



Searching and indexing genomic databases via kernelization

Travis Gagie* and Simon J. Puglisi

Helsinki Institute for Information Technology (HIIT) and Department of Computer Science, University of Helsinki, Helsinki, Finland

Edited by:

Marco Pellegrini, Consiglio Nazionale delle Ricerche, Italy

Reviewed by:

Thierry Lecroq, University of Rouen, France

Sebastian Wandelt, Humboldt University of Berlin, Germany

*Correspondence:

Travis Gagie, Department of Computer Science, University of Helsinki, P.O. Box 68 (Gustaf Hällströmin katu 2b), Helsinki FI-00014, Finland
e-mail: travis.gagie@cs.helsinki.fi

The rapid advance of DNA sequencing technologies has yielded databases of thousands of genomes. To search and index these databases effectively, it is important that we take advantage of the similarity between those genomes. Several authors have recently suggested searching or indexing only one reference genome and the parts of the other genomes where they differ. In this paper, we survey the 20-year history of this idea and discuss its relation to kernelization in parameterized complexity.

Keywords: approximate pattern matching, data compression, genomic databases, indexing, kernelization, random-access reading, string algorithms

1. INTRODUCTION

The Human Genome Project took 13 years and three billion dollars to sequence a human genome, but the latest next-generation sequencing methods take only a few days and a few thousand dollars. With these methods, initiatives such as the 1000 Genomes Project and the 100,000 Genomes Project are now feasible. Advances in sequencing have far outstripped advances in computer processors and random-access memory, however, so it is increasingly challenging to make use of the data available. For example, while modern aligners can easily hold in memory the index for approximate pattern matching on a single human genome, they cannot handle thousands of human genomes. Schneeberger et al. (2009) proposed that we index the common parts of the genomes only once for them all, but we index the parts near variation sites for each genome. Ferrada et al. (2014b) suggested indexing the parts of all the genomes near boundaries between phrases in the LZ77 parse of the database. This is more general and may give better compression but requires the LZ77 parse, which is difficult to compute when the database does not fit in memory. Wandelt et al. (2013) proposed using a modified parse in which phrases must occur in a reference genome, which is easier to compute. (When papers have appeared in journals we cite those versions, although their chronological order may differ from that of previous versions.) Danek et al. (2014) recently showed that with this general approach we can store an index for approximate pattern matching on the database from the 1000 Genomes Project, in the memory of a commodity personal computer. This has so far not been possible with competing approaches, as surveyed by Vyverman et al. (2012).

When we are not given an upper bound on the pattern length, we can use one of the competing indexes that does not require such a bound or we can scan, with an online pattern-matching algorithm, the reference genome, and the parts of the other genomes near phrase boundaries. Wandelt and Leser (2012) and Rahn et al.

(2014) proposed the latter idea specifically for approximate pattern matching in genomic databases, but the general approach has a 20-year history in the field of compressed pattern matching. In this paper, we survey that history and relate it to current research: in Section 2 we discuss some relevant data compression schemes and how they have been augmented to support fast random-access reading; in Section 3 we discuss how they have been used to speed up pattern-matching; in Section 4 we discuss how they have been used in compressed indexing. While writing this survey, we realized that scanning or indexing only parts of the database and then mapping the solution for those parts onto a solution for the whole database, is like kernelization in parameterized complexity (We note that kernels in parameterized complexity bear no relation to operating system kernels nor to kernels in machine learning.). We emphasize this perspective because we feel that computing a pattern-matching kernel is an interesting problem in itself, regardless of how we process it later, and deserving of further study. Of course, the nature and even the existence of the kernel depend on the problem we are trying to solve.

