

# A Probabilistic Line Breaking Algorithm

Remco R. Bouckaert

Computer Science Department, University of Waikato, New Zealand  
rrb@xm.co.nz, remco@cs.waikato.ac.nz

September 17, 2003

## Abstract

We show how a probabilistic interpretation of an ill defined problem, the problem of finding line breaks in a paragraph, can lead to an efficient new algorithm that performs well. The graphical model that results from the probabilistic interpretation has the advantage that it is easy to tune due to the probabilistic approach. Furthermore, the algorithm optimizes the probability a break up is acceptable over the whole paragraph, it does not show threshold effects and it allows for easy incorporation of subtle typographical rules. Thanks to the architecture of the Bayesian network, the algorithm is linear in the number of characters in a paragraph. Empirical evidence suggests that this algorithm performs closer to results published through desk top publishing than a number of existing systems.

## 1 Introduction

With the advance of information systems, the amount of automatically formatted text increases daily. However, judging from a software patent search, (searched [www.spi.com](http://www.spi.com) for "text formatting" 1995-2003) not a lot of attention is spent on increasing the quality of layout. New interest in this area is driven by the standardization of style specifications through CSS and XSL-FO [8]. Limitations of automatic typesetting systems are partly due to the fact that the underlying algorithms have been designed with resource restrictions in mind. For example, once  $\text{\TeX}$  [3] has sufficient material to complete a page, the page is output to free up memory. If many floating objects like figures and tables are present, these can be placed in undesirable locations far from the place where they are referenced.

Likewise,  $\text{\TeX}$ 's line breaking algorithm [4] is optimized on speed in a three phased approach: in the first phase, an attempt is made to break the line without hyphenation. If the resulting line break-up is not acceptable the second phase is started where hyphenation is allowed. If the result is still unacceptable, a third phase is started, where the definition of acceptable break-up is made more tolerable. So, in fact three criteria are used, each more expensive to calculate.

Considerations of limited resources were quite valid in the time that these algorithms were developed. However, they do not apply today and this opens a way to improve various typesetting algorithms. In typesetting a wide range of decisions need to be made in the areas of hyphenation, line breaking, page breaking, floating object placement, footnotes, headers and footers, page numbering, reference resolution, etc. In this paper, we consider line breaking only, but the principles presented should be extendible to page breaking, floating object placement and other areas as well.

A reasonable approach seems to be to ask professional publishers what makes a good line break-up. However, this knowledge is very difficult to elicit. Typically only examples of what is acceptable and what is not can be given, with disagreement on borderline cases. So finding good line break-ups is an ill specified problem in which decisions need to be made with uncertainty of the function to optimize.

A popular line breaking algorithm is the one used in  $\text{\TeX}$  [4]. This algorithm is based on an empirical formula for the acceptability of a line break-up that includes demerits, badness, tolerances and penalties. It is a complicated task to finetune the algorithm because of threshold effects of  $\text{\TeX}$ 's three phases and the number of parameters involved.

In this paper we take a probabilistic approach to line breaking in which we model the probability that a particular break-up is acceptable. The probability of acceptability can be maximized and the break-up with the highest probability chosen. A dynamic graphical model is used to capture acceptability.

The next section introduces concepts of line breaking and related work. In Section 3 a probabilistic algorithm is presented for breaking lines. Section 4 compares the performance of some popular typesetting systems with data published through desk top publishing together with an empirical comparison of the new algorithm. We finish with concluding remarks and suggestions for further extensions in Section 5.

## 2 The line breaking problem

Breaking a paragraph into lines consists of selecting the break points in a paragraph. To determine how much space a line takes, each character (or rather glyph, see comment on ligatures below) in the paragraph is modeled as a box with a specific width, height and elevation from the baseline. These properties are determined by the font being used and the font size. The length of a line is the sum of the widths of the boxes in the line. The width of a box is not only determined by the width of the character, but is also influenced by the following character. If the width of a character would not be *kerned*, the result may look rather unpleasing. For example, compare these unknerved and knerved words:

BAYES  
BAYES

The latter has less space between the A and Y, and the width of the A is modeled less wide then when an A is followed by, say, another A.

Some fonts treat certain combinations of characters as a single character. Examples of such so called *ligatures* is the combination of f and i and the combination of two f's and an i. The following line illustrates the difference between separate characters and their ligatures:

efficient fix  
efficient fix

So, the first step in breaking a line is determining the width of its characters, taking kerning and ligatures in account. The next step is determining possible break points. Typical break points in Latin alphabet languages are whitespaces, hyphen characters and points where hyphenation is allowed.

Finding these hyphenation points by maintaining a wordlist is cumbersome because of the many variations of appearances of a word (e.g. singular and plural). Furthermore, maintaining such a word list in areas where new words are invented regularly (e.g. brand names) is a huge task. An elegant alternative is to use a hyphenation algorithm [3, 6] which essentially works as follows. Some words are stored in an exception list (for example, as-so-ciate). For words not in the exception list, a pattern list is used, which is generated from a word list containing hyphenated words. Patterns contain characters and numbers, for example 'y3thin'. A word is compared with the characters in the patterns and every character is assigned the highest number in the pattern. The word 'anything' matches 'ythin' in the pattern 'y3thin' and the y gets assigned the number 3. Any odd numbered character is a hyphenation point. So 'any-thing' is allowed. It is claimed [3] that more than 95% of the allowable hyphenation points in US English can be found by this algorithm (taking word frequencies in account).

Hyphenation points at the first  $\lambda$  and last  $\rho$  characters are discarded to prevent unacceptable hyphenation points. For example, consider the hyphenation points found when not suppressing the first and last points in the words 'a-ha', 'Bal-i', 's-tu-den-t' or 'A-pu-li-a'. By default,  $\lambda$  is 2 and  $\rho$  is 3 for English.

This example makes the reason obvious. In  $\text{\TeX}$ , the parameters  $\lambda$  and  $\rho$  are kept in `\lefthyphen` and `\righthyphen` respectively. The command `\showhyphens` can be used to show the hyphenation points.

```
D:\>tex
This is TeX, Version 3.14159 (MiKTeX 2)

*\lefthyphenmin=0

*\righthyphenmin=0

*\showhyphens{student}
```

Underfull \hbox (badness 10000) detected at line 0  
[] \tenrm s-tu-den-t

A broken line may not completely fit into the available space between the left and right margin. To make it fit, space may be distributed among the whitespaces in the line, or some whitespace may be taken out if the text exceeds the text width. We assume that there is a single target width for the complete paragraph and the paragraph is adjusted. An indent that applies to the first line of the paragraph can be modeled using a fixed width unbreakable space (possibly with a negative value).

Now that we have a flavor of the problem, we can specify some terms more formally. A *paragraph*  $p$  is a sequence of  $n > 0$  characters  $c_i$ ,  $1 \leq i \leq n$ . A *break point candidate* in  $p$  is an index of a character in  $p$  for which it is allowed to break a line. Typical break point candidates are space and hyphen characters and hyphenation points in words. We are interested in finding a sequence of character indexes  $s_0, s_1, \dots, s_k$ ,  $\forall 1 \leq j \leq k$   $1 \leq s_j \leq n$  and  $s_{j-1} < s_j$  where each  $s_j$  is a break point candidate. Each of the indices  $s_j$  indicates a *break point* in  $p$ . Index  $s_0$  is added for convenience of further definitions and by convention  $s_0 = 0$  and  $s_k = n$ . A *line* in  $p$  with break points  $s_0, \dots, s_k$  is the sub-sequence of characters between breakpoints, more specific for  $1 \leq j \leq k$ ,  $l_j = \{c_{s_{j-1}+1} \dots c_{s_j}\}$  (note  $s_0 = 0$  so  $l_1 = \{c_1 \dots c_{s_1}\}$ ). So, a break point  $s_j$  is the last character on line  $l_j$ .

