

# Open Research Online

---

The Open University's repository of research publications and other research outputs

## Surprises in approximating Levenshtein distances

### Journal Item

How to cite:

Baake, Michael; Grimm, Uwe and Giegerich, Robert (2006). Surprises in approximating Levenshtein distances. *Journal of Theoretical Biology*, 243(2) pp. 279–282.

For guidance on citations see [FAQs](#).

© [\[not recorded\]](#)

Version: [\[not recorded\]](#)

Link(s) to article on publisher's website:

<http://dx.doi.org/doi:10.1016/j.jtbi.2006.06.026>

[http://www.elsevier.com/wps/find/journaldescription.cws\\_home/622904/description#description](http://www.elsevier.com/wps/find/journaldescription.cws_home/622904/description#description)

---

Copyright and Moral Rights for the articles on this site are retained by the individual authors and/or other copyright owners. For more information on Open Research Online's data [policy](#) on reuse of materials please consult the policies page.

---

[oro.open.ac.uk](http://oro.open.ac.uk)

# Surprises in approximating Levenshtein distances

Michael Baake<sup>a</sup>, Uwe Grimm<sup>b</sup> and Robert Giegerich<sup>c</sup>

<sup>a</sup>Fakultät für Mathematik, Universität Bielefeld, Postfach 100131, 33501 Bielefeld, Germany

<sup>b</sup>Department of Mathematics, The Open University, Milton Keynes MK7 6AA, UK

<sup>c</sup>Technische Fakultät, Universität Bielefeld, Postfach 100131, 33501 Bielefeld, Germany

## Abstract

The Levenshtein distance is an important tool for the comparison of symbolic sequences, with many appearances in genome research, linguistics and other areas. For efficient applications, an approximation by a distance of smaller computational complexity is highly desirable. However, our comparison of the Levenshtein with a generic dictionary-based distance indicates their statistical independence. This suggests that a simplification along this line might not be possible without restricting the class of sequences. Several other probabilistic properties are briefly discussed, emphasizing various questions that deserve further investigation.

**Keywords:** global alignment; distance concepts; statistical independence; computational complexity

## 1 Introduction

The Levenshtein (or edit) metric (Levenshtein, 1965) is a standard tool to estimate the distance between two sequences. It is widely used in linguistics and bioinformatics, and for the recognition of text blocks with isolated mistakes. As is well known, its computational complexity, when applied to two sequences of (approximately) the same length  $n$ , is  $\mathcal{O}(n^2)$ . Since this is a hurdle in many practical applications, it is desirable to replace, or to approximate, the Levenshtein (L) distance by some quantity of smaller (preferably linear) computational complexity. Two fast approximation algorithms for edit distances were suggested by Ukkonen (1992), one based on maximal exact matches, the other on suitably restricted subword comparisons between the two sequences; compare also Lippert *et al.* (2002). This would indeed give  $\mathcal{O}(n)$ , due to their computability from the suffix tree; see Gusfield (1999). However, they only provide lower bounds, and hence no complete solution of the problem.

It seems possible to estimate probabilistically, with sublinear complexity, whether the L-distance of two sequences is ‘small’ or ‘large’; see Batu *et al.* (2003). Whether an improvement of this rather coarse result or even a replacement of the L-distance is possible,

with at most linear complexity and a non-probabilistic outcome, seems open. Below, we compare the L-distance with a representative dictionary-based distance. Our findings support the conclusion that such a simplification might be difficult or even impossible. On the way, we highlight some interesting properties that have been neglected so far, but seem relevant for a better understanding of such distance concepts.

## 2 Comparison of two distances

To keep discussion and results transparent, we concentrate on two specific distances, and on binary sequences. We have also tried a number of obvious alternatives, but they did not show any significantly different behaviour. In this sense, the structure of our example is more likely typical than exceptional.

The L-distance  $d_L(u, v)$  of two sequences  $u$  and  $v$  (not necessarily of equal length) is the minimum number of edit operations (insertions, deletions, or substitutions) needed to transform  $u$  into  $v$  or vice versa (Gusfield, 1999, Ch. 11.2). Though  $d_L(u, v)$  is closely related to the longest common subsequence (LCS) (loc. cit., Ch. 11.6.2) of  $u$  and  $v$  (and hence to distances based upon it), one important difference lies in the possibility of substitutions. So, using the LCS in this context requires some care. For sequences of lengths  $m$  and  $n$ , the computational complexity of calculating  $d_L$  (or the LCS) is  $\mathcal{O}(mn)$ , e.g., when based on the Needleman-Wunsch algorithm; see (Ewens and Grant, 2004, Ch. 6.4.2).