2. COMPRESSION WITH RANDOM-ACCESS READING

In general, the best compression of highly repetitive datasets is achieved with the LZ77 algorithm by Ziv and Lempel (1977). Suppose $S[1, \dots, n]$ is a string with $S[n] = \$$, which is an end-of-file symbol that does not occur elsewhere in S . LZ77 works by parsing S into phrases such that, for each phrase $S[i, \dots, j]$, $S[i, \dots, j-1]$ occurs in $S[1, \dots, j-2]$ but $S[i, \dots, j]$ does not occur in $S[1, \dots, j-1]$; that phrase is stored as a triple consisting of a pointer to $S[i, \dots, j]$'s first occurrence in S (which is called the phrase's source), $j-i$, and $S[j]$. The LZ77 encoding of S takes $\mathcal{O}(z \log n)$ bits, where z is the number of phrases in the parse. For example, in the following verses vertical lines indicate phrase boundaries:

919-lb|olt|tle|sl-olfl-belerl-onl-tlhle-lw|al|lll-919-bottles-of-beer-
If-olnle-olf-tlholsel-bottles-slhou|dl-hlap|peln-tol-flall-
981-bottles-of-beer-on-the-wall-

981-bottles-of-beer-on-the-wall-98-bottles-of-beer-
If-one-of-those-bottles-should-happen-to-fall-
971-bottles-of-beer-on-the-wall . . .

(We have displayed the verses with linebreaks to increase readability, but we have not considered them while computing the parse.) Although these verses may be annoyingly similar by the standards of natural language, they are far less similar than human genomes. Indeed, most repetitive biological datasets are much too similar (as well as much too large) for us to use them as informative examples.

One drawback of LZ77 compression is that reading a character in a compressed string can be very slow. Rytter (2003) and Charikar et al. (2005) showed how we can turn that parse into a balanced straight-line program (SLP) for S with $\mathcal{O}(z \log n)$ rules. An SLP for S is a context-free grammar in Chomsky normal form that generates S and only S ; it is balanced if the height of each subtree in the parse tree is logarithmic in that subtree's size. It follows from Rytter's and Charikar et al.'s results that we can store S in $\mathcal{O}(z \log^2 n)$ bits and support random-access reading of any substring of S with length l in $\mathcal{O}(\log n + l)$ time. Verbin and Yu (2013) showed that this is nearly optimal in the worst case. Bille et al. (2011) showed how, given even an unbalanced SLP for S with r rules, we can store S in $\mathcal{O}(r \log n)$ bits and support random-access reading in $\mathcal{O}(\log n + l)$ time. Rytter's, Charikar et al.'s, and Bille et al.'s constructions are not practical, but there are practical grammar-based compressors, such as those by Larsson and Mofat (1999) and Maruyama and Tabei (2014). As far as we know, block graphs by Gagie et al. (2011) and Gagie et al. (2014c) are the most practical grammar-like representations for random-access reading. The LZ78 algorithm by Ziv and Lempel (1978) does not compress repetitive datasets as well as LZ77, but the LZ78 encoding of S can easily be augmented to support random-access reading in $\mathcal{O}(\log \log n + l)$ time. LZ78 also works by parsing S into phrases but then each phrase must extend a previous phrase plus one character. Because of this property, the LZ78 encoding of S has $\Omega(\sqrt{n})$ phrases, even when $S = a^n$.

In the example above, the first verse contains many phrase boundaries but the second verse contains only three. Kuruppu et al. (2010) proposed that, given a set of similar strings (or one string that can easily be divided into similar substrings), we store the first string in plain text as a reference and compress the others with a version of LZ77 that restricts phrases' sources to occur in the reference. They called this scheme Relative Lempel–Ziv (RLZ) and showed it compresses genomic databases very well in practice (although it too uses $\Omega(\sqrt{n})$ phrases, even when $S = a^n$) and there are several implementations of this approach, such as those by Deorowicz and Grabowski (2011), Kuruppu et al. (2012), and Ferrada et al. (2014a). Even when there is no obvious reference, Kuruppu et al. (2011) showed we can often build one by sampling the dataset: intuitively, if a substring is common then it is likely to appear in our sample, and if it is not then we lose little by not compressing it well; this can be formalized using results about SLPs.

