-frequency estimators  $\tilde{p}_1$  and  $\tilde{p}_2$  (on  $f_1$  and  $f_2$  respectively) might be a candidate for the maximum-likelihood estimate  $\hat{p}$  we are looking for

$$\hat{p}(x_1, x_2) = \tilde{p}_1(x_1) \cdot \tilde{p}_2(x_2)$$

Now, note that the marginal distributions of  $\hat{p}$  are identical with the relative-frequency estimators on  $f_1$  and  $f_2$ . For example,  $\hat{p}$ 's marginal distribution for  $x_1$  is calculated as

$$\hat{p}_1(x_1) = \sum_{x_2} \hat{p}(x_1, x_2) = \sum_{x_2} \tilde{p}_1(x_1) \cdot \tilde{p}_2(x_2) = \tilde{p}_1(x_1) \cdot \sum_{x_2} \tilde{p}_2(x_2) = \tilde{p}_1(x_1) \cdot 1 = \tilde{p}_1(x_1)$$

A similar calculation yields  $\hat{p}_2(x_2) = \tilde{p}_2(x_2)$ . Both equations state that  $x_1$  and  $x_2$  are indeed statistically independent given  $\hat{p}$

$$\hat{p}(x_1, x_2) = \hat{p}_1(x_1) \cdot \hat{p}_2(x_2)$$

So, finally, it is guaranteed that  $\hat{p}$  is an instance of the probability model  $\mathcal{M}_{1/2}$  as required for a maximum-likelihood estimate of  $\mathcal{M}_{1/2}$ . *Note:  $\hat{p}$  is even an unique maximum-likelihood estimate since the relative-frequency estimates  $\tilde{p}_i$  are unique maximum-likelihood estimates (see Theorem 1).* The relative-frequency estimates  $\tilde{p}_1$  and  $\tilde{p}_2$  have already been calculated in Example 2. So,  $\hat{p}$  is calculated as follows

$\hat{p}(x_1, x_2)$	$x_2 = 1$	$x_2 = 2$	$x_2 = 3$	$x_2 = 4$	$x_2 = 5$	$x_2 = 6$
$x_1 = 1$	0.0379493	0.0378813	0.0153069	0.0151906	0.0224156	0.0223763
$x_1 = 2$	0.0375198	0.0374526	0.0151337	0.0150187	0.022162	0.0221231
$x_1 = 3$	0.0496113	0.0495224	0.0200109	0.0198587	0.0293041	0.0292527
$x_1 = 4$	0.0251798	0.0251347	0.0101563	0.0100791	0.014873	0.014847
$x_1 = 5$	0.0377609	0.0376932	0.015231	0.0151152	0.0223044	0.0222653
$x_1 = 6$	0.0630989	0.0629859	0.0254511	0.0252577	0.0372709	0.0372055

**Example 4** We shall consider again Experiment 1 in which two loaded dice are rolled. Now, however, we shall assume that we are incompletely informed: the corpus of outcomes (which is given to us) consists only of the sums of the numbers which appear on the first and second die. Nevertheless, we shall compute an estimate for a probability model on the complete-data  $(x_1, x_2) \in \mathcal{X}$ .

If we assume that the corpus which is given to us was calculated on the basis of the corpus given in Example 1, then the occurrence frequency of a sum  $y$  can be calculated as  $f(y) = \sum_{x_1+x_2=y} f(x_1, x_2)$ . These numbers are displayed in the following table

$f(y)$	$y$
3790	2
7508	3
10217	4
10446	5
12003	6
17732	7
13923	8
8595	9
6237	10
5876	11
3673	12

The problem is now, whether this corpus of sums can be used to calculate a good estimate on the outcomes  $(x_1, x_2)$  itself. *Hint: Examples 1 and 3 have shown that a unique relative-frequency estimate  $\tilde{p}(x_1, x_2)$  and a unique maximum-likelihood estimate  $\hat{p}(x_1, x_2)$  can be calculated on the basis of the corpus  $f(x_1, x_2)$ . However, right now, this corpus is not available!* Putting the example in the framework of the EM algorithm (see Definition 12), the set of **incomplete-data types** is

$$\mathcal{Y} = \{2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12\}$$

whereas the set of **complete-data types** is  $\mathcal{X}$ . We also know the **set of analyzes for each incomplete-data type**  $y \in \mathcal{Y}$

$$\mathcal{A}(y) = \{(x_1, x_2) \in \mathcal{X} \mid x_1 + x_2 = y\}$$

As in Example 3, we are especially interested in an estimate of the (slightly restricted) **complete-data model**  $\mathcal{M}_{1/2}$  which assumes that the numbers appearing on the first and second die are statistically independent. So, for this case, a randomly generated **starting instance**  $p_0(x_1, x_2)$  of the complete-data model is simply the product of a randomly generated probability distribution  $p_{01}(x_1)$  for the numbers appearing on the first dice, and a randomly generated probability distribution  $p_{02}(x_2)$  for the numbers appearing on the second dice

$$p_0(x_1, x_2) = p_{01}(x_1) \cdot p_{02}(x_2)$$

The following tables display some randomly generated numbers for  $p_{01}$  and  $p_{02}$

$p_{01}(x_1)$	$x_1$	$p_{02}(x_2)$	$x_2$
0.18	1	0.22	1
0.19	2	0.23	2
0.16	3	0.13	3
0.13	4	0.16	4
0.17	5	0.14	5
0.17	6	0.12	6

Using the random numbers for  $p_{01}(x_1)$  and  $p_{02}(x_2)$ , a starting instance  $p_0$  of the complete-data model  $\mathcal{M}_{1/2}$  is calculated as follows

$p_0(x_1, x_2)$	$x_2 = 1$	$x_2 = 2$	$x_2 = 3$	$x_2 = 4$	$x_2 = 5$	$x_2 = 6$
$x_1 = 1$	0.0396	0.0414	0.0234	0.0288	0.0252	0.0216
$x_1 = 2$	0.0418	0.0437	0.0247	0.0304	0.0266	0.0228
$x_1 = 3$	0.0352	0.0368	0.0208	0.0256	0.0224	0.0192
$x_1 = 4$	0.0286	0.0299	0.0169	0.0208	0.0182	0.0156
$x_1 = 5$	0.0374	0.0391	0.0221	0.0272	0.0238	0.0204
$x_1 = 6$	0.0374	0.0391	0.0221	0.0272	0.0238	0.0204

So, we are ready to start the procedure of the EM algorithm.

**First EM iteration.** In the **E-step**, we shall compute the **complete-data corpus**  $f_q$  expected by  $q := p_0$ . For this purpose, the probability of each incomplete-data type given the starting instance  $p_0$  of the complete-data model has to be computed (see Definition 12.\*)

$$p_0(y) = \sum_{x_1+x_2=y} p_0(x_1, x_2)$$

The above displayed numbers for  $p_0(x_1, x_2)$  yield the following instance of the incomplete-data model

$p_0(y)$	$y$
0.0396	2
0.0832	3
0.1023	4
0.1189	5
0.1437	6
0.1672	7
0.1272	8
0.0867	9
0.0666	10
0.0442	11
0.0204	12

So, the complete-data corpus expected by  $p_0$  is calculated as follows (see line (3) of the EM procedure given in Definition 13)

$f_q(x_1, x_2)$	$x_2 = 1$	$x_2 = 2$	$x_2 = 3$	$x_2 = 4$	$x_2 = 5$	$x_2 = 6$
$x_1 = 1$	3790	3735.95	2337.03	2530.23	2104.91	2290.74
$x_1 = 2$	3772.05	4364.45	2170.03	2539.26	2821	2495.63
$x_1 = 3$	3515.53	3233.08	1737.39	2714.95	2451.85	1903.39
$x_1 = 4$	2512.66	2497.49	1792.29	2276.72	1804.26	1460.92
$x_1 = 5$	3123.95	4146.66	2419.01	2696.47	2228.84	2712
$x_1 = 6$	3966.37	4279.79	2190.88	2547.24	3164	3673

In the **M-step**, we shall compute a maximum-likelihood estimate  $p_1 := \hat{p}$  of the complete-data model  $\mathcal{M}_{1/2}$  on the complete-data corpus  $f_q$ . This can be done along the lines of Examples 2 and 3. *Note: This is more or less the trick of the EM-algorithm! If it appears*

to be difficult to compute a maximum-likelihood estimate of an incomplete-data model then the EM algorithm might solve your problem. It performs a sequence of maximum-likelihood estimations on complete-data corpora. These corpora contain in general more complex data, but nevertheless, it might be well-known, how one has to deal with this data! The following tables display the relative-frequency estimates ( $p_{11}$  and  $p_{12}$ ) on the corpus of outcomes of the first die, as well as on the corpus of outcomes of the second die (where both these corpora are computed on the basis of  $f_q$ )

$p_{11}(x_1)$	$x_1$	$p_{12}(x_2)$	$x_2$
0.170323	1	0.20967	1
0.183442	2	0.225184	2
0.155224	3	0.125066	3
0.116812	4	0.144277	4
0.168342	5	0.142496	5
0.205858	6	0.153307	6

So, the following instance is the maximum-likelihood estimate of the model  $\mathcal{M}_{1/2}$  on  $f_q$

$p_1(x_1, x_2)$	$x_2 = 1$	$x_2 = 2$	$x_2 = 3$	$x_2 = 4$	$x_2 = 5$	$x_2 = 6$
$x_1 = 1$	0.0347204	0.0373677	0.0212322	0.0256952	0.0244696	0.0244038
$x_1 = 2$	0.0375609	0.0404247	0.0229692	0.0277973	0.0264715	0.0264003
$x_1 = 3$	0.0321711	0.034624	0.0196733	0.0238086	0.022673	0.022612
$x_1 = 4$	0.0255287	0.0274752	0.0156113	0.0188928	0.0179917	0.0179433
$x_1 = 5$	0.035833	0.0385651	0.0219126	0.0265186	0.0252538	0.0251858
$x_1 = 6$	0.0409916	0.044117	0.0250672	0.0303363	0.0288893	0.0288116

So, we are ready for the second EM iteration, where an estimate  $p_2$  is calculated. If we continue in this manner, we will arrive finally at the

**1584th EM iteration.** The estimate which is calculated here is

$p_{1584,1}(x_1)$	$x_1$	$p_{1584,2}(x_2)$	$x_2$
0.158396	1	0.239281	1
0.141282	2	0.260559	2
0.204291	3	0.104026	3
0.0785532	4	0.111957	4
0.172207	5	0.134419	5
0.24527	6	0.149758	6

yielding

$p_{1584}(x_1, x_2)$	$x_2 = 1$	$x_2 = 2$	$x_2 = 3$	$x_2 = 4$	$x_2 = 5$	$x_2 = 6$
$x_1 = 1$	0.0379012	0.0412715	0.0164773	0.0177336	0.0212914	0.0237211
$x_1 = 2$	0.0338061	0.0368123	0.014697	0.0158175	0.018991	0.0211581
$x_1 = 3$	0.048883	0.0532299	0.0212516	0.0228718	0.0274606	0.0305942
$x_1 = 4$	0.0187963	0.0204678	0.00817158	0.00879459	0.0105591	0.011764
$x_1 = 5$	0.0412059	0.0448701	0.017914	0.0192798	0.0231479	0.0257894
$x_1 = 6$	0.0586885	0.0639074	0.0255145	0.0274597	0.032969	0.0367312

In this example, more EM iterations will result in exactly the same re-estimates. So, this is a strong reason to quit the EM procedure. Comparing  $p_{1584,1}$  and  $p_{1584,2}$  with the results of Example 2 (*Hint: where we have assumed that a complete-data corpus is given to us!*), we see that the EM algorithm yields pretty similar estimates.