The *character width*  $cw_i$   $1 \leq i \leq n$  of a character  $c_i$  in a paragraph  $p$  with break points  $s_0, \dots, s_k$  is the width of the character taking the font properties in account. Furthermore, the character width takes breakpoints in account with the following rules:

- A character not on a break point may have its width corrected due to kerning.
- A space character at the beginning or end of a line has width 0.
- A non-space and non-hyphen character not at the end of a word that is a break point has width of the character plus the width of a hyphen.
- If  $c_i$  is part of a ligature, and the ligature is not broken,  $cw_i$  is the width of the ligature divided among the characters that make up the ligature.

The *natural line width*  $nlw_i$ ,  $0 \leq i \leq n$  at character  $c_i$  is the sum of widths of characters on to the closest previous breakpoint  $s_j$ . More specific,  $nlw_0$  is the indent, and if  $s_j = 0$ ,  $nlw_i = nlw_0 + \sum_{d=1}^i cw_d$ , otherwise  $nlw_i = \sum_{d=s_j+1}^i cw_d$ . The *natural line width* of a line  $l_j$ ,  $1 \leq j \leq k$  is  $nlw_{s_j}$ . The *number of spaces* in a line  $l_j$  is  $nsp_{s_j} = \sum_{i=s_{j-1}+1}^{s_j-1} \{1|c_i = \text{space}\}$ , so spaces at the start and end of a line are not counted (it is assumed no space follows another space).

The *line breaking problem* for a paragraph  $p$  for a desired *text width* and *indent* is finding a set of break points of  $p$  that look 'nice'. We will consider fixed text width,  $tw$ , throughout this paper with the exception of the first line, which can have a (possibly negative) indent  $in$ .

A simple approach to the line breaking problem is to consider lines one by one, like Unix’s Troff. Simply keep adding words as long as the natural line width does not exceed the text width. If a word makes the line overflow, try hyphenate the word (using a word list or Liang’s hyphenation algorithm) and add as many parts (possibly zero) of the word to the line as long as the text width is not exceeded. Output the line and place the remainder on the following line. In this approach, line breaks are only optimized locally, without regard to the breaks in the remainder of the paragraph.

TeX’s line breaking algorithm [4] on the other hand optimizes line breaks on the level of the paragraph. We give a short description of TeX’s line breaking algorithm here (see [3] for an extensive description). TeX determines character widths taking kerning and ligatures in account and uses a three phased process.

In phase 1, no hyphenation points are considered for line breaking, only white spaces. For a paragraph broken in  $k$  lines, for each line  $j = 1 \dots k$  a badness  $b_j$  is calculated, using  $b_j = 100 \left( \frac{nlw_{s_j} - tw}{nsp_{s_j} \cdot f} \right)^3$  where  $nlw_{s_j}$  the natural line width,  $tw$  the text width,  $nsp_{s_j}$  the number of spaces of line  $j$  and  $f$  is a factor that differs depending on whether the line needs to be stretched or shrunk. Note that  $nlw_{s_j} - tw$  represents the amount that the line needs to be stretched or shrunk. Depending on the value of  $b_j$ , the line is classified as tight, decent, loose or very loose. If none of the line’s badnesses exceed a ‘pretolerance’ limit, the paragraph is accepted.

Otherwise, in phase 2, the hyphenation candidates are calculated for the paragraph and a new set of breakpoints is calculated. If all of the line’s badnesses are below a tolerance level (which differs from the pretolerance level) a demerit for the paragraph is calculated as a combination of line badnesses, roughly as

$$\sum_{j=1}^n (c + b_j)^2 + p \cdot |p| \tag{1}$$

where  $c$  is a constant line penalty and  $p$  a penalty term. A tight line next to a loose line increases the penalty  $p$ . If two consecutive lines are hyphenated, or if the last broken line is hyphenated, the penalty is increased. A paragraph with a demerit below a threshold is accepted.

Otherwise, in phase 3, the steps in phase 2 are repeated but now with an emergency stretch that allows lines to shrink or expand more. If this last phase does not succeed, TeX outputs overly long lines, due to the thresholds in the algorithm.

### 3 A probabilistic line breaking algorithm

We start with deriving the line breaking algorithm and illustrate how to design a simple but efficient probabilistic algorithm based on a dynamic Bayesian network. Then, this model is generalized and we show how the behavior of the network can be influenced by tuning network parameters. Because the network is easy to comprehend, the resulting change in behavior is intuitively satisfying.

### 3.1 A simple probabilistic line breaking algorithm

We approach the line breaking problem by defining a probability distribution that represents the acceptability of a set of breakpoints.

Let  $p$  be a paragraph with  $n$  characters  $c_1, \dots, c_n$ , text width  $tw$ , indent  $in$  and a set of break points  $S = \{s_0, \dots, s_k\}$ . Let  $a_i$ ,  $1 \leq i \leq n$  be a boolean representing whether the break-up up to character  $i$  is acceptable, and let  $a$  denote  $a_n$ . Then we are interested in  $P(a, S, p, in)$  where we leave  $tw$  and  $in$  implicit. We make the following two assumptions

1. the acceptability of the first  $i$  characters in the break-up is independent of the break-up of the remainder, provided  $nlw_{i-1}$  is known for the remainder.
2. we have no initial preference for any breakup, that is,  $P(S|p, in)$  is constant for any  $S$ .

The first assumption is the base for the efficiency of  $\text{T}_{\text{E}}\text{X}$ 's line breaking algorithm. The last assumption is for the sake of simplifying the derivation of the probabilistic algorithm, and will be relaxed in the following section. Now,  $P(a, S, p, in) = P(a|S, p, in)P(S|p, in)P(p, in)$ . Since  $p$  and  $in$  are given, and using the second assumption,  $P(p, in)$  and  $P(S|p, in)$  are constant, so  $P(a, S, p, in) \propto P(a|S, p, in)$ . The first assumption lets us decompose  $P(a|S, p, in)$  as  $P(a_i|S \setminus \{i+1, \dots, n\}, p, in) \cdot P(a|S \setminus \{1, \dots, i\}, p, nlw_i, in)$ , where the first term is the acceptability of the break-up of the first  $i$  characters, and the second term the acceptability of the remainder given the natural line width  $nlw_i$  of the last line of the first  $i$  characters. Now note  $P(a_i|S \setminus \{i+1, \dots, n\}, p, in) = P(a_i|S \setminus \{i+1, \dots, n\}, p \setminus \{c_{i+1}, \dots, c_n\}, in)$ . Therefore, we can write  $P(a, S|p, in)$  as the product of the acceptability of the individual characters, that is

$$\begin{aligned}
 P(a, S|p, in) &= P(a_1, S \setminus \{2, \dots, n\}|p, in) \\
 &\quad P(a, S \setminus \{2, \dots, n\}|p \setminus \{c_2, \dots, c_n\}, nlw_1, in) \\
 &= P(a_1, 1 \in S|p, in) \\
 &\quad P(a, S \setminus \{c_2, \dots, c_n\}|p \setminus \{c_2, \dots, c_n\}, nlw_1) \\
 &= P(a_1, 1 \in S|p, in) \cdot P(a_2, S \setminus \{3, \dots, n\}|p, nlw_1) \\
 &\quad P(a, S \setminus \{3, \dots, n\}|p \setminus \{c_3, \dots, c_n\}, nlw_2) \\
 &= \prod_{i=1}^n P(a_i, i \in S|p, nlw_{i-1})
 \end{aligned}$$