3. SEARCHING

Farach and Thorup (1998) observed that the first occurrence of any pattern $P[1, \dots, m]$ in S must cross or end at a phrase boundary in the LZ77 parse. Kärkkäinen and Ukkonen (1996) showed how, if we already know the locations of P 's occurrences in S that cross or end at phrase boundaries, then we can deduce the locations of all its other occurrences from the structure of the parse. By the same arguments, LZ78 also has these properties and (Karpinski et al., 1997) simultaneously proved similar results for SLPs. Bille et al. (2009) observed that any substring of S within edit distance k of P (i.e., any of P 's approximate matches) has length at most $m + k$, and any such substring that does not cross or end at an LZ78 phrase boundary must be an exact copy of an earlier one that does. They gave an algorithm for approximate pattern matching in LZ78 strings that works by extracting the $m + k$ and $m + k - 1$ characters before and after each LZ78 phrase boundary, respectively, using a technique similar to those discussed in Section 2; scanning the resulting substrings with any online algorithm for approximate pattern matching in uncompressed strings; and then deducing the locations of the other approximate matches from the structure of the parse.

Bille et al. (2011) extended this approach to show how we can find all P 's approximate matches in S from an SLP for S . Recently, Gagie et al. (2014b) extended it further to show how we can preprocess the LZ77 parse of S in $\mathcal{O}(z \log n)$ time such that later, given P and k , we can find all P 's *occ* approximate matches in $\mathcal{O}(z \min(mk, m + k^4) + occ)$ time. Their algorithm works by extracting the $m + k$ and $m + k - 1$ characters before and after each LZ77 phrase boundary, respectively, and then continuing as with the algorithm by Bille et al. (2009). The set of substrings we extract is like a kernel in parameterized complexity: the total length of the substrings can be much smaller than n , but a solution on them can quickly be mapped to a solution on all of S . For our example from Section 2 with $m = 4$ and $k = 1$, the kernel is:

99-bottles-of-beer-on-the-wall-99-bo
eer-If-one-of-those-bottles-should-happen-to-fall-98-bot
ll-98-bot
eer-If-on
ll-97-bot . . .

If we want a kernel consisting of only a single string, we can concatenate the substrings with $k + 1$ copies of \$ between each consecutive pair. Notice that if we are careful, we can avoid scanning the fourth substring "eer-If-on," since it occurs in the second substring.

We do not wish to leave the impression that kernelization is the only approach used in compressed pattern matching, nor even that the papers mentioned above are the only ones that use it. We have focused on those papers because we feel they are the most relevant to the practical bioinformatics papers by Wandelt and Leser (2012) and Rahn et al. (2014) mentioned in Section 1. Those authors were apparently unaware of the field of compressed pattern matching and re-invented kernelization specifically for approximate pattern matching in genomic databases, with kernels based on RLZ instead of LZ77, LZ78, or SLPs. This may be because the earlier researchers using kernelization for pattern matching did not perform convincing experiments on large enough datasets,

publicize their ideas in interdisciplinary forums or implement their ideas in tools usable by other scientists.

4. INDEXING

Kärkkäinen and Ukkonen (1996) gave the first LZ-based index, which supported exact pattern matching and stored S separately and uncompressed. They used Patricia trees and range reporting to find a set of candidate matches crossing or ending at LZ77 phrase boundaries; verified them by checking S ; and then used more range reporting to find the other matches. We can obtain various time-space tradeoffs by compressing S and use the methods discussed in Section 1 to extract the characters needed to verify candidate matches. Claude and Navarro (2012) modified Kärkkäinen and Ukkonen's index to use a grammar-compressed encoding of S , and Kreft and Navarro (2013) modified it to use the encoding of S produced by a version of LZ77 they called LZ-End, which supports fast random-access reads starting at phrase boundaries. Arroyuelo et al. (2012) and Do et al. (2014) gave indexes based on LZ78 and RLZ, respectively, and Maruyama et al. (2013) and Takabatake et al. (2014) gave indexes based on the edit-sensitive parsing by Cormode and Muthukrishnan (2007). Gagie et al. (2014a) recently gave a version of Kärkkäinen and Ukkonen's index that uses a total of $\mathcal{O}(z \log^2 n)$ bits and returns the locations of all P 's occurrences in S in $\mathcal{O}(m \log m + occ \log \log n)$ time. These indexes require no assumptions about the pattern.