## 5 Probabilistic Context-Free Grammars

This Section provides a more substantial example based on the **context-free grammar** or **CFG** formalism, and it is organized as follows: First, we will give some background information about CFGs, thereby motivating that treating CFGs as generators leads quite naturally to the notion of a probabilistic context-free grammar (PCFG). Second, we provide some additional background information about ambiguity resolution by probabilistic CFGs, thereby focusing on the fact that probabilistic CFGs can resolve ambiguities, if the underlying CFG has a sufficiently high expressive power. For other cases, we are pin-pointing to some useful grammar-transformation techniques. Third, we will investigate the standard probability model of CFGs, thereby proving that this model is restricted in almost all cases of interest. Furthermore, we will give a new formal proof that maximum-likelihood estimation of a CFG's probability model on a corpus of trees is equal to the well-known and especially simple treebank-training method. Finally, we will present the EM algorithm for training a (manually written) CFG on a corpus of sentences, thereby pin-pointing to the fact that EM training simply consists of an iterative sequence of treebank-training steps. Small toy examples will accompany all proofs that are given in this Section.

### Background: Probabilistic Modeling of CFGs

Being a bit sloppy (see e.g. Hopcroft and Ullman (1979) for a formal definition), a CFG simply consists of a finite set of **rules**, where in turn, each rule consists of two parts being separated by a special symbol “ $\rightarrow$ ”, the so-called **rewriting symbol**. The two parts of a rule are made up of so-called **terminal** and **non-terminal** symbols: a rule's left-hand side simply consists of a single non-terminal symbol, whereas the right-hand side is a finite sequence of terminal and non-terminal symbols<sup>7</sup>. Finally, the set of non-terminal symbols contains at least one so-called **starting symbol**. CFGs are also called **phrase-structure grammars**, and the formalism is equivalent to **Backus-Naur forms** or **BNF** introduced by Backus (1959). In computational linguistics, a CFG is usually used in two ways

- as a **generator**: a device for generating sentences, or
- as a **parser**: a device for assigning structure to a given sentence


In the following, we will briefly discuss these two issues. First of all, note that in natural language, words do not occur in any order. Instead, languages have constraints on word order<sup>8</sup>. The central idea underlying phrase-structure grammars is that words are organized into **phrases**, i.e., grouping of words that form a unit. Phrases can be detected, for example,

---


<sup>7</sup>As a consequence, the terminal and non-terminal symbols of a given CFG form two finite and disjoint sets.

<sup>8</sup>Note, however, that so-called **free-word-order languages** (like Czech, German, or Russian) permit many different ways of ordering the words in a sentence (without a change in meaning). Instead of word order, these languages use case markings to indicate who did what to whom.

by their ability (i) to stand alone (e.g. as an answer of a wh-question), (ii) to occur in various sentence positions, or by their ability (iii) to show uniform syntactic possibilities for expansion or substitution. As an example, here is the very first context-free grammar parse tree presented by Chomsky (1956):




As being displayed, Chomsky identified for the sentence “the man took the book” (encoded in the leaf nodes of the parse tree) the following phrases: two **noun phrases**, “the man” and “the book” (the figure displays them as NP subtrees), and one **verb phrase**, “took the book” (displayed as VP subtree). The following list of sentences, where these three phrases have been substituted or expanded, bears some evidence for Chomsky’s analysis:



Chomsky’s parse tree is based on the following CFG:

- Sentence  $\rightarrow$  NP VP
- NP  $\rightarrow$  the man
- NP  $\rightarrow$  the book
- VP  $\rightarrow$  Verb NP
- Verb  $\rightarrow$  took


The CFG’s terminal symbols are {the, man, took, book}, its non-terminal symbols are {Sentence, NP, VP, Verb}, and its starting symbol is “Sentence”. Now, we are coming back to the beginning of the section, where we mentioned that a CFG is usually thought of in two ways: as a generator **or** as a parser. As a generator, the example CFG might produce the following series of intermediate parse trees (only the last one will be submitted to the generator’s output):




Starting with the starting symbol, each of these intermediate parse trees is generated by applying one rule of the CFG to a suitable non-terminal leaf node of the previous parse tree, thereby adding the CFG rule as a local tree. The generator stops, if all leaf nodes of the current parse tree are terminal nodes. The whole generation process, of course, is non-deterministic, and this fact will lead us later on directly to probabilistic CFGs. As a parser, instead, the example CFG has to deal with an input sentence like

“the man took the book”


Usually, the parser starts processing the input sentence by assigning the words some local trees:



Then, the parser tries to add more local trees, by processing all the non-terminal nodes found in previous steps:



Doing this recursively, the parser provides us with a parse tree of the input sentence:



The example CFG is unambiguous for the given input sentence. Note, however, that this is far away from being the common situation. Usually, the parser stops, if all parse trees of the input sentence have been generated (and submitted to the output).

Now, we demonstrate that the fact that we can understand CFGs as generators leads directly to the **probabilistic context-free grammar** or **PCFG** formalism. As we already demonstrated for the generation process, the rules of the CFG serve as local trees that are incrementally used to build up a full parse tree (i.e. a parse tree without any non-terminal leaf nodes). This process, however, is non-deterministic: At most of its steps, some sort of **random choice** is involved that selects one of the different CFG rules which can potentially be appended to one of the non-terminal leaf nodes of the current parse tree<sup>9</sup>. Here is an example in the context of the generation process displayed above. For the CFG underlying Chomsky's very first parse tree, the non-terminal symbol NP is the left-hand side of two rules:

---

<sup>9</sup>Clearly, the final output of the generator is directly affected by the specific rule that has been selected by this random choice. Note also that there is another type of uncertainty in the generation process, playing, however, only a minor role: the specific place at which a CFG rule is to be appended does obviously not affect the generator's final output. So, these places can be deterministically chosen. For the generation process displayed above, for example, we decided to append the local trees always to the left-most non-terminal node of the actual partial-parse tree.

NP  $\rightarrow$  the man  
 NP  $\rightarrow$  the book


Clearly, when using the underlying CFG as a generator, we have to select either the first or the second rule, whenever a local NP tree shall be appended to the partial-parse tree given in the actual generation step. The choice might be either **fair** (both rules are chosen with probability 0.5) or **unfair** (the first rule is chosen, for example, with probability 0.9 and the second one with probability 0.1). In either case, a random choice between competing rules can be described by probability values which are directly allocated to the rules:

$$0 \leq p(\text{NP} \rightarrow \text{the man}) \leq 1 \quad \text{and} \quad 0 \leq p(\text{NP} \rightarrow \text{the book}) \leq 1$$


such that

$$p(\text{NP} \rightarrow \text{the man}) + p(\text{NP} \rightarrow \text{the book}) = 1$$

Now, having these probabilities at hand, it turns out that it is even possible to predict how often the generator will produce the one or the other of the following alternate partial-parse trees:



In turn, having this result at hand, we can also predict how often the generator will produce full-parse trees, for example, Chomsky's very first parse tree, or the parse tree of the sentence "the book took the book":




So, if  $p(\text{NP} \rightarrow \text{the man}) = 0.9$  and  $p(\text{NP} \rightarrow \text{the book}) = 0.1$ , then it is nine times more likely that the generator produces Chomsky's very first parse tree. In the following, we are trying to generalize this result even a bit more. As we saw, there are three rules in the CFG, which cause no problems in terms of uncertainty. These are:

Sentence  $\rightarrow$  NP VP  
 VP  $\rightarrow$  Verb NP  
 Verb  $\rightarrow$  took

To be more specific, we saw that these three rules have been always deterministically added to the partial-parse trees of the generation process. In terms of probability theory, determinism is expressed by the fact that certain events occur with a probability of one. In other words, a generator selects each of these rules with a probability of 100%, either when starting the generation process, or when expanding a VP or a Verb non-terminal node. So, we let

$$\begin{aligned} p(\text{Sentence} \rightarrow \text{NP VP}) &= 1 \\ p(\text{VP} \rightarrow \text{Verb NP}) &= 1 \\ p(\text{Verb} \rightarrow \text{took}) &= 1 \end{aligned}$$

The question is now: Have we won something by treating also the deterministic choices as probabilistic events? The answer is yes: A closer look at our example reveals that we can now predict easily how often the generator will produce a specific parse tree: The likelihood of a CFG's parse tree can be simply calculated as the product of the probabilities of all rules occurring in the tree. For example:



$$p(S \rightarrow \text{NP VP}) \cdot p(\text{NP} \rightarrow \text{the man}) \cdot p(\text{VP} \rightarrow \text{Verb NP}) \cdot p(\text{Verb} \rightarrow \text{took}) \cdot p(\text{NP} \rightarrow \text{the book})$$

To wrap up, we investigated the small CFG underlying Chomsky's very first parse tree. Motivated by the fact that a CFG can be used as a generator, we assigned each of its rules a weight (a non-negative real number) such that the weights of all rules with the same left-hand side sum up to one. In other words, all CFG fragments (comprising the CFG rules with the same left-hand side) have been assigned a probability distribution, as displayed in the following table:

CFG rule	Rule probability
Sentence $\rightarrow$ NP VP	$p(\text{Sentence} \rightarrow \text{NP VP}) = 1$
NP $\rightarrow$ the man	$p(\text{NP} \rightarrow \text{the man})$
NP $\rightarrow$ the book	$p(\text{NP} \rightarrow \text{the book})$
VP $\rightarrow$ Verb NP	$p(\text{VP} \rightarrow \text{Verb NP}) = 1$
Verb $\rightarrow$ took	$p(\text{Verb} \rightarrow \text{took}) = 1$

} **summing to 1**

As a result, the likelihood of each of the grammar's parse trees (when using the CFG as a generator) can be calculated by multiplying the probabilities of all rules occurring in the tree. This observation leads directly to the standard definition of a probabilistic context-free grammar, as well as to the definition of probabilities for parse-trees.