Further, the natural line width at character  $i$ ,  $0 \leq i \leq n$  is calculated as

$$nlw_i = \begin{cases} cw_i & |i-1 \in S \\ nlw_{i-1} + cw_i & |i-1 \notin S \\ in & |i = 0 \end{cases}$$

Now we are left with determining  $P(a_i|i \in S, p, nlw_{i-1})$ . If  $c_i$  is not a break point candidate (i.e., it is not a space, hyphen, or hyphenation point) and  $i \in S$ , the break-up is not acceptable under any circumstances, and  $P(a_i|i \in S, p, nlw_{i-1}) = 0$ . If not breaking at  $c_i$  means that the natural line width remains within the text width, the line break-up is just as acceptable as the acceptability of the remainder of the paragraph,  $P(a_i|i \in S, p, nlw_{i-1}) = 1$ . However, not breaking may be unacceptable as well, if this means extending the natural length beyond the text width. Breaking before the natural line width has reached the text width also decreases acceptability. In summary,  $P(a_i|i \in S, p, nlw_{i-1})$  is defined as

$$\begin{array}{ll} 0 & |i \text{ is not a break point candidate and } i \in S \\ 1 & |nlw_{i-1} + cw_i \leq tw \text{ and } i \notin S \\ P^a(cw_i|nlw_{i-1} - tw) & |nlw_{i-1} + cw_i > tw \\ P^a(cw_i|nlw_{i-1} - tw) & |nlw_{i-1} + cw_i \leq tw \text{ and } i \in S \end{array}$$

where  $P^a$  a distribution representing acceptability for breaking not earlier than at  $c_i$  instead of  $c_{i-1}$ . For the choice of this distribution, we let ourselves inspire by the empirical distribution of what seems to be acceptable as shown in Figure 3 by the graph for published line lengths. This shows a distribution which is roughly normal, with different variances for lines exceeding and undercutting the text width. So, we define

$$P^a(cw_i|x) = \begin{cases} N(x + cw_i, \sigma^+)/P^a(x) & |x + cw_i \geq 0 \\ N(x + cw_i, \sigma^-)/P^a(x) & |x + cw_i < 0 \end{cases}$$

where  $N(x, \sigma)$  the standard normal distribution  $N(x, \sigma) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{x^2}{2\sigma^2}}$ ,  $P^a(x) = N(x, \sigma^+)$  if  $x \geq 0$  and  $N(x, \sigma^-)$  otherwise and  $\sigma^+$  and  $\sigma^-$  two variances, which are the two parameters to choose. Note that the contribution of line  $j$  to  $P(a|S, p, in)$  is  $\prod_{i=s_{j-1}+1}^{s_j} P(i \in S|p, nlw_{i-1})$  and that this reduces to  $N(nlw_{s_j} - tw, \sigma^+)$  if line  $j$ 's natural width exceeds the text width ( $nlw_{s_j} > tw$ ) and  $N(nlw_{s_j} - tw, \sigma^-)$  if  $nlw_{s_j} \leq tw$ .

What we are after is of course selecting a set of break points that maximizes  $P(a|S, p, in)$ . Note that what we have done so far is derive a dynamic Bayes network with a structure as pictured in Figure 1. The character widths  $cw_i$  are given (refer to the rules as described in Section 2 for the actual details). Node  $b_i$  indicates whether breaking is allowed at character  $c_i$ . The conditional distribution for  $nlw_i$  and  $a_i$  are described in this section above.<sup>1</sup> What remains is selecting a value for the boolean variable indicated by  $i \in S$  in the figure. Straightforward inference e.g. as described by Pearl [7] or Lauritzen [5], for finding the most likely configuration for  $b_i$  suffices to find the breakpoints.

The following lemma provides some insight in how this algorithm relates to TeX's line breaking algorithm.

<sup>1</sup>Instead of implementing one slice per character, one slice per break point candidate can be used. Of course, a correction for the characters widths between break point candidates need to be implemented then.

LEMMA 3.1 For a paragraph  $p$ , a set of break points  $S = \{s_0, \dots, s_k\}$  maximizes  $P(a, S, p, in)$  iff  $S$  maximizes

$$\sum_{j \in S^+} -(nlw_{s_j} - tw)^2 + \sum_{j \in S^-} (-(nlw_{s_j} - tw)^2 C_1 + C_2) - k.C_3 \quad (2)$$

where  $S^+$  the breakpoints  $j \in S$  such that  $nlw_{s_j} > tw$ ,  $S^- = S \setminus S^+$ ,  $C_1 = \frac{\sigma^{+2}}{\sigma^{-2}}$ ,  $C_2 = \sigma^{+2} \log \frac{\sigma^+}{\sigma^-}$ , and  $C_3 = \sigma^{+2} \log(2\pi\sigma^+)$ .

**Proof:** Let  $p$  be a paragraph with  $n$  characters  $c_1, \dots, c_n$  for which the natural line width  $nlw_j$  can be calculated for a given set of breakpoints  $S$ , then

$$\arg \max_S P(a, S, p, in) = \arg \max_S \prod_i P(a_i | i \in S, nlw_{i-1}) = \arg \max_S \prod_{j=1}^k \prod_{i=s_{j-1}+1}^{s_j} P(a_i | i \in S, nlw_{i-1})$$

Now, observe that the term  $\prod_{i=s_{j-1}+1}^{s_j} P(a_i | i \in S, nlw_{i-1})$  equals  $N(nlw_{s_j} - tw, \sigma^+)$  if  $nlw_{s_j} > tw$  and that it equals  $N(nlw_{s_j} - tw, \sigma^-)$  if  $nlw_{s_j} < tw$ . So, we can write the product above as

$$\arg \max_S \prod_{j \in S^+} N(nlw_{s_j} - tw, \sigma^+) \prod_{j \in S^-} N(nlw_{s_j} - tw, \sigma^-)$$

Maximizing this term equals maximizing the log of this term (log is an ascending function and all terms  $N(.,.)$  are strictly positive). Using  $\log N(m, \sigma) = \log \frac{1}{2\pi\sigma} e^{-m^2/\sigma^2} = -\log(2\pi\sigma) - m^2/\sigma^2$  gives,

$$\begin{aligned} \arg \max_S \sum_{j \in S^+} -\log(2\pi\sigma^+) - \frac{(nlw_{s_j} - tw)^2}{\sigma^{+2}} \\ + \sum_{j \in S^-} -\log(2\pi\sigma^-) - \frac{(nlw_{s_j} - tw)^2}{\sigma^{-2}} \end{aligned}$$


We can bring the constant value  $k$  times  $\log(2\pi\sigma^+)$  outside the sum and multiply with the positive constant  $\sigma^{+2}$  and still get the same  $S$ , giving

$$\begin{aligned} \arg \max_S \sum_{j \in S^+} -(nlw_{s_j} - tw)^2 + \\ \sum_{j \in S^-} \left( \sigma^{+2} \log \frac{\sigma^+}{\sigma^-} - (nlw_{s_j} - tw)^2 \frac{\sigma^{+2}}{\sigma^{-2}} \right) - k\sigma^{+2} \log(2\pi\sigma^+) \end{aligned}$$

Substituting  $C_1, C_2$ , and  $C_3$  as defined above gives the stated result.  $\square$

In words, maximizing the probability of acceptability  $P(a, S, p, in)$  by choosing a set of break points  $S$  results in the same set of breakpoints as maximizing the sum over the lines longer than the text width  $tw$  (first sum term) and lines shorter than  $tw$  (second sum term). Comparing this with the formula that  $\text{\TeX}$  maximizes, equation (1), we see that  $\text{\TeX}$  uses the square of a sum of

Figure 1: Network representing simple probabilistic line breaking algorithm.



cubic functions, where we derived a quadratic one. Both have a term penalizing the number of lines (the term  $-kC_3$  in (2), taking the reasonable assumption that  $\log(2\pi\sigma^+) > 0$ ). Further,  $\text{\TeX}$  is concerned with the amount that spaces need to be expanded or squeezed, instead of looking at the whole line width. In Section 4 we will see that this does not matter too much.