A generic choice for a dictionary-based metric is

$$d_D(u, v) = \text{card}(\mathcal{A}(u) \Delta \mathcal{A}(v)),$$

where  $\mathcal{A}(u)$  is the full dictionary of  $u$ , i.e., the set of all non-empty subwords of  $u$ , and  $A \Delta B = (A \cup B) \setminus (A \cap B)$  is the symmetric difference of  $A$  and  $B$ . This choice actually disregards the goal of computational simplification, but focuses on the full dictionary information instead, and thus, in some sense, represents the optimal information on the sequences to be compared. It is well known that, using the suffix tree structure, the calculation of closely related dictionary-based distances is possible with linear complexity, e.g., by means of Ukkonen's algorithm; compare (Gusfield, 1999, Ch. 6). On the other hand, further restrictions are likely to reduce the usefulness in relation to the L-distance.

Both  $d_L$  and  $d_D$  define a *metric*, i.e., for arbitrary sequences  $u, v$  and  $w$ , the distance  $d \in \{d_L, d_D\}$  satisfies the axioms of a metric (Schechter, 1997, Ch. 2.11):

- (i)  $0 \leq d(u, v) < \infty$  (positivity);
- (ii)  $d(u, v) = 0$  if and only if  $u = v$  (non-degeneracy);
- (iii)  $d(u, v) = d(v, u)$  (symmetry);
- (iv)  $d(u, v) \leq d(u, w) + d(w, v)$  (triangle inequality).

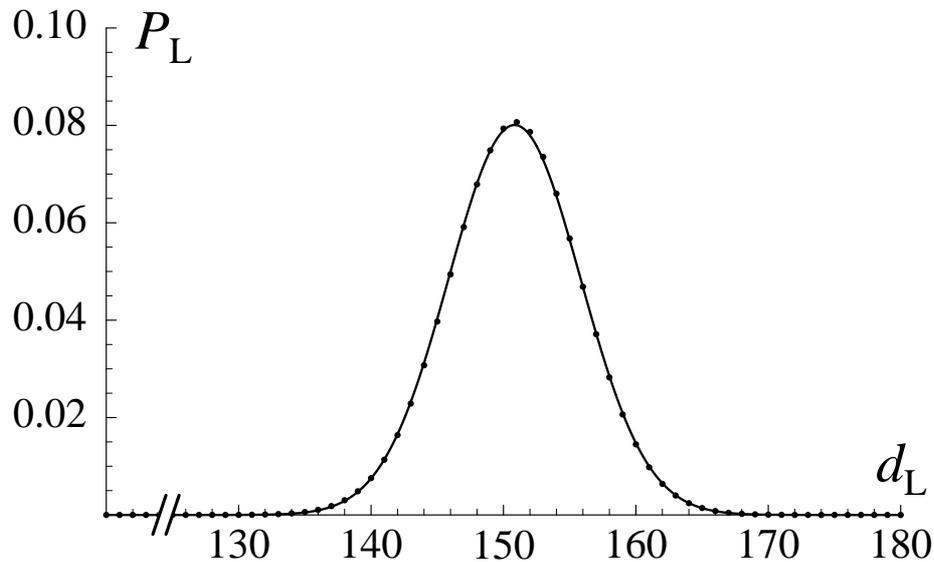


Figure 1: Simulated probability distribution  $P_L$  of the L-distance  $d_L$  between two random sequences of length 500 (dots) and Gaussian approximation (line), with mean 150.84 and variance 24.82.

Less clear is the relation between  $d_L$  and  $d_D$ . Since one can easily construct pairs of sequences that are close in one, but not in the other distance, they are certainly not equivalent in the strong sense as also used for norms, compare (Werner, 1995, Ch. I.2). They are equivalent in the weaker sense of generating the same topology (Schechter, 1997, Ch. 22.5), which is the discrete topology here. However, this is of little use for the question addressed above. The situation does not improve if one replaces  $d(u, v)$  by the quotient  $d(u, v)/(1 + d(u, v))$ , which is another metric, with range in  $[0, 1]$ . As we shall see below, the situation is actually much worse.