**Definition 16** A pair  $\langle G, p \rangle$  consisting of a context-free grammar  $G$  and a probability distribution  $p: \mathcal{X} \rightarrow [0, 1]$  on the set  $\mathcal{X}$  of all finite full-parse trees of  $G$  is called a **probabilistic context-free grammar** or **PCFG**, if for all parse trees  $x \in \mathcal{X}$

$$p(x) = \prod_{r \in G} p(r)^{f_r(x)}$$

Here,  $f_r(x)$  is the number of occurrences of the rule  $r$  in the tree  $x$ , and  $p(r)$  is a probability allocated to the rule  $r$ , such that for all non-terminal symbols  $A$

$$\sum_{r \in G_A} p(r) = 1$$

where  $G_A = \{ r \in G \mid \text{lhs}(r) = A \}$  is the grammar fragment comprising all rules with the left-hand side  $A$ . In other words, a probabilistic context-free grammar is defined by a context-free grammar  $G$  and some probability distributions on the grammar fragments  $G_A$ , thereby inducing a probability distribution on the set of all full-parse trees.

So far, we did not check for our example that the probabilities of all full-parse trees are summing up to one. According to Definition 16, however, this is **the fundamental property** of PCFGs (and it should be really checked for every PCFG which is accidentally given to us). Obviously, the example grammar has four full-parse trees, and the sum of their probabilities can be calculated as follows (by omitting all rules with a probability of one):

$$\begin{aligned} p(\mathcal{X}) &= p(\text{NP} \rightarrow \text{the man}) \cdot p(\text{NP} \rightarrow \text{the book}) + \\ &\quad p(\text{NP} \rightarrow \text{the book}) \cdot p(\text{NP} \rightarrow \text{the book}) + \\ &\quad p(\text{NP} \rightarrow \text{the man}) \cdot p(\text{NP} \rightarrow \text{the man}) + \\ &\quad p(\text{NP} \rightarrow \text{the book}) \cdot p(\text{NP} \rightarrow \text{the man}) \\ &= \left( p(\text{NP} \rightarrow \text{the man}) + p(\text{NP} \rightarrow \text{the book}) \right) \cdot p(\text{NP} \rightarrow \text{the book}) + \\ &\quad \left( p(\text{NP} \rightarrow \text{the man}) + p(\text{NP} \rightarrow \text{the book}) \right) \cdot p(\text{NP} \rightarrow \text{the man}) \\ &= 1 \end{aligned}$$

For the last equation, we are using three times that  $p$  is a probability distribution on the grammar fragment  $G_{\text{NP}}$ , i.e., we are exploiting that  $p(\text{NP} \rightarrow \text{the man}) + p(\text{NP} \rightarrow \text{the book}) = 1$ .


The following examples show that we really have to do this kind of “probabilistic grammar checking”. We are presenting two non-standard PCFGs: The first one consists of the rules

$$\begin{aligned} \text{S} &\rightarrow \text{NP sleeps} & (1.0) \\ \text{S} &\rightarrow \text{John sleeps} & (0.7) \\ \text{NP} &\rightarrow \text{John} & (0.3) \end{aligned}$$

The second one is a well-known highly-recursive grammar (Chi and Geman 1998), and it is given by

$$\begin{aligned} \text{S} &\rightarrow \text{S S} & (q) \\ \text{S} &\rightarrow \text{a} & (1-q) \end{aligned} \quad \text{with } 0.5 < q \leq 1$$

What is wrong with these grammars? Well, the first grammar provides us with a probability distribution on its full-parse trees, as can be seen here



On each of its grammar fragments, however, the rule probabilities do not form a probability distribution (neither on  $G_S$  nor on  $G_{NP}$ ). The second grammar is even worse: We do have a probability distribution on  $G_S$ , but even so, we do not have a probability distribution on the set of full-parse trees (because their probabilities are summing to less than one<sup>10</sup>). In a certain sense, however, we are able to repair both grammars. For example,

$$\begin{aligned} S &\rightarrow NP \text{ sleeps} && (0.3) \\ S &\rightarrow \text{John sleeps} && (0.7) \\ NP &\rightarrow \text{John} && (1.0) \end{aligned}$$

is the standard-PCFG counterpart of the first grammar, where

$$\begin{aligned} S &\rightarrow S S && (1-q) \\ S &\rightarrow a && (q) \end{aligned} \quad \text{with } 0.5 < q \leq 1$$

is a standard-PCFG counterpart of the second grammar: The first grammar and its counterpart provide us with exactly the same parse-tree probabilities, while the second grammar and its counterpart produce parse-tree probabilities, which are proportional to each other. Especially for the second example, this interesting result is a special case of an important general theorem recently proven by Nederhof and Satta (2003). Sloppily formulated, their Theorem 7 states that:

There is for each **weighted CFG** a standard PCFG with the same symbolic backbone, such that (i) the parse-tree probabilities (produced by the PCFG) are summing to one, and (ii) the parse-tree weights (produced by the weighted CFG) are proportional to the parse-tree probabilities (produced by the PCFG).

As a consequence, we are getting what we really want: Applied to ambiguity resolution, the original grammars and their counterparts provide us with exactly the same maximum-probability-parse trees.

## Background: Resolving Ambiguities with PCFGs

A property of most formalizations of natural language in terms of CFGs is **ambiguity**: the fact that sentences have more than one possible phrase structure (and therefore more than one meaning). Here are two prominent types of ambiguity:

---


<sup>10</sup>This can be proven as follows: Let  $\mathcal{T}$  be the set of all finite full-parse trees that can be generated by the given context-free grammar. Then, it is easy to verify that  $\pi := p(\mathcal{T})$  is a solution of the following equation

$$\pi = 1 - q + q \cdot \pi \cdot \pi$$

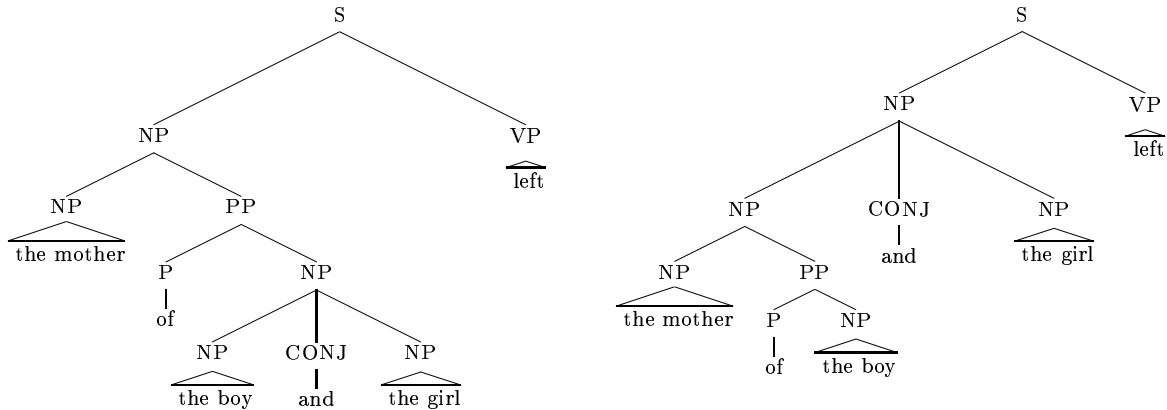
Here,  $1 - q$  is the probability of the tree  $\begin{array}{c} S \\ | \\ a \end{array}$ , whereas  $q \cdot \pi \cdot \pi$  corresponds to the forest  $\begin{array}{c} S \\ / \quad \backslash \\ \mathcal{T} \quad \mathcal{T} \end{array}$ .

It is easy to check that the derived quadratic equation has two solutions:  $\pi_1 = 1$  and  $\pi_2 = \frac{1-q}{q}$ . Note that it is quite natural that **two** solutions arise: The set of all “infinite full-parse trees” matches also our under-specified approach of calculating  $\pi$ . Now, in the case of  $0.5 < q \leq 1$ , it turns out that the set of infinite trees is allocated a proper probability  $\pi_1 = 1$ . (For the special case  $q = 1$ , this can be intuitively verified: The generator will never touch the rule  $S \rightarrow a$ , and therefore, this special PCFG produces infinite parse trees only.) As a consequence, all finite full-parse trees is allocated the total probability  $\pi_2$ . In other words,  $p(\mathcal{T}) = \frac{1-q}{q} < 1$ .


*Ambiguity caused by prepositional-phrase attachment:*




*Ambiguity caused by conjunctions:*



As usual in computational linguistics, some phrase structures have been displayed in abbreviated form: For example, the term  $\begin{array}{c} \text{NP} \\ \text{the mother} \end{array}$  is used as a short form for the parse



, and the term  $\begin{array}{c} \text{PP} \\ \text{of the boy} \end{array}$  is a place holder for the even more




In both examples, the ambiguity is caused by the fact that the underlying CFG contains **recursive rules**, i.e., rules that can be applied an arbitrary number of times. Clearly, the rules  $\text{NP} \rightarrow \text{NP CONJ NP}$  and  $\text{NP} \rightarrow \text{NP PP}$  belong to this type, since they can be used to generate nominal phrases of an arbitrary length. The rules  $\text{VP} \rightarrow \text{V NP}$  and  $\text{PP} \rightarrow \text{P NP}$ , however, might be also called (indirectly) recursive, since they can generate verbal and prepositional phrases of an arbitrary length (in combination with  $\text{NP} \rightarrow \text{NP PP}$ ). Besides ambiguity, recursivity makes it also possible that two words that are generated by the same CFG rule (i.e. which are syntactically linked) can occur far apart in a sentence:

The **bird** with the nice brown eyes and the beautiful tail feathers **catches** a worm.