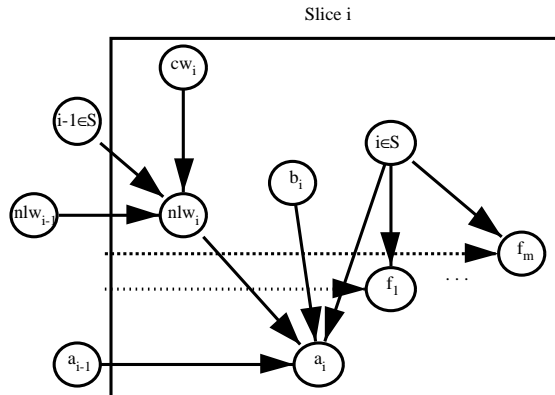
Benefits of this simple probabilistic approach are that it allows for tuning with intuitively attractive parameters. Further, by choosing a smooth probability distribution for  $P^a$  with infinite domain, no threshold effects leading to overly long lines occur (as happens in  $\text{\TeX}$ 's approach, resulting in overfull hboxes). Finally, note that when instantiating the network with values of  $S$ ,  $cw$  etc., the functional dependencies between  $nlw_i$  and  $nlw_{i-1}$  have no impact on the inference of  $a$ . So, we are left with a network that has a poly-tree structure for which finding the most likely configuration of  $b_i$  is linear in the size of the network [7].

### 3.2 Extending the probabilistic line breaking algorithm

The second assumption in the previous section can be relaxed to take interline effects in account. This allows discouraging multiple hyphenated lines in a row and preventing tight lines following loose lines. This information can be incorporated by defining new metrics and updating  $P(S|p, in)$ , the preference of a break-up.

There are two issues that need to be considered when factoring in more metrics. First, to ensure it is possible to find the break points  $S$  that are most likely to be accepted efficiently, the network structure that would model the newly created metrics should not create large cycles. It would be sufficient to ensure that the metrics are calculable for each character  $c_i$  from the information in slice  $i$  and  $i - 1$  in the network. So when we have  $m$  metrics  $f_1, \dots, f_m$ , we can write  $P(S, f_1, \dots, f_m|p) = \prod_{i=1}^n P(i \in S, f_{i,1}, \dots, f_{i,m}|p)$ , where  $f_{i,j}$

Figure 2: Network representing full probabilistic line breaking algorithm.




is the metric  $f_j$  calculated at character  $i$ . Second, a choice need to be made modeling the interoperations of the metrics on the probability of acceptability. A pragmatic choice is to assume that the metrics have independent impact on out preference for a break-up  $S$  at a particular character  $c_i$ . Then we can write  $P(i \in S, f_{i,1}, \dots, f_{i,m} | p) = P(i \in S | p) \prod_{j=1}^m P(f_{i,j} | i \in S, p) \propto \prod_{j=1}^m P(f_{i,j} | i \in S, p)$ .

This gives the network as shown in Figure 2. Some metrics that could be calculated for each character are: tightness of previous line (one of 'tight', 'normal', 'loose'), number of previously hyphenated lines (one of 0, 1, 2, many), the amount spaces in a line need to be expanded to make the line fit the text width, and desirability of breakpoint (distinguishing desirability of spaces, hyphens, and hyphenation points). The latter requires the number of spaces at character  $i$  to be calculated as well.

A similar analysis as in the proof of Lemma 3.1 reveals that each metric contributes as a separate sum, i.e.,  $\arg \max_S P(a, S, p, in) = \arg \max_S \sum_{j=1}^k m_{1,s_j} + \sum_{j=1}^k m_{2,s_j} + \dots + \sum_{j=1}^k m_{f,s_j}$ . There is a similarity with (1) in that subtle effects impact additive. The difference is that  $\text{\TeX}$ 's penalties are difficult to interpret, especially because effects are multiplied (the term  $p \cdot |p|$  in (1)).

This extension has all the benefits of the simple probabilistic approach: intuitively tunable parameters, complexity linear in the number of character, and no threshold effects. Furthermore, it captures subtle interline effects and can be extended easily to capture even more factors.

Figure 3: Line lengths for published paragraphs and  $\text{\TeX}$  and DLPager output (204pt is optimal). Horizontal axis shows the line length in points, vertical axis shows the number of lines having the particular line length.




## 4 Empirical results

To get an impression of the behavior of various line breaking algorithms, we looked at a set of existing publications and compared these with the output of some popular typesetting systems. Data was obtained from various scientific journals in the pharmaceutical area. Because the information was provided in PDF format, quite a bit of processing was required to get it into a useable format. Using the XPDF library<sup>2</sup>, text, font and location information was extracted from the PDF files and strings glued together to lines, taking super and subscripts in account. Using simple pattern matching rules, any non-paragraph text, like headers and footers, titles, references etc., was removed and text lines glued to paragraphs. Because of tables and figures floating through the text, this was not 100% successful and a manual post-processing step was required to ensure broken paragraphs were glued or split. The database contained 5.459 paragraphs with 74.120 lines in total. Natural line widths were calculated taking kerning in account.

We compared the published paragraphs with  $\text{\TeX}$  to establish two benchmarks. The  $\text{\TeX}$  input was obtained by collecting text from each of the paragraphs, mapping special characters (like  $\leq$ ,  $\pm$ ,  $\copyright$  etc) and generating a  $\text{\TeX}$  file with a pagebreak after each paragraph. The  $\text{\TeX}$  input uses the same font and paragraph settings as in the published paragraphs. Then  $\text{\LaTeX}$  was run to generate PDF and the same process as for published PDF was applied. However, this time paragraphs could be reliably extracted because each of them was out-

<sup>2</sup> Available from <http://www.foolabs.com/xpdf>.

Figure 4: Line lengths for published paragraphs and output of various XSL-FO processors. Axis as in Figure 3.



put on its own page.  $\text{\TeX}$  was run both with default settings and with sloppy paragraphs. The same process was repeated using DL Pager, a commercial composition engine from Datalogics<sup>3</sup>.

An informal experiment<sup>4</sup> was performed where publication managers were asked to rank randomly selected paragraphs comparing published text with Pager and our algorithm. Paragraphs were presented side by side on a page to make comparison easier and the paragraphs were randomly ordered on the page. The result of this experiment was that publication managers tended to prefer the published result, with our new algorithm close by and Pager last. However, all three types of paragraphs were ranked first and last some of the time. Overall, this result is encouraging because it shows that Bayesian networks can actually help in increasing the typesetting quality.