### 3 Concrete results

To get a first impression of the L-distance, we computed the discrete probability distribution of the values  $d_L(u, v)$  for sequences  $u \neq v$  of the same length, under uniform distribution on sequence space. This has long been known to be a reasonable first approach for the comparison of sequences from data bases (Reich *et al.*, 1984). Up to length 20, this was done using all possible pairs; for longer sequences, the distribution was estimated from a sufficiently large random selection of pairs. For length  $n = 500$ , the result obtained from  $4 \times 10^8$  pairs is shown in Figure 1. For large  $n$ , the distributions seem to be well described by Gaussian (or normal) distributions. This qualitative behaviour does not change much and seems to improve with sequence length. One could add weight to this finding by per-

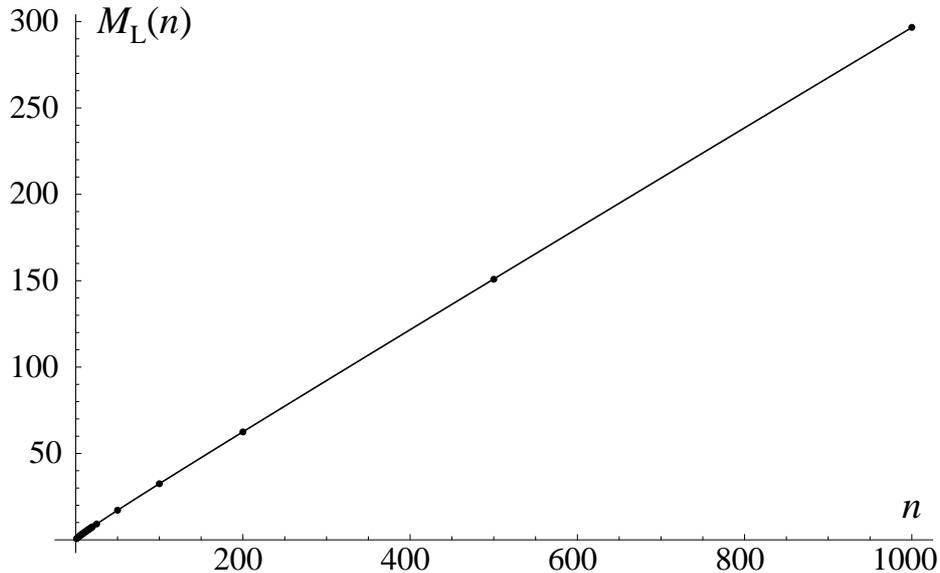


Figure 2: Mean  $M_L(n)$  of the probability distribution  $P_L$  as a function of sequence length  $n$ , calculated exactly for  $n \leq 20$  and by simulation otherwise. The solid line shows the least squares fit  $M_L(n) = 0.413\sqrt{n} + 0.283n$ .

forming a statistical test on Gaussianity, which would score well. However, we think that one should not over-interpret this observation, in particular in view of a recent numerical investigation by Pang *et al.* (2005) which indicates that a gamma distribution might give an even better description.

Note that, if extrema over *local* alignments are taken, one obtains an extremal value distribution (Pearson and Wood, 2001, Ch. 2.3.2). However, this implies nothing for the *global* alignment considered here. The possible (or approximate) Gaussian nature of this case has been observed before by Dayhoff, see (Mount, 2001, Ch. 3) and references given there; a more detailed investigation of tail probabilities can be found in Waterman (1994). Still, it seems to be hardly noted, although it is a relevant phenomenon that deserved further attention, with exact results presently not in sight.

For this reason, we could only investigate our findings numerically. Beyond checking the Gaussian behaviour qualitatively, means and variances were calculated for different  $n$ , both by exact enumeration (for  $n \leq 20$ ) and by simulation (for larger  $n$ , up to  $n = 1000$ ). It is an interesting question whether the mean and the variance, as functions of sequence length, show power-law behaviour, at least asymptotically. Our data, see Figures 2 and 3, are compatible with an asymptotically linear growth of the mean and an asymptotic  $n^{2/3}$  power law for the variance, both with a square-root correction term (for which we do not have any particular justification). Such predictions and conjectures are presently discussed by various people (Matzinger, 2004). In particular, the  $n^{2/3}$  power law for the