These types of phenomena are called **non-local dependencies**, and it is important to note that non-local phenomena (which can be handled by CFGs) are beyond the scope of many popular models that focus on modeling local dependencies (such as n-gram, Markov, and hidden Markov models<sup>11</sup>). So, a part-of-speech tagger (based on a HMM model) might have difficulties with sentences like the one we mentioned, because it will not expect that a singular verb occurs after a plural noun.

Having this at hand, of course, the central question is now: Can PCFGs handle ambiguity? The somewhat surprising answer is: Yes, but the symbolic backbone of the PCFG plays a major role in solving this difficult task. To be a bit more specific, the CFG underlying the given PCFG has to have some good properties, or the other way round, probabilistic modeling of some “weak” CFGs may result in PCFGs which can not resolve the CFG’s ambiguities. From a probabilistic modeler’s point of view, there is really some non-trivial relation between such tasks as “writing a formal grammar” and “modeling a probabilistic grammar”. So, we are convinced that formal-grammar writers should help probabilistic-grammar modelers, and the other way round.

To exemplify this, we will have a closer look at the examples above, where we presented two common types of ambiguity. In general, a PCFG resolves ambiguity (i) by calculating all the full parse-trees of a given sentence (using the symbolic backbone of the CFG), and (ii) by allocating probabilities to all these trees (using the rule probabilities of the PCFG), and finally (iii) by choosing the most probable parse as the analysis of the given sentence. According to this procedure, we are calculating, for example



$$\begin{aligned}
 & p(S \rightarrow NP VP) \cdot p\left(\frac{NP}{\text{Peter}}\right) \cdot \\
 & p(VP \rightarrow V NP PP) \cdot \\
 & p(V \rightarrow \text{saw}) \cdot p\left(\frac{NP}{\text{Mary}}\right) \cdot p\left(\frac{PP}{\text{with a telescope}}\right)
 \end{aligned}$$

$$\begin{aligned}
 & p(S \rightarrow NP VP) \cdot p\left(\frac{NP}{\text{Peter}}\right) \cdot \\
 & p(VP \rightarrow V NP) \cdot p(NP \rightarrow NP PP) \cdot \\
 & p(V \rightarrow \text{saw}) \cdot p\left(\frac{NP}{\text{Mary}}\right) \cdot p\left(\frac{PP}{\text{with a telescope}}\right)
 \end{aligned}$$

Comparing the probabilities for these two analyzes, we are choosing the analysis at the left-hand side of this figure, if

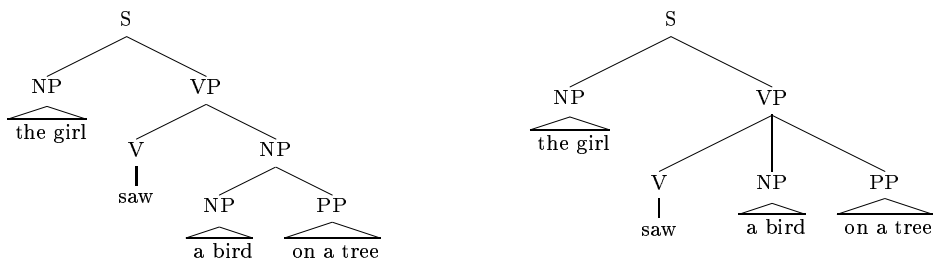
$$p(VP \rightarrow V NP PP) > p(VP \rightarrow V NP) \cdot p(NP \rightarrow NP PP)$$

<sup>11</sup>It is well-known, however, that a CFG without **center-embeddings** can be transformed to a regular grammar (the symbolic backbone of a hidden Markov model).

So, in principle, the PP-attachment ambiguity encoded in this CFG **can** be solved by a probabilistic model built on top of this CFG. Moreover, it is especially nice that such a PCFG resolves the ambiguity by looking only at those rules of the CFG, which directly cause the PP-attachment ambiguity.

So far, so good: We are able to use PCFGs in order to select between competing analyzes of a sentence. Looking at the second example (ambiguity caused by conjunctions), however, we are faced with a serious problem: Both trees have a different structure, but exactly the same context-free rules are used for generating these different structures. As a consequence, both trees are allocated the same probability (independently from the specific rule probabilities which might have been offered to you by the very best estimation methods). So, any PCFG based on the given CFG is unable to resolve the ambiguity manifested in the two trees.

Here is another problem. Using the grammar underlying our first example, the sentence “the girl saw a bird on a tree” has the following two analyzes



Comparing the probabilities for these two analyzes, we are choosing the analysis at the left-hand side of this figure, if

$$p(\mathbf{VP} \rightarrow \mathbf{V NP}) \cdot p(\mathbf{NP} \rightarrow \mathbf{NP PP}) > p(\mathbf{VP} \rightarrow \mathbf{V NP PP})$$

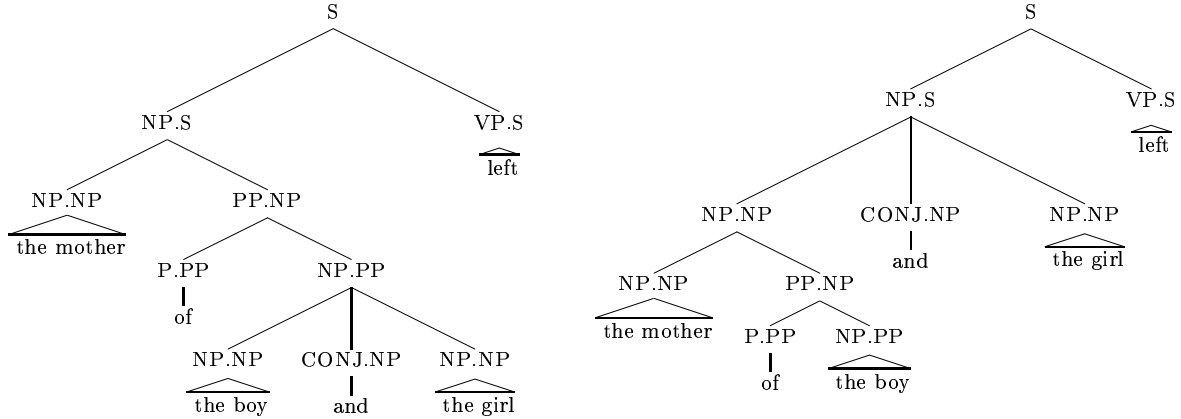
Relating this result to the disambiguation result of the sentence “John saw Mary with a telescope”, the prepositional phrases are attached in both cases either to the verbal phrase or to the nominal phrase. Instead, it seems to be more plausible that the PP “with the telescope” is attached to the verbal phrase, whereas the PP “on a tree” is attached to the noun phrase.

Obviously, there is only one possible solution to this problem: We have to re-write the given CFG in such a way that a probabilistic model will be able to assign different probabilities to different analyzes. For our last example, it is sufficient to enrich the underlying CFG with some simple PP markup, enabling in principle that

$$\begin{aligned} p(\mathbf{VP} \rightarrow \mathbf{V NP PP-ON}) &< p(\mathbf{VP} \rightarrow \mathbf{V NP}) \cdot p(\mathbf{NP} \rightarrow \mathbf{NP PP-ON}) \\ p(\mathbf{VP} \rightarrow \mathbf{V NP PP-WITH}) &> p(\mathbf{VP} \rightarrow \mathbf{V NP}) \cdot p(\mathbf{NP} \rightarrow \mathbf{NP PP-WITH}) \end{aligned}$$

Of course, other linguistically more sophisticated modifications of the original CFG (that handle e.g. agreement information, sub-cat frames, selectional preferences, etc) are also welcome. Our only request is that the modified CFGs lead to PCFGs which are able to resolve the different types of ambiguities encoded in the original CFG. Now, writing and re-writing a formal grammar is a job that grammar writers can do probably much better than modelers of probabilistic grammars. In the past, however, writers of formal grammars seemed to be uninterested in this specific task, or they are still unaware of its existence. So, modelers of PCFGs regard it nowadays also as an important part of their job to transform a given CFG

in such a way that probabilistic versions of the modified CFG are able to resolve ambiguities of the original CFG. During the last years, a bunch of automatic grammar-transformation techniques have been developed, which offer some interesting solutions to this quite complex problem. Where the work of Klein and Manning (2003) describes one of the latest approaches to semi-automatic grammar-transformation, the **parent-encoding technique** introduced by Johnson (1998) is the earliest and purely automatic grammar-transformation technique: For each local tree, the parent's category is appended to all daughter categories. Using the example above, where we showed that conjunctions cause ambiguities, the parent-encoded trees are looking as follows:



Clearly, parent-encoding of the original trees may result in different probabilities of the transformed trees: In this example, we will choose the analysis at the left-hand side, if

$$p(\mathbf{NP.S} \rightarrow \mathbf{NP.NP} \mathbf{PP.NP}) \cdot p(\mathbf{NP.PP} \rightarrow \mathbf{NP.NP} \mathbf{CONJ.NP} \mathbf{NP.NP}) \cdot p\left(\frac{\mathbf{NP.NP}}{\text{the boy}}\right)$$

is more likely than

$$p(\mathbf{NP.NP} \rightarrow \mathbf{NP.NP} \mathbf{PP.NP}) \cdot p(\mathbf{NP.S} \rightarrow \mathbf{NP.NP} \mathbf{CONJ.NP} \mathbf{NP.NP}) \cdot p\left(\frac{\mathbf{NP.PP}}{\text{the boy}}\right)$$

As in the example before, it is again nice to see that these probabilities are pin-pointing exactly at those rules of the underlying grammar which have introduced the ambiguity.

In the rest of the section, we will present the notion of treebank grammars, which can be informally described as PCFGs that are constructed on the basis of a corpus of full-parse trees (Charniak 1996). We will demonstrate that treebank grammars can resolve the ambiguous sentences of the treebank (as well as ambiguous similar sentences), if the treebank mark-up is rich enough to distinguish between the different types of ambiguities that are encoded in the treebank.