An instant quantitative impression was obtained by measuring the distribution of the natural line widths and space widths of published material and lines generated by the typesetting systems. Though this does not show subtle inter-line behavior (such as appearance of multiple hyphenated lines), it still gives a good initial assessment of a system since looseness of lines is mentioned by publication managers as the primary indication of the quality of a line.

Figure 3 shows the distribution of the natural line widths for the published paragraphs and the  $\text{\TeX}$  output. Only middle lines are shown, not the first line (due to its indent) or last line. The ideal line is 204 points wide, if it is more, spaces have to be shrunk, or if it is less they have to be stretched. Figure 3 show clearly that handcrafted lines tend to be closer to the ideal length than

<sup>3</sup>See <http://www.datalogics.com> for details

<sup>4</sup>See appendix of this article for details of the experiment.

$\text{\TeX}$  lines. In fact default  $\text{\TeX}$  outputs a large number of lines that exceed the line length so much that there is not enough space in the line to shrink them, and the line exceeds the right boundary. Making paragraphs sloppy largely solves this problem, but there are even more looser lines than with the default setting. DL PAGER output tends to have more lines closer to the ideal output compared to default  $\text{\TeX}$ . However, in cases where lines are loose, they tend to be very loose, making the tail on the left in the left side graph rather thick.

Figure 4 shows natural line widths of three XSL-FO processors. FOP version 0.20.4<sup>5</sup> is an open source processor that uses a line breaking algorithm inspired by  $\text{\TeX}$ . XEP version 3.2<sup>6</sup> and XSL Formatter version 2.4<sup>7</sup> are commercial XSL-FO processors using unknown line breaking algorithms. All three XSL-FO processors use hyphenation rules inspired by  $\text{\TeX}$ 's hyphenation files, indicating that all of them use Liang's hyphenation algorithm [6]. The XSL-FO standard [8] leaves the implementation of the line breaking rules to the XSL-FO processor, within bounds of common sense line breaking rules as specified in the Unicode standard [2]. Consequently, there is very little a user can do to influence the layout of paragraphs when using XSL-FO. Figure 4 shows that all three processors are quite far from the published output. FOP and XSL-Formatter both are very reluctant to expand a space, as indicated by the sharp drop in the number of lines longer than the optimal 204 point width.

Figure 5 shows the distribution of space widths for the non-FO systems under investigation. Lines can be made a bit longer or shorter to justify by expanding or squeezing spaces. For example, in a line with 3 spaces and a natural line width of 201 point 1 point may be added to each of the spaces to expand the line to 204 point. The space width for a line is calculated as (natural line width - text width)/number of spaces where the paragraph width is 204 point and lines without spaces are ignored. The distribution for published text is more peaked and has tails that are less fat than any of the systems. Sloppy  $\text{\TeX}$  has a wider distribution than plain  $\text{\TeX}$  and in fact any of the systems. The distribution of DLPAGER comes closest to the published results. Figure 6 shows the same as Figure 5 for the XSL-FO processors. It illustrates once more the reluctance of Fop and Antenna House's tendency to squeeze spaces. XEP shows least hesitation in squeezing spaces of all the systems.

Table 1 shows the number of lines for published results and the typesetting systems. Systems with many loose lines, like sloppy  $\text{\TeX}$  and Fop tend to produce more lines, while systems with many tight lines, like XEP, produce fewer lines. Plain  $\text{\TeX}$  also produces few lines, but this is mainly due to the many lines that do not fit at all. This table and Figures 3 and 4 show that there is considerable difference in behavior and indicate that there is room for improvement in automatic line breaking.

Figure 7 shows the distribution of natural line widths (left) and spaces (right) of our probabilistic line breaking algorithm for different settings of  $\sigma^-$  and  $\sigma^+$ . The published text is shown as reference. Interestingly, for  $\sigma^- = 1.4$  and

<sup>5</sup> Available from <http://xml.apache.org>.

<sup>6</sup> Available from <http://www.renderx.com>.

<sup>7</sup> Available from <http://www.antennahouse.com>.

Figure 5: Space widths for published paragraphs and T<sub>E</sub>X and DLPager output (250/1000pt is optimal). Horizontal axis shows the space width in 1/100th of points, vertical axis shows the number of lines having the particular space width.




Figure 6: Space widths for published paragraphs and XSL-FO processors (250/1000pt is optimal). Axis as Figure 5.




Table 1: Number of lines per rendering engine

Algorithm	Number of lines
Published	74120
T <sub>E</sub> X	72561
T <sub>E</sub> X.sloppy	74719
DLPager	74456
Fop	78505
XEP	72476
Antenna	75730

$\sigma^+ = 0.3$  we get a distribution that is especially close to the published text, closer than any of the other systems examined. By changing the values to  $\sigma^- = 0.3$  and  $\sigma^+ = 1.4$ , we choose to accept tighter lines over loose lines and the line width distribution shifts to the right (so spaces get squeezed and its distribution goes to the left). Setting variances to  $\sigma^- = 10$  and  $\sigma^+ = 0.3$  gives a result left from the two. Remarkably, the probability of acceptability depends on the natural line widths only, not on the space width, like in T<sub>E</sub>X’s algorithm. Figure 7 shows that this does not result in overly tight lines though.


We compared the simple probabilistic algorithm with the one extended with some of the terms of Section 3.2 by creating graphs similar to Figure 7. The graphs showed that the distributions of natural line widths and space widths are very close. Informal inspection of the output shows that the choice of break points indeed is influenced as desired by tuning parameters of the network.

## 5 Conclusions

We showed how a probabilistic interpretation of an ill defined problem, the problem of finding line breaks in a paragraph, can lead to an efficient new algorithm. The graphical model that results from the probabilistic interpretation has the advantage that it is easy to tune. Thanks to the architecture of the network structure being a polytree, the algorithm is linear in the number of characters in a paragraph. Line breaks are optimized taking the whole paragraph in account. We demonstrated how adding metrics to capture subtle typographical rules can be done without impacting the efficiency. The resulting algorithm does not show threshold effects (like T<sub>E</sub>X) and it allows for easy incorporation of subtle typesetting properties. Empirical evidence suggests that this algorithm performs closer to results published through desk top publishing (DTP) than a number of existing systems.

There are a few limitations to our work. We only considered English text with US-English hyphenation rules in our experiments. However, different languages have other hyphenation rules and some even have different breaking rules, like east Asian languages [2]. Further, we only considered batch composi-

Figure 7: Natural line widths (left) and space widths (right) for published paragraphs and simple algorithm with various values of  $\sigma^-$  and  $\sigma^+$ . Axis as in Figure 3.



tion systems, but line breaking is routinely applied in word processors and DTP packages as well. These systems deal with incremental line breaking and the algorithm can be easily extended to deal with these limitations.

A bigger task is dealing with page breaks. It has been observed that page breaking can be interpreted as vertical version of the line breaking problem [1], so a probabilistic interpretation of this problem could be applied straightforwardly to find a new algorithm. An extra complication, however, is placement of floating objects, like figures, tables and footnotes. Also, there is the interaction between line breaks and paragraph breaks, making this a more challenging problem.