Kärkkäinen and Sutinen (1998) gave an index based on a version of LZ77 that allows phrases to overlap by $q - 1$ characters, where q is a parameter. If P has length exactly q , then their index returns the locations of all P 's occurrences in S in optimal $\mathcal{O}(m + occ)$ time. If we are given an upper bound M on the pattern length at construction time, then even with Kärkkäinen and Ukkonen's original version, we need keep only a kernel of the text and can use $\mathcal{O}(z \log n + zM \log \sigma)$ bits in total, where σ is the size of the alphabet. We suspect this escaped investigation for so long because it seemed too obvious and inelegant to be theoretically interesting, and the need to index massive, highly repetitive datasets in practice has become pressing only since the development of next-generation sequencing methods.

The use of kernelization for indexing was eventually investigated by Schneeberger et al. (2009), although they did not present kernelization as a separate process because their work was application-driven. As noted in Section 1, they proposed that, given a database of genomes from the same species, we index the common parts of the genomes only once for them all, but we index the parts near variation sites for each genome. Wandelt et al. (2013) and Danek et al. (2014) gave similar results, essentially using a kernel based on the RLZ parse. Like Schneeberger et al., these authors indexed the kernels using specific methods based on q -grams or seeds. Danek et al.'s index for the database for the 1000 Genomes Project is the first one to fit in the memory of a commodity personal computer. Ferrada et al. (2014b) emphasized kernelization (albeit not under that name) in terms of the LZ77 parse, which is more general and may give better compression, and pointed out that we can use any index for approximate pattern matching to store the kernel. One point they did not comment on, and which we hope to have clarified in this paper, is that we

can consider kernels based on LZ77, LZ78, RLZ, other compression schemes, or possibly other algorithms entirely. These kernels may be easier to compute when the database does not fit in memory, or have other useful properties that make them preferable in some situations. One interesting problem is how we can best maintain a dynamic kernel for an expanding database. This could allow us to align reads against a genomic database and then add the newly assembled genome, which could be useful when dealing with mutating cancer genomes or changing strains of a disease during an outbreak.

ACKNOWLEDGMENTS

Many thanks to Fabio Cunial, Paweł Gawrychowski, Szymon Grabowski, Juha Kärkkäinen, Veli Mäkinen, Gonzalo Navarro, Esa Pitkänen, Yasuo Tabei, and Niko Välimäki, for helpful discussions, and to the anonymous reviewers for their comments. *Funding:* The authors are funded by Academy of Finland grants 268324, 258308 and 250345 (CoECGR).