**Definition 17** For a given *treebank*, i.e., for a non-empty and finite corpus of full-parse trees, the *treebank grammar*  $\langle G, p \rangle$  is a PCFG defined by

- (i)  $G$  is the context-free grammar read off from the treebank, and


(ii)  $p$  is the probability distribution on the set of full-parse trees of  $G$ , induced by the following specific probability distributions on the grammar fragments  $G_A$ :

$$p(r) = \frac{f(r)}{\sum_{r \in G_A} f(r)} \quad \text{for all } r \in G_A$$

Here,  $f(r)$  is the number of times a rule  $r \in G$  occurs in the treebank.

**Note:** Later on (see Theorem 10), we will show that each treebank grammar is the unique maximum-likelihood estimate of  $G$ 's probability model on the given treebank. So, it is especially guaranteed that  $p$  is a probability distribution on the set of full-parse trees of  $G$ , and that  $\langle G, P \rangle$  is a standard PCFG (see Definition 16).

**Example 5** We shall now consider a treebank given by the following 210 full-parse trees:



We shall (i) generate the treebank grammar and (ii) using this treebank grammar, we shall resolve the ambiguities of the sentences occurring in the treebank.

Ad (i): The following table displays the rules  $r$  of the CFG encoded in the given treebank (thereby assuming for the sake of simplicity that the  $NP$  and  $PP$  non-terminals expand directly to terminal symbols), the rule frequencies  $f(r)$  i.e. the number of times a rule occurs in the treebank, as well as the rule probabilities  $p(r)$  as defined in Definition 17.

CFG rule	Rule frequency	Rule probability
S $\rightarrow$ NP VP	100 + 5 + 100 + 5	$\frac{210}{210} = 1.000$
VP $\rightarrow$ V NP PP-WITH	100	$\frac{100}{210} \approx 0.476$
VP $\rightarrow$ V NP PP-ON	5	$\frac{5}{210} \approx 0.024$
VP $\rightarrow$ V NP	5 + 100	$\frac{105}{210} = 0.500$
NP $\rightarrow$ Peter	100 + 5	$\frac{105}{525} = 0.200$
NP $\rightarrow$ Mary	100 + 5 + 100 + 5	$\frac{210}{525} = 0.400$
NP $\rightarrow$ a bird	100 + 5	$\frac{105}{525} = 0.200$
NP $\rightarrow$ NP PP-WITH	5	$\frac{5}{525} \approx 0.010$
NP $\rightarrow$ NP PP-ON	100	$\frac{100}{525} \approx 0.190$
PP-WITH $\rightarrow$ with a telescope	100 + 5	$\frac{105}{105} = 1.000$
PP-ON $\rightarrow$ on a tree	100 + 5	$\frac{105}{105} = 1.000$
V $\rightarrow$ saw	100 + 5 + 100 + 5	$\frac{210}{210} = 1.000$

Ad (ii): As we have already seen, the treebank grammar selects the full-parse tree  $t_1$  of the sentence “Peter saw Mary with a telescope”, if

$$p(\text{VP} \rightarrow \text{V NP PP-WITH}) > p(\text{VP} \rightarrow \text{V NP}) \cdot p(\text{NP} \rightarrow \text{NP PP-WITH})$$

Using the approximate probabilities for these rules, this is indeed true:  $0.476 > 0.500 \cdot 0.010$ . (Note that exactly the same argument can be applied to similar but more complex sentences like “Peter saw a bird on a tree with a telescope”.) For the second sentence occurring in the treebank, “Mary saw a bird on a tree”, the treebank grammar selects the full-parse tree  $t_3$ , if

$$p(\text{VP} \rightarrow \text{V NP PP-ON}) < p(\text{VP} \rightarrow \text{V NP}) \cdot p(\text{NP} \rightarrow \text{NP PP-ON})$$

Indeed, this is the case:  $0.024 < 0.500 \cdot 0.190$ .

## Maximum-Likelihood Estimation of PCFGs

So far, we have seen that probabilistic context-free grammars can be used to resolve the ambiguities that are caused by their underlying context-free backbone, and we noted already that certain grammar-transformations are sometimes necessary to achieve this goal. All these application features are displayed in a “horizontal view” in Figure 1. In what follows next, we will concentrate on the “vertical view” of this figure. To be more specific, we will focus on two questions. First, how to characterize the probability model of a given context-free grammar, and second, how to estimate an appropriate instance of the context-free grammar’s probability model, if a corpus of input data is also given. The latter question is a tough one: It is true that the treebank-training method, which we defined in the previous section more or less heuristically, leads to PCFGs that are able to resolve ambiguities. Right now, however, we have no clear idea how the treebank-training method is related to the maximum-likelihood estimation and the EM training. On the other hand, we are well-prepared to answer the first question.

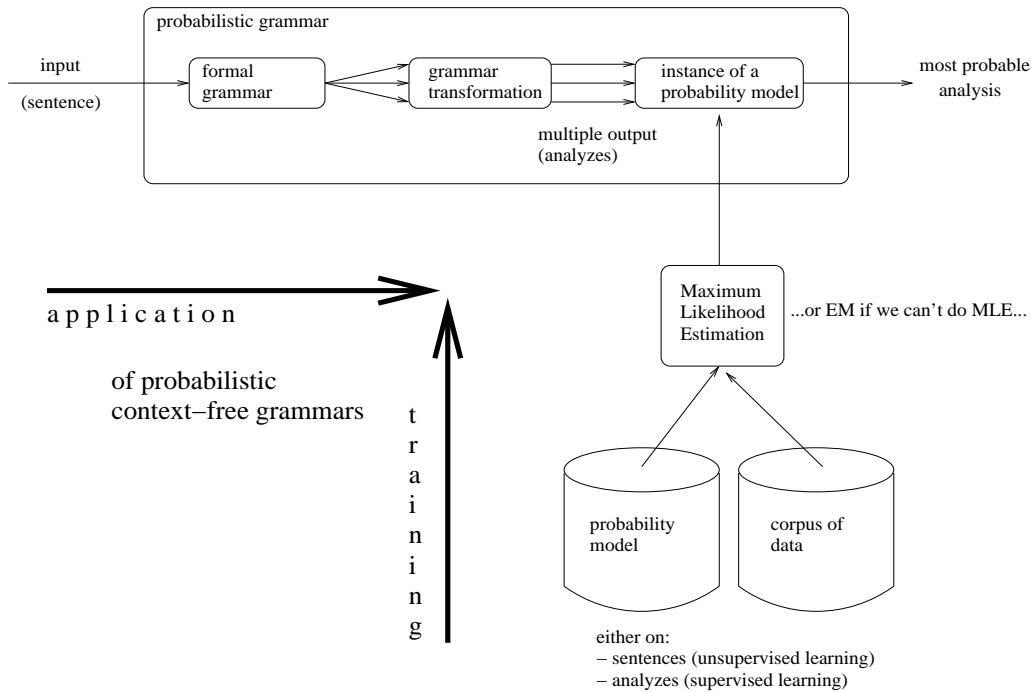


Figure 1: Application and training of probabilistic context-free grammars

**Definition 18** Let  $G$  be a context-free grammar, and let  $\mathcal{X}$  be the set of full-parse trees of  $G$ . Then, the **probability model of  $G$**  is defined by

$$\mathcal{M}_G = \left\{ p \in \mathcal{M}(\mathcal{X}) \mid p(x) = \prod_{r \in G} p(r)^{f_r(x)} \text{ with } \sum_{r \in G_A} p(r) = 1 \text{ for all grammar fragments } G_A \right\}$$

In other words, each instance  $p$  of the probability model  $\mathcal{M}_G$  is associated to a probabilistic context-free grammar  $\langle G, p \rangle$  having the context-free backbone  $G$ . (See Definition 16 for the meaning of the terms  $p(r)$ ,  $f_r(x)$  and  $G_A$ .)

As we have already seen, there are some non-standard PCFGs (like  $s \rightarrow s s$  (0.9),  $s \rightarrow a$  (0.1)) which do not induce a probability distribution on the set of full-parse trees. So, we know already that there are some non-standard PCFGs  $\langle G, p \rangle$  such that  $p \notin \mathcal{M}(\mathcal{X})$ . This gives us a rough idea that it might be quite difficult to characterize those PCFGs  $\langle G, p \rangle$  which are associated to an instance  $p$  of the unrestricted probability model  $\mathcal{M}(\mathcal{X})$ . In other words, it might be quite difficult to characterize  $G$ 's probability model  $\mathcal{M}_G$ . For calculating a maximum-likelihood estimate of  $\mathcal{M}_G$  on a corpus  $f_{\mathcal{T}}$  of full-parse trees, however, we have to solve this task. For example, if we are targeting to exploit the powerful Theorem 1, we have to prove either that  $\mathcal{M}_G$  equals the unrestricted probability model  $\mathcal{M}(\mathcal{X})$ , or that the relative-frequency estimate  $\tilde{p}$  on  $f_{\mathcal{T}}$  is an instance of  $\mathcal{M}_G$ . For most context-free grammars  $G$ , however, the following theorems show that this is a too simplistic approach of finding a maximum-likelihood estimate of  $\mathcal{M}_G$ .

**Theorem 8** Let  $G$  be a context-free grammar, and let  $\tilde{p}$  be the relative-frequency estimate on a non-empty and finite corpus  $f_{\mathcal{T}}: \mathcal{X} \rightarrow \mathcal{R}$  of full-parse trees of  $G$ . Then  $\tilde{p} \notin \mathcal{M}_G$  if

- (i)  $G$  can be read off from  $f_{\mathcal{T}}$ , and
- (ii)  $G$  has a full-parse tree  $x_{\infty} \in \mathcal{X}$  that is not in  $f_{\mathcal{T}}$ .

**Proof** Assume that  $\tilde{p} \in \mathcal{M}_G$ . In what follows, we will show that this assumption leads to a contradiction. First, by definition of  $\mathcal{M}_G$ , it follows that there are some weights  $0 \leq \pi(r) \leq 1$  such that

$$\tilde{p}(x) = \prod_{r \in G} \pi(r)^{f_r(x)} \quad \text{for all } x \in \mathcal{X}$$

We will show next that  $\pi(r) > 0$  for all  $r \in G$ .