## References

- [1] Jonathan Fine. Line breaking and page breaking. TUGBoat, vol 21 no 3, 210–221, 2000.
- [2] Asmus Freytag. Line Breaking Properties Unicode Standard Annex #14 (part of the Unicode Standard). Technical Report, 2002.
- [3] D.E. Knuth. Computers & Typesetting Volume A, The TeXbook. Reading, Massachusetts: Addison-Wesley, 1984.
- [4] D.E Knuth and M.F. Plass. Breaking Paragraphs into Lines. Software—Practice and Experience 11, p.1119–1184, 1981.
- [5] S.L. Lauritzen and D.J. Spiegelhalter. Local computations with probabilities on graphical structures and their applications to expert systems (with discussion). Journal of the Royal Statistical Society B, 50:157–224, 1988.

- [6] Franklin M. Liang. Word Hy-phen-a-tion by Com-put-er. Ph.D. Thesis, Department of Computer Science, Stanford University, August 1983.
- [7] J. Pearl. Probabilistic Reasoning in Intelligent Systems, Networks of Plausible Inference. Morgan Kaufmann, 1998.
- [8] Extensible Stylesheet Language (XSL). Version 1.0, W3C Recommendation, 15 October 2001.

## Appendix: Data used in informal evaluation

The following pages show the published paragraphs side by side with DLPager output and output of our probabilistic algorithm mixed up as follows.

page	left	middle	right
1	published	pager	probabilistic
2	pager	probabilistic	published
3	probabilistic	published	pager
4	published	pager	probabilistic
5	pager	probabilistic	published
6	probabilistic	published	pager
7	published	pager	probabilistic
8	pager	probabilistic	published
9	probabilistic	published	pager
10	published	pager	probabilistic

The following table shows the rankings of the publishing managers of the paragraphs. The final row shows the sums of the rankings, providing a summary of the ranking. A lower final ranking indicates the particular method is preferred most.

Publication manager 1

Page	Live	Pager	Bayes
1	2	1	3
10	3	2	1
20	2	3	1
30	3	1	2
40	2	1	3
50	1	2	3
60	3	2	1
70	3	1	2
80	1	3	2
90	1	2	3
100	1	2	1
	22	20	22

Publication manager 2

Page	Live	Pager	Bayes
1	2	1	3
10	2	3	1
20	2	1	3
30	1	3	2
40	1	3	2
50	1	3	3
60	2	3	1
70	3	1	2
80	1	3	2
90	1	2	3
100	1	3	1
	17	26	23

Publication manager 3

Page	Live	Pager	Bayes
1	1	2	3
10	3	2	1
20	2	3	1
30	2	3	1
40	2	1	3
50	1	2	3
60	3	2	1
70	3	1	2
80	1	3	2
90	1	3	2
100	1	3	1
	20	25	20

Test pages used follow on next 11 pages.

### **1 a**

On 31 August 1998, the US Food and Drug Administration licensed the tetravalent rhesus rotavirus (RRV-TV) vaccine manufactured by Wyeth Lederle Vaccines and Pediatrics. On 16 July 1999, initial reports linking vaccination to intussusception appeared and use of the vaccine was suspended. Shortly thereafter, the Advisory Committee on Immunization Practices (ACIP) and the American Academy of Pediatrics withdrew their recommendations for the use of the vaccine. In this article, I will provide a background for understanding rotavirus vaccines, a brief discussion of the RRV-TV vaccine and a discussion of the other vaccines in clinical and preclinical development.

### **1 b**

On 31 August 1998, the US Food and Drug Administration licensed the tetravalent rhesus rotavirus (RRV-TV) vaccine manufactured by Wyeth Lederle Vaccines and Pediatrics. On 16 July 1999, initial reports linking vaccination to intussusception appeared and use of the vaccine was suspended. Shortly thereafter, the Advisory Committee on Immunization Practices (ACIP) and the American Academy of Pediatrics withdrew their recommendations for the use of the vaccine. In this article, I will provide a background for understanding rotavirus vaccines, a brief discussion of the RRV-TV vaccine and a discussion of the other vaccines in clinical and preclinical development.

### **1 c**

On 31 August 1998, the US Food and Drug Administration licensed the tetravalent rhesus rotavirus (RRV-TV) vaccine manufactured by Wyeth Lederle Vaccines and Pediatrics. On 16 July 1999, initial reports linking vaccination to intussusception appeared and use of the vaccine was suspended. Shortly thereafter, the Advisory Committee on Immunization Practices (ACIP) and the American Academy of Pediatrics withdrew their recommendations for the use of the vaccine. In this article, I will provide a background for understanding rotavirus vaccines, a brief discussion of the RRV-TV vaccine and a discussion of the other vaccines in clinical and preclinical development.

### 10 a

Several large efficacy studies were conducted with this vaccine.<sup>[32,39-41]</sup> In these trials, including more than 10 000 infants, the vaccine was generally well tolerated but caused low grade fever in 14 to 29% of recipients. During these trials, intussusception was identified in 5 of 10 054 vaccine recipients (2 of 8240 given the licensed dose) and 1 of 4633 placebo recipients.<sup>[42]</sup> Interestingly, none of these cases occurred after the first vaccination. Furthermore, the rates of intussusception were not significantly different and did not exceed the expected rates based on large surveillance databases from New York.<sup>[42]</sup>

### 10 b

Several large efficacy studies were conducted with this vaccine.<sup>[32,39-41]</sup> In these trials, including more than 10 000 infants, the vaccine was generally well tolerated but caused low grade fever in 14 to 29% of recipients. During these trials, intussusception was identified in 5 of 10 054 vaccine recipients (2 of 8240 given the licensed dose) and 1 of 4633 placebo recipients.<sup>[42]</sup> Interestingly, none of these cases occurred after the first vaccination. Furthermore, the rates of intussusception were not significantly different and did not exceed the expected rates based on large surveillance databases from New York.<sup>[42]</sup>

### 10 c

Several large efficacy studies were conducted with this vaccine.<sup>[32,39-41]</sup> In these trials, including more than 10 000 infants, the vaccine was generally well tolerated but caused low grade fever in 14 to 29% of recipients. During these trials, intussusception was identified in 5 of 10 054 vaccine recipients (2 of 8240 given the licensed dose) and 1 of 4633 placebo recipients.<sup>[42]</sup> Interestingly, none of these cases occurred after the first vaccination. Furthermore, the rates of intussusception were not significantly different and did not exceed the expected rates based on large surveillance databases from New York.<sup>[42]</sup>

## 20 a

The first attenuated human rotavirus vaccine candidate, M37, recovered from a stool of an asymptomatic neonate, was evaluated for safety and immunogenicity in several trials. The vaccine was found to be well tolerated but poorly immunogenic, especially for induction of serotype specific neutralising antibody.<sup>[60-62]</sup> No further development of this vaccine candidate is planned at present. Similarly, another neonatal strain, RV3, was also shown to be well tolerated but replicated poorly and induced an immune response in only 2 of 5 recipients.<sup>[63]</sup> Most recently, 2 new vaccine candidates have been developed from neonatal strains isolated in India. The isolate from Delhi (116E) is a G<sub>9</sub>P[11] strain, whereas the one from Bangalore (I321) is a G<sub>10</sub>P[11] strain. Both appear to be natural reassortants of human and bovine rotavirus strains.<sup>[64,65]</sup> Both vaccine strains were adapted to tissue culture, passaged multiple times and appeared well tolerated in trials in adults (my unpublished work). Further development of these vaccines is being planned in India.