REFERENCES

- Arroyuelo, D., Navarro, G., and Sadakane, K. (2012). Stronger Lempel-Ziv based compressed text indexing. *Algorithmica* 62, 54–101. doi:10.1007/s00453-010-9443-8
- Bille, P., Fagerberg, R., and Gørtz, I. L. (2009). Improved approximate string matching and regular expression matching on Ziv-Lempel compressed texts. *ACM Trans. Algorithms* 6:3. doi:10.1145/1644015.1644018
- Bille, P., Landau, G. M., Raman, R., Sadakane, K., Satti, S. R., and Weimann, O. (2011). "Random access to grammar-compressed strings," in *Proceedings of the 22nd Symposium on Discrete Algorithms (SODA)* (Philadelphia: Society for Industrial and Applied Mathematics (SIAM)), 373–389.
- Charikar, M., Lehman, E., Liu, D., Panigrahy, R., Prabhakaran, M., Sahai, A., et al. (2005). The smallest grammar problem. *IEEE Trans. Inf. Theory* 51, 2554–2576. doi:10.1109/TIT.2005.850116
- Claude, F., and Navarro, G. (2012). "Improved grammar-based compressed indexes," in *Proceedings of the 19th Symposium on String Processing and Information Retrieval (SPIRE)* (Berlin: Springer-Verlag), 180–192.
- Cormode, G., and Muthukrishnan, S. (2007). The string edit distance matching problem with moves. *ACM Trans. Algorithms* 3:2. doi:10.1145/1186810.1186812
- Danek, D. A., Deorowicz, S., and Grabowski, S. (2014). Indexes of large genome collections on a PC. *PLoS ONE* 9:e109384. doi:10.1371/journal.pone.0109384
- Deorowicz, S., and Grabowski, S. (2011). Robust relative compression of genomes with random access. *Bioinformatics* 27, 2979–2986. doi:10.1093/bioinformatics/btr505
- Do, H. H., Jansson, J., Sadakane, K., and Sung, W. K. (2014). Fast relative Lempel-Ziv self-index for similar sequences. *Theor. Comp. Sci.* 532, 14–30. doi:10.1016/j.tcs.2013.07.024
- Farach, M., and Thorup, M. (1998). String matching in Lempel-Ziv compressed strings. *Algorithmica* 20, 388–404. doi:10.1007/PL00009202
- Ferrada, H., Gagie, T., Gog, S., and Puglisi, S. J. (2014a). "Relative Lempel-Ziv with constant-time random access," in *Proceedings of the 21st Symposium on String Processing and Information Retrieval (SPIRE)* (Berlin: Springer-Verlag), 13–17.
- Ferrada, H., Gagie, T., Hirvola, T., and Puglisi, S. J. (2014b). Hybrid indexes for repetitive datasets. *Philos. Trans. R. Soc. A* 327, 2016. doi:10.1098/rsta.2013.0137
- Gagie, T., Gawrychowski, P., Kärkkäinen, J., Nekrich, Y., and Puglisi, S. J. (2014a). "LZ77-based self-indexing with faster pattern matching," in *Proceedings of the 11th Latin American Symposium on Theoretical Informatics (LATIN)* (Berlin: Springer-Verlag), 731–742.
- Gagie, T., Gawrychowski, P., and Puglisi, S. J. (2014b). Faster approximate pattern matching in compressed repetitive texts. *J. Discrete Algorithms*. doi:10.1016/j.jda.2014.10.003
- Gagie, T., Hoobin, C., and Puglisi, S. J. (2014c). "Block graphs in practice," in *Proceedings of the 2nd International Conference on Algorithms for Big Data (ICABD)* (Aachen: CEUR-WS.org), 30–36.
- Gagie, T., Gawrychowski, P., and Puglisi, S. J. (2011). "Faster approximate pattern matching in compressed repetitive texts," in *Proceedings of the 22nd International*

- Symposium on Algorithms and Computation (ISAAC)* (Berlin: Springer-Verlag), 653–662.
- Kärkkäinen, J., and Sutinen, E. (1998). Lempel-Ziv index for q -grams. *Algorithmica* 21, 137–154. doi:10.1007/PL00009205
- Kärkkäinen, J., and Ukkonen, E. (1996). “Lempel-Ziv parsing and sublinear-size index structures for string matching,” in *Proceedings of the 3rd South American Workshop on String Processing (WSP)* (Ottawa: Carleton University Press), 141–155.
- Karpinski, M., Rytter, W., and Shinohara, A. (1997). An efficient pattern-matching algorithm for strings with short descriptions. *Nordic J. Comput.* 4, 172–186.
- Kreft, S., and Navarro, G. (2013). On compressing and indexing repetitive sequences. *Theor. Comp. Sci.* 483, 115–133. doi:10.1016/j.tcs.2012.02.006
- Kuruppu, S., Beresford-Smith, B., Conway, T. C., and Zobel, J. (2012). Iterative dictionary construction for compression of large DNA data sets. *IEEE/ACM Trans. Comput. Biol. Bioinform.* 9, 137–149. doi:10.1109/TCBB.2011.82
- Kuruppu, S., Puglisi, S. J., and Zobel, J. (2010). “Relative Lempel-Ziv compression of genomes for large-scale storage and retrieval,” in *Proceedings of the 17th Symposium on String Processing and Information Retrieval (SPIRE)* (Berlin: Springer-Verlag), 201–206.
- Kuruppu, S., Puglisi, S. J., and Zobel, J. (2011). “Reference sequence construction for relative compression of genomes,” in *Proceedings of the 18th Symposium on String Processing and Information Retrieval (SPIRE)* (Berlin: Springer-Verlag), 420–425.
- Larsson, N. J., and Moffat, A. (1999). “Offline dictionary-based compression,” in *Proceedings of the Data Compression Conference (DCC)* (Hoboken, NJ: IEEE Press), 296–305.
- Maruyama, S., Nakahara, M., Kishiue, N., and Sakamoto, H. (2013). ESP-index: a compressed index based on edit-sensitive parsing. *J. Discrete Algorithms* 18, 100–112. doi:10.1016/j.jda.2012.07.009
- Maruyama, S., and Tabei, Y. (2014). “Fully online grammar compression in constant space,” in *Proceedings of the Data Compression Conference (DCC)* (Hoboken, NJ: IEEE Press), 173–182.
- Rahn, R., Weese, D., and Reinert, K. (2014). Journalized string tree – a scalable data structure for analyzing thousands of similar genomes on your laptop. *Bioinformatics* 30, 3499–3505. doi:10.1093/bioinformatics/btu438
- Rytter, W. (2003). Application of Lempel-Ziv factorization to the approximation of grammar-based compression. *Theor. Comp. Sci.* 302, 211–222. doi:10.1016/S0304-3975(02)00777-6
- Schneeberger, K., Hagmann, J., Ossowski, S., Warthmann, N., Gesing, S., Kohlbacher, O. D., et al. (2009). Simultaneous alignment of short reads against multiple genomes. *Genome Biol.* 10, R98. doi:10.1186/gb-2009-10-9-r98
- Takabatake, Y., Tabei, Y., and Sakamoto, H. (2014). “Improved ESP-index: a practical self-index for highly repetitive texts,” in *Proceedings of the 13th Symposium on Experimental Algorithms (SEA)* (Berlin: Springer-Verlag), 338–350.
- Verbin, E., and Yu, W. (2013). “Data structure lower bounds on random access to grammar-compressed strings,” in *Proceedings of the 24th Symposium on Combinatorial Pattern Matching (CPM)* (Berlin: Springer-Verlag), 247–258.
- Vyverman, M., Baets, B. D., Fack, V., and Dawyndt, P. (2012). Prospects and limitations of full-text index structures in genome analysis. *Nucleic Acids Res.* 40, 6993–7015. doi:10.1093/nar/gks408
- Wandelt, S., and Leser, U. (2012). “String searching in referentially compressed genomes,” in *Proceedings of the Conference on Knowledge Discovery and Information Retrieval (KDIR)* (SciTePress), 95–102.
- Wandelt, S., Starlinger, J., Bux, M., and Leser, U. (2013). “RCSI: scalable similarity search in thousand(s) of genomes,” in *Proceedings of the VLDB Endowment*, Vol. 6 (San Jose, CA: VLDB Endowment), 1534–1545. doi:10.14778/2536258.2536265
- Ziv, J., and Lempel, A. (1977). A universal algorithm for sequential data compression. *IEEE Trans. Inf. Theory* 23, 337–343. doi:10.1109/83.663496
- Ziv, J., and Lempel, A. (1978). Compression of individual sequences via variable-rate coding. *IEEE Trans. Inf. Theory* 24, 530–536. doi:10.1109/TIT.1978.1055911

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Received: 11 December 2014; accepted: 22 January 2015; published online: 09 February 2015.

Citation: Gagie T and Puglisi SJ (2015) Searching and indexing genomic databases via kernelization. *Front. Bioeng. Biotechnol.* 3:12. doi: 10.3389/fbioe.2015.00012

This article was submitted to *Bioinformatics and Computational Biology*, a section of the journal *Frontiers in Bioengineering and Biotechnology*.

Copyright © 2015 Gagie and Puglisi. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) or licensor are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.