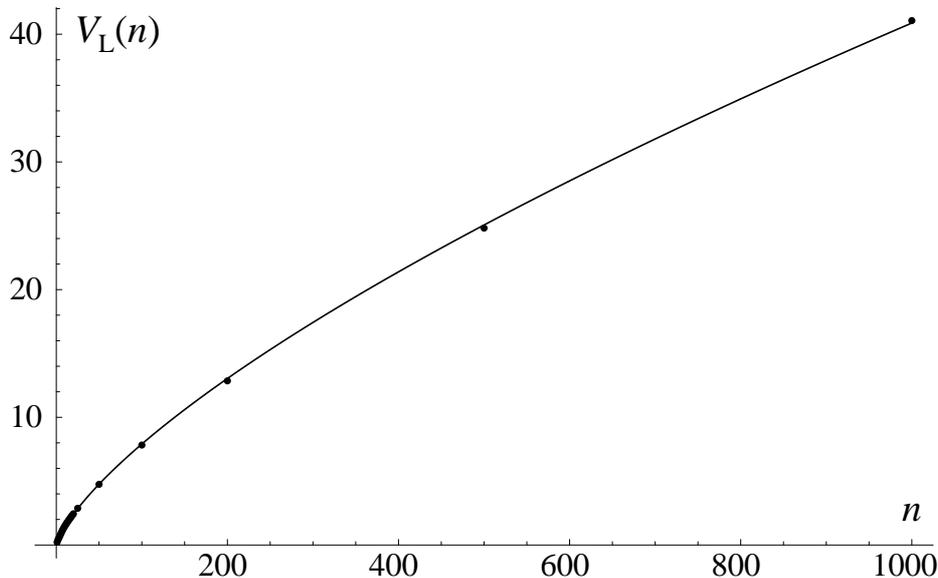


Figure 3: Variance  $V_L(n)$  of the probability distribution  $P_L$  as a function of sequence length  $n$ , calculated exactly for  $n \leq 20$  and by simulation otherwise. The solid line shows the least squares fit  $V_L(n) = -0.283\sqrt{n} + 0.498 n^{2/3}$ .

variance would be in line with analogous observations for the LCS, compare Hwa and Lässig (1996). Since there has recently been some doubt in the correctness of this finding (Matzinger, 2004), it requires further corroboration and investigation.

A similar finding (though with larger fluctuations) applies to the distribution of the values  $d_D(u, v)$  for random pairs  $u \neq v$ . However, there is no compelling reason to investigate this specific distance in detail, as it was mainly selected for illustrative purposes and does not seem to be closely related to one of the standard problems of probability theory.

More interesting, and also more relevant, is the question for the *joint* distribution of  $d_D(u, v)$  and  $d_L(u, v)$ . A necessary requirement for a useful relation between the two distances would be a strong correlation. However, as Figure 4 shows for sequences of length 100, there is little correlation at all – the joint distribution is rather well described by the product of the two Gaussians needed for the marginal distributions. This observation could be quantified with some effort, but we refrain from doing so because it would not contribute to the interpretation at this stage.

Our finding means that, at least on the level of the full sequence space or for the alignment of two random sequences (as analyzed in our simulations), the distances  $d_D(u, v)$  and  $d_L(u, v)$  are closer to being statistically independent of each other than to being useful approximations of one another.

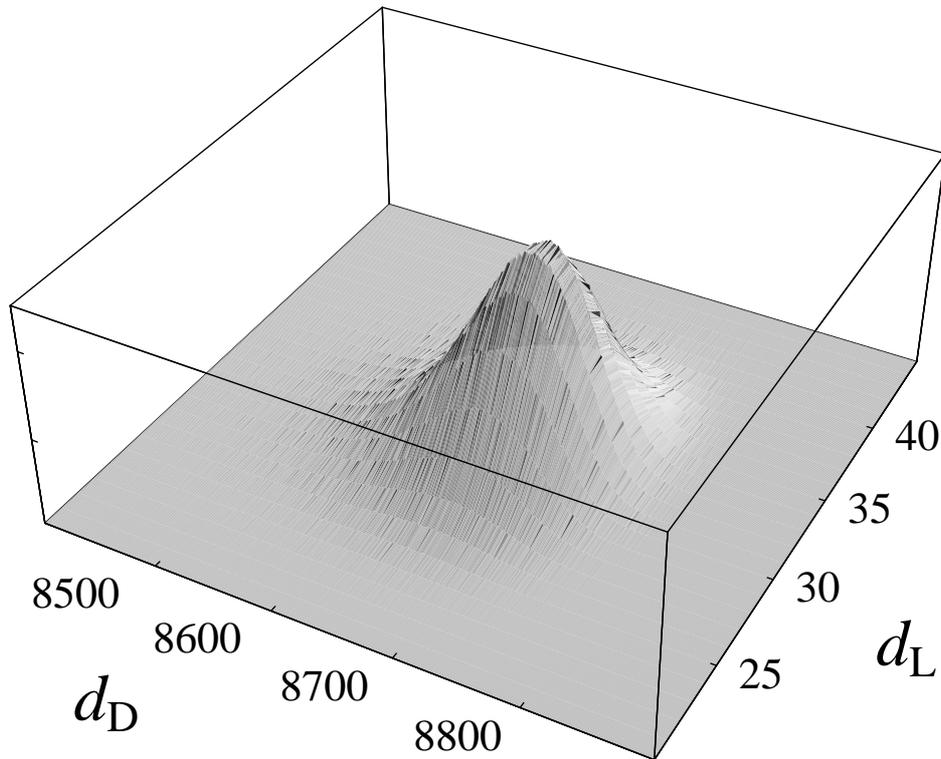


Figure 4: Numerical approximation of the joint probability distribution  $P$  for  $d_D$  and  $d_L$ , obtained from a simulation with  $10^7$  random pairs of sequences of length 100. Both marginal distributions are approximately Gaussian, with slightly larger fluctuations for the distribution of  $d_D$ -values.

## 4 Concluding remarks

Our findings are to be interpreted with care. They do not rule out a simplified approach to L-type distances, at least when restricted to (possibly relevant) subsets of sequences. However, they seem to indicate that subword comparison leads to statistically independent information, at least when viewed on the full sequence space. Clearly, different distance concepts can and should be tried. Moreover, a rigorous stochastic analysis of the various limit distributions is necessary to clarify the picture obtained from the simulations.

As long as analytic results (e.g., via limit theorems) are unavailable, it would also help to perform a more detailed statistical analysis of the various distributions, including clear-cut statistical tests. In particular, it would be extremely relevant to also consider suitable subspaces of the full sequence space, such as those extractable from existing data bases. Though this is clearly far beyond the scope of this short note, we believe that it would be a rewarding task for future investigations.

## Acknowledgements

It is our pleasure to thank E. Baake and D. Lenz for helpful discussions, and F. Merkl and M. Vingron for useful hints on the literature. Financial support from British Council (ARC 1213) and DAAD is gratefully acknowledged.

## References

- Batu, T., Ergün, F., Kilian, J., Magen, A., Raskhodnikova, S., Rubinfeld, R., and Sami, R. 2003. A sublinear algorithm for weakly approximating edit distance, 316–324. In *Proc. of the 35th Annual ACM Symp. on Theory of Computing*. ACM Press, New York, NY.
- Ewens, W.J., and Grant, G.R. 2004. *Statistical Methods in Bioinformatics*. 2nd ed., Springer, New York.
- Gusfield, D. 1999. *Algorithms on Strings, Trees, and Sequences*. Corr. reprint, Cambridge University Press, New York, NY.
- Hwa, T., and Lässig, M. 1996. Similarity detection and localization. *Phys. Rev. Lett.* 76, 2591–2594.
- Levenshtein, V.I. 1965. Binary codes capable of correcting deletions, insertions, and reversals. *Soviet Physics Dokl.* 10, 707–710.
- Lippert, R.A., Huang, H., and Waterman, M.S. 2002. Distributional regimes for the number of  $k$ -word matches between two random sequences. *Proc. Natl. Acad. Sci. USA* 99, 13980–13989.
- Matzinger, H. 2004. Private communication.
- Mount, D.W. 2001. *Bioinformatics – Sequence and Genome Analysis*. Cold Spring Harbor Laboratory Press, Cold Spring Harbor, NY.
- Pang, H., Tang, J., Chen, S.-S., and Tao, S. 2005. Statistical distributions of optimal global alignment scores of random protein sequences. *BMC Bioinformatics* 6, 257 (9 pages).
- Pearson, W.R., and Wood, T.C. 2001. Statistical significance in biological sequence comparison, 39–65. In Balding, D.J., Bishop, M., and Cannings, C., eds., *Handbook of Statistical Genetics*. Wiley, Chichester.
- Reich, J.G., Drabsch, H., and Däumler, A. 1984. On the statistical assessment of similarities in DNA sequences. *Nucleic Acids Res.* 12, 5529–5543.
- Schechter, E. 1997. *Handbook of Analysis and its Foundations*. Academic Press, San Diego, CA.
- Ukkonen, E. 1992. Approximate string-matching with  $q$ -grams and maximal matches. *Theor. Comp. Sci.* 92, 191–211.
- Waterman, M. 1994. Estimating statistical significance of sequence alignments. *Philos. Trans. R. Soc. London B* 344, 383–390.
- Werner, D. 1995. *Funktionalanalysis*. Springer, Berlin.