## 20 b

The first attenuated human rotavirus vaccine candidate, M37, recovered from a stool of an asymptomatic neonate, was evaluated for safety and immunogenicity in several trials. The vaccine was found to be well tolerated but poorly immunogenic, especially for induction of serotype specific neutralising antibody.<sup>[60-62]</sup> No further development of this vaccine candidate is planned at present. Similarly, another neonatal strain, RV3, was also shown to be well tolerated but replicated poorly and induced an immune response in only 2 of 5 recipients.<sup>[63]</sup> Most recently, 2 new vaccine candidates have been developed from neonatal strains isolated in India. The isolate from Delhi (116E) is a G<sub>9</sub>P[11] strain, whereas the one from Bangalore (I321) is a G<sub>10</sub>P[11] strain. Both appear to be natural reassortants of human and bovine rotavirus strains.<sup>[64,65]</sup> Both vaccine strains were adapted to tissue culture, passaged multiple times and appeared well tolerated in trials in adults (my unpublished work). Further development of these vaccines is being planned in India.

## 20 c

The first attenuated human rotavirus vaccine candidate, M37, recovered from a stool of an asymptomatic neonate, was evaluated for safety and immunogenicity in several trials. The vaccine was found to be well tolerated but poorly immunogenic, especially for induction of serotype specific neutralising antibody.<sup>[60-62]</sup> No further development of this vaccine candidate is planned at present. Similarly, another neonatal strain, RV3, was also shown to be well tolerated but replicated poorly and induced an immune response in only 2 of 5 recipients.<sup>[63]</sup> Most recently, 2 new vaccine candidates have been developed from neonatal strains isolated in India. The isolate from Delhi (116E) is a G<sub>9</sub>P[11] strain, whereas the one from Bangalore (I321) is a G<sub>10</sub>P[11] strain. Both appear to be natural reassortants of human and bovine rotavirus strains.<sup>[64,65]</sup> Both vaccine strains were adapted to tissue culture, passaged multiple times and appeared well tolerated in trials in adults (my unpublished work). Further development of these vaccines is being planned in India.

### **30 a**

Although the incidence and severity of acute rejection may be one factor that increases the risk of chronic rejection, it may not be the principal one in many patients. It is important to recognise that the frequency of rejection also varies according to the primary underlying liver disease. Thus, patients with hepatitis B or alcoholic liver disease have relatively low rejection rates compared with patients with primary biliary cirrhosis (PBC) or primary sclerosing cholangitis (PSC), whose rates range between 20 and 70%.

### **30 b**

Although the incidence and severity of acute rejection may be one factor that increases the risk of chronic rejection, it may not be the principal one in many patients. It is important to recognise that the frequency of rejection also varies according to the primary underlying liver disease. Thus, patients with hepatitis B or alcoholic liver disease have relatively low rejection rates compared with patients with primary biliary cirrhosis (PBC) or primary sclerosing cholangitis (PSC), whose rates range between 20 and 70%.

### **30 c**

Although the incidence and severity of acute rejection may be one factor that increases the risk of chronic rejection, it may not be the principal one in many patients. It is important to recognise that the frequency of rejection also varies according to the primary underlying liver disease. Thus, patients with hepatitis B or alcoholic liver disease have relatively low rejection rates compared with patients with primary biliary cirrhosis (PBC) or primary sclerosing cholangitis (PSC), whose rates range between 20 and 70%.

#### 40 a

PBC remains a common indication for liver transplantation, particularly in Northern Europe, and many centres report 5-year survival >80%. However, a number of studies have shown that patients undergoing transplant for PBC are at greater risk of developing both acute and chronic rejection.<sup>[26,27]</sup> In our own series between 1994 and 1997, acute early rejection (<28 days) occurred in 43% of PBC transplant patients, compared with only 20% of those receiving grafts for alcoholic liver disease and 12% in hepatitis B disease.<sup>[28]</sup> One hypothesis suggests that the genetic susceptibility that predisposes patients to develop PBC might be the same one that allows rejection, with the biliary epithelial cells being the immune target in patients already sensitised to biliary antigens.<sup>[29]</sup> If that is the case, then the level of immunosuppressive therapy could affect the pattern of recurrence and chronic rejection.

#### 40 b

PBC remains a common indication for liver transplantation, particularly in Northern Europe, and many centres report 5-year survival >80%. However, a number of studies have shown that patients undergoing transplant for PBC are at greater risk of developing both acute and chronic rejection.<sup>[26,27]</sup> In our own series between 1994 and 1997, acute early rejection (<28 days) occurred in 43% of PBC transplant patients, compared with only 20% of those receiving grafts for alcoholic liver disease and 12% in hepatitis B disease.<sup>[28]</sup> One hypothesis suggests that the genetic susceptibility that predisposes patients to develop PBC might be the same one that allows rejection, with the biliary epithelial cells being the immune target in patients already sensitised to biliary antigens.<sup>[29]</sup> If that is the case, then the level of immunosuppressive therapy could affect the pattern of recurrence and chronic rejection.

#### 40 c

PBC remains a common indication for liver transplantation, particularly in Northern Europe, and many centres report 5-year survival >80%. However, a number of studies have shown that patients undergoing transplant for PBC are at greater risk of developing both acute and chronic rejection.<sup>[26,27]</sup> In our own series between 1994 and 1997, acute early rejection (<28 days) occurred in 43% of PBC transplant patients, compared with only 20% of those receiving grafts for alcoholic liver disease and 12% in hepatitis B disease.<sup>[28]</sup> One hypothesis suggests that the genetic susceptibility that predisposes patients to develop PBC might be the same one that allows rejection, with the biliary epithelial cells being the immune target in patients already sensitised to biliary antigens.<sup>[29]</sup> If that is the case, then the level of immunosuppressive therapy could affect the pattern of recurrence and chronic rejection.

### **50 a**

The 2 main abnormalities in chronic rejection are loss of bile ducts and an obliterative arteriopathy affecting large and medium sized arteries. Less specific are changes in the liver parenchyma.<sup>[40-42]</sup>

### **50 b**

The 2 main abnormalities in chronic rejection are loss of bile ducts and an obliterative arteriopathy affecting large and medium sized arteries. Less specific are changes in the liver parenchyma.<sup>[40-42]</sup>

### **50 c**

The 2 main abnormalities in chronic rejection are loss of bile ducts and an obliterative arteriopathy affecting large and medium sized arteries. Less specific are changes in the liver parenchyma.<sup>[40-42]</sup>

### 60 a

After the 1980s, the decade of cyclosporin triple therapy, the introduction of tacrolimus into clinical liver transplantation in the early 1990s had a further significant impact on the options for control of rejection. Like cyclosporin, tacrolimus has inhibitory effects on T lymphocyte activation, binding to FK506 binding protein and inhibiting calcineurin. Tacrolimus may be given intravenously (with a higher incidence of adverse events) at 0.01 to 0.1 mg/kg/day (0.03 to 0.1 mg/kg/day in children) or orally at 0.1 to 0.3 mg/kg/day (0.15 to 0.3 mg/kg/day in children). It has confirmed efficacy as a primary or rescue therapy, with chronic rejection rates significantly lower than cyclosporin, but with graft and patient survival similar. The adverse effects of tacrolimus are dosage related. Tacrolimus is mainly metabolised in the liver and the main excretory pathway of metabolites is biliary.