Assume that there is a rule  $r_0 \in G$  with  $\pi(r_0) = 0$ . By (i),  $G$  can be read off from  $f_{\mathcal{T}}$ . So,  $f_{\mathcal{T}}$  contains a full-parse tree  $x_0$  such that the rule  $r_0$  occurs in  $x_0$ , i.e.,

$$f_{\mathcal{T}}(x_0) > 0 \quad \text{and} \quad f_{r_0}(x_0) > 0$$

It follows both  $\tilde{p}(x_0) = \frac{f_{\mathcal{T}}(x_0)}{|f_{\mathcal{T}}|} > 0$  and  $\tilde{p}(x_0) = \prod_{r \in G} \pi(r)^{f_r(x_0)} = \dots \pi(r_0)^{f_{r_0}(x_0)} \dots = 0$ , which is a contradiction.

Therefore

$$\tilde{p}(x) = \prod_{r \in G} \pi(r)^{f_r(x)} > 0 \quad \text{for all } x \in \mathcal{X}$$


On the other hand by (ii), there is the full-parse tree  $x_{\infty} \in \mathcal{X}$  which is not in  $f_{\mathcal{T}}$ . So,  $\tilde{p}(x_{\infty}) = \frac{f_{\mathcal{T}}(x_{\infty})}{|f_{\mathcal{T}}|} = 0$ , which is a contradiction to the last inequality **q.e.d.**

**Example 6** The treebank grammar described in Example 5 is a context-free grammar of the type described in Theorem 8.

The relative-frequency estimate  $\tilde{p}$  on the treebank is given by:

$$\tilde{p}(t_1) = \tilde{p}(t_3) = \frac{100}{210} \quad \text{and} \quad \tilde{p}(t_2) = \tilde{p}(t_4) = \frac{5}{210}$$

All other full-parse trees of the treebank grammar get allocated a probability of zero by the relative-frequency estimate. So, for example,  $\tilde{p}(x_{\infty}) = 0$  for



As a consequence,  $\tilde{p}$  can not be an instance of the probability model of the treebank grammar: Otherwise,  $\tilde{p}$  would allocate both full-parse trees  $t_1$  and  $x_{\infty}$  exactly the same probabilities (because  $t_1$  and  $x_{\infty}$  contain exactly the same rules).

**Theorem 9** *For each context-free grammar  $G$  with an infinite set  $\mathcal{X}$  of full-parse trees, the probability model of  $G$  is restricted*

$$\mathcal{M}_G \neq \mathcal{M}(\mathcal{X})$$

**Proof** First, each context-free grammar consists of a finite number of rules. Thus it is possible to construct a treebank, such that  $G$  is encoded by the treebank. (Without loss of generality, we are assuming here that all non-terminal symbols of  $G$  are reachable and productive.) So, let  $f_{\mathcal{T}}$  be a non-empty and finite corpus of full-parse trees of  $G$  such that  $G$  can be read off from  $f_{\mathcal{T}}$ , and let  $\tilde{p}$  be the relative-frequency estimate on  $f_{\mathcal{T}}$ . Second, using  $|\mathcal{X}| = \infty$  and  $|f_{\mathcal{T}}| < \infty$ , it follows that there is at least one full-parse tree  $x_{\infty} \in \mathcal{X}$  which is not in  $f_{\mathcal{T}}$ . So, the conditions of Theorem 8 are met, and we are concluding that  $\tilde{p} \notin \mathcal{M}_G$ . On the other hand,  $\tilde{p} \in \mathcal{M}(\mathcal{X})$ . So, clearly,  $\mathcal{M}_G \neq \mathcal{M}(\mathcal{X})$ . In other words,  $\mathcal{M}_G$  is a restricted probability model **q.e.d.**

After all, we might recognize that the previous results are not bad. Yes, the probability model of a given CFG is restricted in most of the cases. The missing distributions, however, are the relative-frequency estimates on each treebank encoding the given CFG. These relative-frequency estimates lack the ability of any generalization power: They allocate each full-parse tree not being in the treebank a zero-probability. Obviously, however, we surely want a probability model that can be learned by maximum-likelihood estimation on a corpus of full-parse trees, but that is at the same time able to deal with full-parse trees not seen in the treebank. The following theorem shows that we have already found one.

**Theorem 10** *Let  $f_{\mathcal{T}}: \mathcal{X} \rightarrow \mathcal{R}$  be a non-empty and finite corpus of full-parse trees, and let  $\langle G, p_{\mathcal{T}} \rangle$  be the treebank grammar read off from  $f_{\mathcal{T}}$ . Then,  $p_{\mathcal{T}}$  is a maximum-likelihood estimate of  $\mathcal{M}_G$  on  $f_{\mathcal{T}}$ , i.e.,*

$$L(f_{\mathcal{T}}; p_{\mathcal{T}}) = \max_{p \in \mathcal{M}_G} L(f_{\mathcal{T}}; p)$$

*Moreover, maximum-likelihood estimation of  $\mathcal{M}_G$  on  $f_{\mathcal{T}}$  yields a unique estimate.*

This theorem is well-known. The following proof, however, is especially simple and (to the best of my knowledge) was given first by Prescher (2002).

**Proof** First step: We will show that for all model instances  $p \in \mathcal{M}_G$

$$L(f_{\mathcal{T}}; p) = L(f; p)$$

At the right-hand side of this equation,  $f$  refers to the corpus of rules that are read off from the treebank  $f_{\mathcal{T}}$ , i.e.,  $f(r)$  is the number of times a rule  $r \in G$  occurs in the treebank; Similarly,  $p$  refers to the probabilities  $p(r)$  of the rules  $r \in G$ . In contrast, at the left-hand side of the equation,  $p$  refers to the probabilities  $p(x)$  of the full-parse trees  $x \in \mathcal{X}$ . The proof of the equation is relatively straight-forward:

$$\begin{aligned} L(f_{\mathcal{T}}; p) &= \prod_{x \in \mathcal{X}} p(x)^{f_{\mathcal{T}}(x)} \\ &= \prod_{x \in \mathcal{X}} \left( \prod_{r \in G} p(r)^{f_r(x)} \right)^{f_{\mathcal{T}}(x)} \end{aligned}$$

$$\begin{aligned}
&= \prod_{x \in \mathcal{X}} \prod_{r \in G} p(r)^{f_{\mathcal{T}}(x) \cdot f_r(x)} \\
&= \prod_{r \in G} \prod_{x \in \mathcal{X}} p(r)^{f_{\mathcal{T}}(x) \cdot f_r(x)} \\
&= \prod_{r \in G} p(r)^{\sum_{x \in \mathcal{X}} f_{\mathcal{T}}(x) \cdot f_r(x)} \\
&= \prod_{r \in G} p(r)^{f(r)} \\
&= L(f; p)
\end{aligned}$$

In the 6<sup>th</sup> equation, we simply used that  $f(r)$  (the number of times a specific rule occurs in the treebank) can be calculated by summing up all the  $f_r(x)$  (the number of times this rule occurs in a specific full-parse tree  $x \in \mathcal{X}$ ):

$$f(r) = \sum_{x \in \mathcal{X}} f_{\mathcal{T}}(x) \cdot f_r(x)$$

So, maximizing  $L(f_{\mathcal{T}}; p)$  is equivalent to maximizing  $L(f; p)$ . Unfortunately, Theorem 1 can not be applied to maximize the term  $L(f; p)$ , because the rule probabilities  $p(r)$  do not form a probability distribution on the set of all grammar rules. They do form, however, probability distributions on the grammar fragments  $G_A$ . So, we have to refine our result a bit more.

Second step: We are showing here that

$$L(f; p) = \prod_A L(f_A; p)$$

Here, each  $f_A$  is a corpus of rules, read off from the given treebank, thereby filtering out all rules not having the left-hand side  $A$ . To be specific, we define

$$f_A(r) = \begin{cases} f(r) & \text{if } r \in G_A \\ 0 & \text{else} \end{cases}$$

Again, the proof is easy:

$$\begin{aligned}
L(f; p) &= \prod_{r \in G} p(r)^{f(r)} \\
&= \prod_A \prod_{r \in G_A} p(r)^{f(r)} \\
&= \prod_A \prod_{r \in G_A} p(r)^{f_A(r)} \\
&= \prod_A L(f_A; p)
\end{aligned}$$

Third step: Combining the first and second step, we conclude that

$$L(f_{\mathcal{T}}; p) = \prod_A L(f_A; p)$$

So, maximizing  $L(f_{\mathcal{T}}; p)$  is equivalent to maximizing  $\prod_A L(f_A; p)$ . Now, a product is maximized, if all its factors are maximized. So, in what follows, we are focusing on how to maximize

the terms  $L(f_A; p)$ . First of all, the corpus  $f_A$  comprises only rules with the left-hand side  $A$ . So, the value of  $L(f_A; p)$  depends only on the values  $p(r)$  of the rules  $r \in G_A$ . These values, however, form a probability distribution on  $G_A$ , and all these probability distributions on  $G_A$  have to be considered for maximizing  $L(f_A; p)$ . It follows that we have to calculate an instance  $\hat{p}_A$  of the unrestricted probability model  $\mathcal{M}(G_A)$  such that

$$L(f_A; \hat{p}_A) = \max_{p \in \mathcal{M}(G_A)} L(f_A; p)$$

In other words, we have to calculate a maximum-likelihood estimate of the unrestricted probability model  $\mathcal{M}(G_A)$  on the corpus  $f_A$ . Fortunately, this task can be easily solved. According to Theorem 1, the relative-frequency estimate on the corpus  $f_A$  is our unique solution. This yields for the rules  $r \in G_A$

$$\hat{p}_A(r) = \frac{f_A(r)}{|f_A|} = \frac{f(r)}{\sum_{r \in G_A} f(r)}$$

Comparing these formulas to the formulas given in Definition 17, we conclude that for all non-terminal symbols  $A$

$$\hat{p}_A(r) = p_{\mathcal{T}}(r) \quad \text{for all } r \in G_A$$

So, clearly, the treebank grammar  $p_{\mathcal{T}}$  is a serious candidate for a maximum-likelihood estimate of the probability model  $\mathcal{M}_G$  on  $f_{\mathcal{T}}$ . Now, as Chi and Geman (1998) verified, the treebank grammar  $p_{\mathcal{T}}$  is indeed an instance of the probability model  $\mathcal{M}_G$ . So, combining all the results, it follows that

$$L(f_{\mathcal{T}}; p_{\mathcal{T}}) = \max_{p \in \mathcal{M}_G} L(f_{\mathcal{T}}; p)$$

Finally, since all  $\hat{p}_A \in \mathcal{M}(G_A)$  are unique maximum-likelihood estimates,  $p_{\mathcal{T}} \in \mathcal{M}_G$  is also a unique maximum-likelihood estimate **q.e.d.**

## EM Training of PCFGs

We are starting this section with an overview of some theoretical and empirical work on using the EM algorithm for training of probabilistic context-free grammars.

- From 1966 to 1972, a group around Leonard E. Baum invented the forward-backward algorithm for probabilistic regular grammars (or hidden Markov models) and formally proved that this algorithm has some good convergence properties. See, for example, the overview presented in Baum (1972).
- Booth and Thompson (1973) discovered a (still nameless) constraint for PCFGs. For those PCFGs fulfilling the constraint, the probabilities of all full-parse trees sum to one.
- Baker (1979) invented the inside-outside algorithm (as a generalization of the forward-backward algorithm). The training corpus of this algorithm, however, is not allowed to contain more than one single sentence.
- Lari and Young (1990) generalized Baker's inside-outside algorithm (or in other words, they invented the modern form of the inside-outside algorithm): The training corpus of their algorithm may contain arbitrary many sentences.

- Pereira and Schabes (1992) used the inside-outside algorithm for estimating a PCFG for English on a partially bracketed corpus.
- Sanchez and Benedi (1997) and Chi and Geman (1998) formally proved that for treebank grammars and grammars estimated by the EM algorithm, the probabilities of all full-parse trees sum to one.
- Beil et al. (1999) used the inside-outside algorithm for estimating a PCFG for German on a corpus of sentences. The context-free backbone of this PCFG is a manually written grammar fragment of German, containing several thousands of rules.
- Prescher (2001) formally proved that the inside-outside algorithm is a dynamic programming instance of the EM algorithm for PCFGs. So, the inside-outside algorithm inherits all the convergence properties of the general EM algorithm (no formal proofs of these properties have been given by Baker and Lari and Young).
- Nederhof and Satta (2003) discovered that the PCFG standard-constraints (“the probabilities of the rules with the same left-hand side are summing to one”) are dispensable<sup>12</sup>.

Clearly, this overview points to the fact that the forward-backward algorithm is a special instance of the inside-outside algorithm, and moreover that both algorithms can be regarded as simple instances of the EM algorithm. In the following, however, we will not focus on the affinity of these three algorithms. Moreover, in the light of the early results of Booth and Thompson (1973) and the late results of Nederhof and Satta (2003), it might be worthy to reflect on our notion of a standard probability model for context-free grammars. Surely, this topic is a highly interesting one, and some fruitful results may come out of it. Here, however, we would like to limit our selfs and to refer the interested reader to some of the references given above. The rest of this paper is dedicated to the EM algorithm for PCFGs. We would like to present its procedure and its properties, thereby motivating that the EM algorithm can be successfully used to train a manually written grammar on a plain text corpus. The following theorem shows that the EM algorithm is not only related to the inside-outside algorithm, but that the EM algorithm is also strongly connected with the treebank-training method on which we we have focused so far.

**Theorem 11** *Let  $\langle G, p_0 \rangle$  be a probabilistic context-free grammar, where  $p_0$  is an arbitrary starting instance of the probability model  $\mathcal{M}_G$ . Let  $f : \mathcal{Y} \rightarrow \mathcal{R}$  be a non-empty and finite corpus of **sentences of  $G$**  (terminal strings that have at least one full-parse tree  $x \in \mathcal{X}$ ). Then, applied to the PCFG  $\langle G, p_0 \rangle$  and the sentence corpus  $f$ , the EM Algorithm performs the following procedure*

---

<sup>12</sup>In our terms, their result can be expressed as follows. Let  $G$  be a context-free grammar, and let  $\mathcal{M}_G^*$  be the probability model that disregards the PCFG standard-constraints

$$\mathcal{M}_G^* = \left\{ p \in \mathcal{M}(\mathcal{X}) \mid p(x) = \prod_{r \in G} p(r)^{f_r(x)} \right\}$$

Obviously, it follows then that  $\mathcal{M}_G \subseteq \mathcal{M}_G^*$ . Exploiting their Corollary 8, however, it follows somewhat surprisingly that both models are even equal:  $\mathcal{M}_G^* = \mathcal{M}_G$ . As a trivial consequence, a maximum-likelihood estimate of the standard probability model  $\mathcal{M}_G$  (on a corpus of trees or on a corpus of sentences) is also a maximum-likelihood estimate of the probability model  $\mathcal{M}_G^*$  — as well as the other way round.

- (1) for each  $i = 1, 2, 3, \dots$  do
- (2)  $q := p_{i-1}$
- (3) **E-step (PCFGs):** generate the **treebank**  $f_{\mathcal{T}_q} : \mathcal{X} \rightarrow \mathcal{R}$  defined by
$$f_{\mathcal{T}_q}(x) := f(y) \cdot q(x|y) \quad \text{where } y = \text{yield}(x)$$
- (4) **M-step (PCFGs):** read off the **treebank grammar**  $\langle G, p_{\mathcal{T}_q} \rangle$
- (5)  $p_i := p_{\mathcal{T}_q}$
- (6) end // for each  $i$
- (7) print  $p_0, p_1, p_2, p_3, \dots$

Moreover, these EM re-estimates allocate the corpus  $f$  a sequence of corpus probabilities that is monotonic increasing

$$L(f; p_0) \leq L(f; p_1) \leq L(f; p_2) \leq L(f; p_3) \leq \dots$$

**Proof** See Definition 13, and theorems 5 and 10.

To demonstrate in greater detail **how** the EM algorithm is applied to PCFGs, we are presenting now an example. Remind first that we showed in our Example 5 that a treebank grammar is (in principle) able to disambiguate correctly different forms of PP attachment, and remind also that we had to introduce some simple PP mark-up to achieve this result. Now, however, we are choosing a slightly simpler example, so that the number of EM calculations is kept small. We also would like to give you an intuition about **why** we should (in principle) be able to learn good PCFGs by the EM algorithm. Again, we are dealing with a CFG with a PP-attachment ambiguity. This time, however, the training corpus is not made up of full-parse trees of the grammar but of sentences. So, in this scenario, learning a PCFG seems to be difficult. To make it easier for us, only two types of sentences shall occur in our corpus: Whereas the sentences of the first type are ambiguous, the sentences of the second type are not. Our simple goal is to demonstrate that the EM algorithm is able to learn from the unambiguous sentences how to disambiguate the ambiguous sentences. Think about it: Exactly this unimposing feature enables the EM algorithm to learn also highly ambiguous grammars from sentence corpora. For real-world corpora and real-world grammars, it is almost guaranteed that the corpus sentences have on average thousands of analyzes. So, no sentence of the corpus will hardly be completely unambiguous. On the other hand, almost all sentences of the corpus clearly contain partly unambiguous information — represented for example in small sub-trees that the analyzes of this sentence share. Now, by weighting the unambiguous information high, and at the same time, by weighting the ambiguous information low (or even better: indifferently/randomly), and by repeating this process over and over again, the EM algorithm will finally output a PCFG that might have “learned” something: a PCFG that is able to select for each sentence the one single analysis that fits best the information encoded in the corpus. Here is our toy example:

**Example 7** We shall now consider an experiment in which a manually written CFG is estimated by the EM algorithm. We shall assume that the following corpus of 15 sentences is given to us

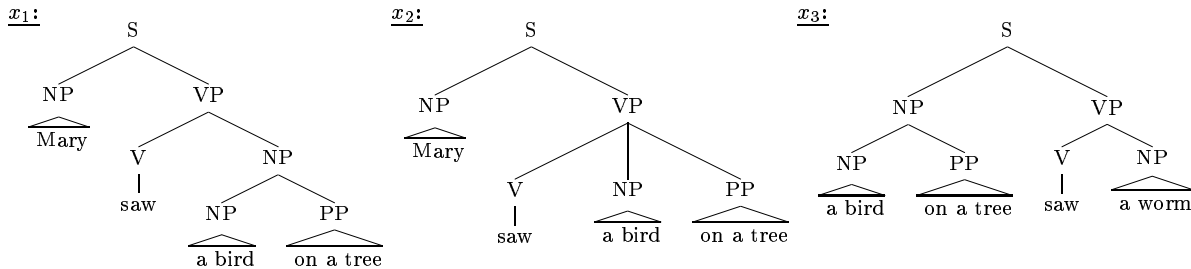
$f(y)$	$y$	
5	$y_1$	$y_1 = \text{“Mary saw a bird on a tree”}$
10	$y_2$	$y_2 = \text{“a bird on a tree saw a worm”}$

Using this text corpus, we shall compute the EM re-estimates of the following PCFG (that is able to parse all the corpus sentences)

$S \rightarrow NP VP$	(1.00)
$VP \rightarrow V NP$	(0.50)
$VP \rightarrow V NP PP$	(0.50)
$NP \rightarrow NP PP$	(0.25)
$NP \rightarrow Mary$	(0.25)
$NP \rightarrow a\ bird$	(0.25)
$NP \rightarrow a\ worm$	(0.25)
$PP \rightarrow on\ a\ tree$	(1.00)
$V \rightarrow saw$	(1.00)

**Warning:** As being displayed, the uniform distributions on the grammar fragments ( $G_S, G_{VP}, G_{NP} \dots$ ) serve in this example as a starting instance  $p_0 \in \mathcal{M}_G$ . On the one hand, this is not bad, since this is (or was?) the common practice for EM training of probabilistic grammars. On the other hand, the EM algorithm gives us the freedom to experiment with a huge range of starting instances. Now, from my current state of experience, it is very very very important that this freedom is really used when training real-world grammars. Clearly, the convergence proposition of Theorem 11 permits that some starting instances of the EM algorithm may lead to better re-estimates than other starting instances. So, more simply, if we are experimenting with more than one single starting instance, then we can seriously hope that our efforts are re-compensated by getting a better probabilistic grammar.

First of all, note that the first sentence is ambiguous, whereas the second is not. The following figure displays the full-parse trees of both.



**First EM iteration.** In the **E-step** of the EM algorithm for PCFGs, a treebank  $f_{\mathcal{T}_q}$  has to be generated on the basis of the starting instance  $q := p_0$ . The generation process is very easy. First, we have to calculate the probabilities of the full-parse trees  $x_1, x_2$  and  $x_3$ , which in turn provide us with the probabilities of the sentences  $y_1$  and  $y_2$ . Here are the results

$p_0(x)$	$x$	$p_0(y)$	$y$
0.0078125	$x_1$	0.0390625	$y_1$
0.0312500	$x_2$	0.0078125	$y_2$
0.0078125	$x_3$		

For example,  $p_0(x_1)$  and  $p_0(y_1)$  are calculated as follows (using definitions 16 and 12.\*):

$$p_0(x_1) = \prod_{r \in G} p(r)^{f_r(x_1)}$$

$$\begin{aligned}
&= p(S \rightarrow \text{NP VP}) \cdot p\left(\begin{array}{c} \text{NP} \\ \hline \text{Mary} \end{array}\right) \cdot p(\text{VP} \rightarrow \text{V NP}) \cdot p(\text{V} \rightarrow \text{saw}) \cdot p\left(\begin{array}{c} \text{NP} \\ \hline \text{a bird} \end{array}\right) \cdot p\left(\begin{array}{c} \text{PP} \\ \hline \text{on a tree} \end{array}\right) \\
&= 1.00 \cdot 0.25 \cdot 0.50 \cdot 1.00 \cdot 0.25 \cdot 0.25 \cdot 1.00 \\
&= 0.0078125
\end{aligned}$$

$$\begin{aligned}
p_0(y_1) &= \sum_{\text{yield}(x)=y_1} p(x) \\
&= p(x_1) + p(x_2) \\
&= 0.0078125 + 0.0312500 \\
&= 0.0390625
\end{aligned}$$

Now, using the probability distribution  $q := p_0$ , we have to generate the treebank  $f_{\mathcal{T}_q}$  by distributing the frequencies of the sentences to the full-parse trees. The result is

$f_{\mathcal{T}_q}(x)$	$x$
1	$x_1$
4	$x_2$
10	$x_3$

For example,  $f_{\mathcal{T}_q}(x_1)$  is calculated as follows (see line (3) of the EM procedure in Theorem 11):

$$\begin{aligned}
f_{\mathcal{T}_q}(x_1) &= f(\text{yield}(x_1)) \cdot q(x_1 | \text{yield}(x_1)) \\
&= f(y_1) \cdot q(x_1 | y_1) \\
&= f(y_1) \cdot \frac{q(x_1)}{q(y_1)} \\
&= 5 \cdot \frac{0.0078125}{0.0390625} \\
&= 1
\end{aligned}$$

In the **E-step** of the EM algorithm for PCFGs, we have to read off the treebank grammar  $\langle G, p_{\mathcal{T}_q} \rangle$  from the treebank  $f_{\mathcal{T}_q}$ . Here is the result

S $\rightarrow$ NP VP	(1.000)
VP $\rightarrow$ V NP	(0.733 $\approx$ $\frac{1+10}{15}$ )
VP $\rightarrow$ V NP PP	(0.267 $\approx$ $\frac{4}{15}$ )
NP $\rightarrow$ NP PP	(0.268 $\approx$ $\frac{1+10}{41}$ )
NP $\rightarrow$ Mary	(0.122 $\approx$ $\frac{1+4}{41}$ )
NP $\rightarrow$ a bird	(0.366 $\approx$ $\frac{1+4+10}{41}$ )
NP $\rightarrow$ a worm	(0.244 $\approx$ $\frac{10}{41}$ )
PP $\rightarrow$ on a tree	(1.000)
V $\rightarrow$ saw	(1.000)

The probabilities of the rules of this grammar form our first EM re-estimate  $p_1$ . So, we are ready for the second EM iteration. The following table displays the rules of our manually written CFG, as well as their probabilities allocated by the different EM re-estimates

CFG rule	$p_0$	$p_1$	$p_2$	$p_3$	...	$p_{18}$
S $\rightarrow$ NP VP	1.000	1.000	1.000	1.000		1.000
VP $\rightarrow$ V NP	0.500	0.733	0.807	0.850		0.967
VP $\rightarrow$ V NP PP	0.500	0.267	0.193	0.150		0.033
NP $\rightarrow$ NP PP	0.250	0.268	0.287	0.298		0.326
NP $\rightarrow$ Mary	0.250	0.122	0.118	0.117		0.112
NP $\rightarrow$ a bird	0.250	0.366	0.357	0.351		0.337
NP $\rightarrow$ a worm	0.250	0.244	0.238	0.234		0.225
PP $\rightarrow$ on a tree	1.000	1.000	1.000	1.000		1.000
V $\rightarrow$ saw	1.000	1.000	1.000	1.000		1.000

In the 19<sup>th</sup> iteration, nothing new happens. So, we can quit the algorithm and discuss the results. After all, of course, the most interesting thing for us is how these different PCFGs perform, if they are applied to ambiguity resolution. So, we will have a look at the probabilities these PCFGs allocate to the two analyzes of the first sentence. We have already noted several times that  $p$  prefers  $x_1$  to  $x_2$ , if

$$p(\text{VP} \rightarrow \text{V NP}) \cdot p(\text{NP} \rightarrow \text{NP PP}) > p(\text{VP} \rightarrow \text{V NP PP})$$

The following table displays the values of these terms for our re-estimated PCFGs

$p$	$p(\text{VP} \rightarrow \text{V NP}) \cdot p(\text{NP} \rightarrow \text{NP PP})$	$p(\text{VP} \rightarrow \text{V NP PP})$
$p_0$	$0.500 \cdot 0.250 = 0.125$	0.500
$p_1$	$0.733 \cdot 0.268 = 0.196$	0.267
$p_2$	$0.807 \cdot 0.287 = \mathbf{0.232}$	0.193
$p_3$	$0.850 \cdot 0.298 = \mathbf{0.253}$	0.150
$\vdots$		
$p_{18}$	$0.967 \cdot 0.326 = \mathbf{0.315}$	0.033

So, the EM re-estimates prefer  $x_1$  to  $x_2$  starting with the second EM iteration, due to the fact that the term  $p(\text{VP} \rightarrow \text{V NP}) \cdot p(\text{NP} \rightarrow \text{NP PP})$  is monotonic increasing (within a range from 0.125 to 0.315), while at the same time the term  $p(\text{VP} \rightarrow \text{V NP PP})$  is drastically monotonic decreasing (from 0.500 to 0.033).

## Acknowledgments

Parts of the paper cover parts of the teaching material of two courses at ESSLLI 2003 in Vienna. One of them has been sponsored by the *European Chapter of the Association for Computational Linguistics (EACL)*, and both have been co-lectured by Khalil Sima'an and me. Various improvements of the paper have been suggested by Gabriel Infante-Lopez, Karin Müller, Mark-Jan Nederhof, Breannán Ó Nualláin, and Khalil Sima'an.

## References

- Backus, J. W. (1959). The syntax and semantics of the proposed international algebraic language of the Zürich ACM-GAMM Conference. In *Proceedings of the International Conference on Information Processing*, Paris.

- Baker, J. K. (1979). Trainable grammars for speech recognition. In D. Klatt and J. Wolf (Eds.), *Speech Communication Papers for ASA '97*, pp. 547–550.
- Baum, L. E. (1972). An inequality and associated maximization technique in statistical estimation for probabilistic functions of Markov processes. *Inequalities III*, 1–8.
- Beil, F., G. Carroll, D. Prescher, S. Riezler, and M. Rooth (1999). Inside-outside estimation of a lexicalized PCFG for German. In *Proceedings of ACL'99*, College Park, MD.
- Booth, T. L. and R. A. Thompson (1973). Applying probability measures to abstract languages. *IEEE Transactions on Computers C-22*(5), 442–450.
- Charniak, E. (1996). Tree-bank grammars. Technical Report CS-96-02, Brown University.
- Chi, Z. (1999). Statistical properties of probabilistic context-free grammars. *Computational Linguistics* 25(1).
- Chi, Z. and S. Geman (1998). Squibs and discussions: Estimation of probabilistic context-free grammars. *Computational Linguistics* 24(2).
- Chomsky, N. (1956). Three models for the description of language. *IRE Transactions on Information Theory*.
- Cover, T. M. and J. A. Thomas (1991). *Elements of Information Theory*. New York: Wiley.
- DeGroot, M. H. (1989). *Probability and statistics* (2 ed.). Addison-Wesley.
- Dempster, A. P., N. M. Laird, and D. B. Rubin (1977). Maximum likelihood from incomplete data via the *EM* algorithm. *J. Royal Statist. Soc.* 39(B), 1–38.
- Duda, R. O., P. E. Hart, and D. G. Stork (2001). *Pattern Classification — 2nd ed.* New York: Wiley.
- Hopcroft, J. E. and J. D. Ullman (1979). *Introduction to Automata Theory, Languages, and Computation*. Reading, MA: Addison-Wesley.
- Jaynes, E. T. (1957). Information theory and statistical mechanics. *Physical Review* 106, 620–630.
- Johnson, M. (1998). PCFG models of linguistic tree representations. *Computational Linguistics* 24(4).
- Klein, D. and C. D. Manning (2003). Accurate unlexicalized parsing. In *Proceedings of ACL-03*, Sapporo, Japan.
- Lari, K. and S. J. Young (1990). The estimation of stochastic context-free grammars using the inside-outside algorithm. *Computer Speech and Language* 4, 35–56.
- McLachlan, G. J. and T. Krishnan (1997). *The EM Algorithm and Extensions*. New York: Wiley.
- Nederhof, M.-J. and G. Satta (2003). Probabilistic parsing as intersection. In *Proceedings of the 8th International Workshop on Parsing Technologies (IWPT-03)*, Nancy, France.
- Pereira, F. and Y. Schabes (1992). Inside-outside reestimation from partially bracketed corpora. In *Proceedings of ACL'92*, Newark, Delaware.
- Prescher, D. (2001). Inside-outside estimation meets dynamic EM. In *Proceedings of IWPT-2001*, Beijing.

- Prescher, D. (2002). *EM-basierte maschinelle Lernverfahren für natürliche Sprachen*. Ph. D. thesis, IMS, University of Stuttgart.
- Ratnaparkhi, A. (1997). A simple introduction to maximum-entropy models for natural language processing. Technical report, University of Pennsylvania.
- Sanchez, J. A. and J. M. Benedi (1997). Consistency of stochastic context-free grammars from probabilistic estimation based on growth transformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 19.
- Wu, C. F. J. (1983). On the convergence properties of the EM algorithm. *The Annals of Statistics* 11(1), 95–103.