### 60 b

After the 1980s, the decade of cyclosporin triple therapy, the introduction of tacrolimus into clinical liver transplantation in the early 1990s had a further significant impact on the options for control of rejection. Like cyclosporin, tacrolimus has inhibitory effects on T lymphocyte activation, binding to FK506 binding protein and inhibiting calcineurin. Tacrolimus may be given intravenously (with a higher incidence of adverse events) at 0.01 to 0.1 mg/kg/day (0.03 to 0.1 mg/kg/day in children) or orally at 0.1 to 0.3 mg/kg/day (0.15 to 0.3 mg/kg/day in children). It has confirmed efficacy as a primary or rescue therapy, with chronic rejection rates significantly lower than cyclosporin, but with graft and patient survival similar. The adverse effects of tacrolimus are dosage related. Tacrolimus is mainly metabolised in the liver and the main excretory pathway of metabolites is biliary.

### 60 c

After the 1980s, the decade of cyclosporin triple therapy, the introduction of tacrolimus into clinical liver transplantation in the early 1990s had a further significant impact on the options for control of rejection. Like cyclosporin, tacrolimus has inhibitory effects on T lymphocyte activation, binding to FK506 binding protein and inhibiting calcineurin. Tacrolimus may be given intravenously (with a higher incidence of adverse events) at 0.01 to 0.1 mg/kg/day (0.03 to 0.1 mg/kg/day in children) or orally at 0.1 to 0.3 mg/kg/day (0.15 to 0.3 mg/kg/day in children). It has confirmed efficacy as a primary or rescue therapy, with chronic rejection rates significantly lower than cyclosporin, but with graft and patient survival similar. The adverse effects of tacrolimus are dosage related. Tacrolimus is mainly metabolised in the liver and the main excretory pathway of metabolites is biliary.

### **70 a**

In addition, the use of tacrolimus with the antiproliferative agent sirolimus (rapamycin) is being explored, as well as its use with mycophenolate mofetil.

### **70 b**

In addition, the use of tacrolimus with the antiproliferative agent sirolimus (rapamycin) is being explored, as well as its use with mycophenolate mofetil.

### **70 c**

In addition, the use of tacrolimus with the antiproliferative agent sirolimus (rapamycin) is being explored, as well as its use with mycophenolate mofetil.

### 80 a

Patients undergoing liver transplantation often do so in an advanced stage of liver disease, with many in intensive care before transplantation. For most, a return to full activities can be expected, with high quality of life. A reduction in incapacitating symptoms such as fatigue, pruritus and osteopenia is combined with general improvement in the sense of well-being and nutrition in most patients. Gross et al.,<sup>[91]</sup> analysing 157 cholestatic patients, showed statistically significant improvement by 1 year in quality of life after liver grafting. Patients reported fewer health problems or limitations in mobility and enhanced social and sexual life, and more objective rating evaluations showed a reduction in poor quality of life and Karnofsky scoring. The majority of patients rated their quality of life and health status as good or excellent, with only 7% reporting significant limitation in life activities because of health problems.<sup>[91]</sup> Although some lifestyle limitations and reductions in quality of life, such as hypertension, bodyweight gain and obesity, headaches and other comorbidity, relate to the immunosuppressive drugs,<sup>[92]</sup> many of the underlying limitations represent the long term effects of advanced liver disease present before transplantation.

### 80 b

Patients undergoing liver transplantation often do so in an advanced stage of liver disease, with many in intensive care before transplantation. For most, a return to full activities can be expected, with high quality of life. A reduction in incapacitating symptoms such as fatigue, pruritus and osteopenia is combined with general improvement in the sense of well-being and nutrition in most patients. Gross et al.,<sup>[91]</sup> analysing 157 cholestatic patients, showed statistically significant improvement by 1 year in quality of life after liver grafting. Patients reported fewer health problems or limitations in mobility and enhanced social and sexual life, and more objective rating evaluations showed a reduction in poor quality of life and Karnofsky scoring. The majority of patients rated their quality of life and health status as good or excellent, with only 7% reporting significant limitation in life activities because of health problems.<sup>[91]</sup> Although some lifestyle limitations and reductions in quality of life, such as hypertension, bodyweight gain and obesity, headaches and other comorbidity, relate to the immunosuppressive drugs,<sup>[92]</sup> many of the underlying limitations represent the long term effects of advanced liver disease present before transplantation.

### 80 c

Patients undergoing liver transplantation often do so in an advanced stage of liver disease, with many in intensive care before transplantation. For most, a return to full activities can be expected, with high quality of life. A reduction in incapacitating symptoms such as fatigue, pruritus and osteopenia is combined with general improvement in the sense of well-being and nutrition in most patients. Gross et al.,<sup>[91]</sup> analysing 157 cholestatic patients, showed statistically significant improvement by 1 year in quality of life after liver grafting. Patients reported fewer health problems or limitations in mobility and enhanced social and sexual life, and more objective rating evaluations showed a reduction in poor quality of life and Karnofsky scoring. The majority of patients rated their quality of life and health status as good or excellent, with only 7% reporting significant limitation in life activities because of health problems.<sup>[91]</sup> Although some lifestyle limitations and reductions in quality of life, such as hypertension, bodyweight gain and obesity, headaches and other comorbidity, relate to the immunosuppressive drugs,<sup>[92]</sup> many of the underlying limitations represent the long term effects of advanced liver disease present before transplantation.

### 90 a

The bioavailability of rIL-2 might differ according to the route of administration or dose schedule. Kirchner et al.<sup>[41]</sup> recently demonstrated that twice-daily subcutaneous administration of rIL-2 provides better bioavailability compared with once-daily administration, and concluded that it might be the preferred way of administration. Available data, however, do not confirm the superiority of any particular dose level or schedule of rIL-2. Randomised trials are therefore required to document any improvement in response rates and durations with different treatment schedules. Until then, the decision regarding the treatment schedule and route of administration must be individualised based on patients' clinical status, comorbidities and experience, the familiarity of the treating physician with rIL-2, and whether inpatient or outpatient therapy is desired.

### 90 b

The bioavailability of rIL-2 might differ according to the route of administration or dose schedule. Kirchner et al.<sup>[41]</sup> recently demonstrated that twice-daily subcutaneous administration of rIL-2 provides better bioavailability compared with once-daily administration, and concluded that it might be the preferred way of administration. Available data, however, do not confirm the superiority of any particular dose level or schedule of rIL-2. Randomised trials are therefore required to document any improvement in response rates and durations with different treatment schedules. Until then, the decision regarding the treatment schedule and route of administration must be individualised based on patients' clinical status, comorbidities and experience, the familiarity of the treating physician with rIL-2, and whether inpatient or outpatient therapy is desired.

### 90 c

The bioavailability of rIL-2 might differ according to the route of administration or dose schedule. Kirchner et al.<sup>[41]</sup> recently demonstrated that twice-daily subcutaneous administration of rIL-2 provides better bioavailability compared with once-daily administration, and concluded that it might be the preferred way of administration. Available data, however, do not confirm the superiority of any particular dose level or schedule of rIL-2. Randomised trials are therefore required to document any improvement in response rates and durations with different treatment schedules. Until then, the decision regarding the treatment schedule and route of administration must be individualised based on patients' clinical status, comorbidities and experience, the familiarity of the treating physician with rIL-2, and whether inpatient or outpatient therapy is desired.

**100 a**

A variety of rIL-2 regimens are being explored in the post SCT setting. In theory, rIL-2 should be more effective after SCT than after chemotherapy for a variety of reasons, including:

**100 b**

A variety of rIL-2 regimens are being explored in the post SCT setting. In theory, rIL-2 should be more effective after SCT than after chemotherapy for a variety of reasons, including:

**100 c**

A variety of rIL-2 regimens are being explored in the post SCT setting. In theory, rIL-2 should be more effective after SCT than after chemotherapy for a variety of reasons, including: