

Using Sequence Ensembles for Seeding Alignments of MinION Sequencing Data

Rastislav Rabatin, Broňa Brejová, and Tomáš Vinař

Faculty of Mathematics, Physics and Informatics, Comenius University,
Mlynská dolina, 842 48 Bratislava, Slovakia

Abstract: Sequence similarity search is in bioinformatics often solved by seed-and-extend heuristics: we first locate short exact matches (hits) by hashing or other efficient indexing techniques and then extend these hits to longer sequence alignments. Such approaches are effective at finding very similar sequences, but they quickly lose sensitivity when trying to locate weaker similarities.

In this paper, we develop seeding strategies for data from MinION DNA sequencer. This recent technology produces sequencing reads which are prone to high error rates of up to 30%. Since most of these errors are insertions or deletions, it is difficult to adapt seed-and-extend algorithms to this type of data. We propose to represent each read by an ensemble of sequences sampled from a probabilistic model, instead of a single sequence. Using this extended representation, we were able to design a seeding strategy with 99.9% sensitivity and very low false positive rate. Our technique can be used to locate the part of the genome corresponding to a particular read, or even to find overlaps between pairs of reads.

1 Introduction and Background

Sequence alignment is one of the basic problems in bioinformatics. The goal is to find regions of sequence similarity between two sequences (usually DNA sequences or proteins) or between a query sequence and a large sequence database. Although the highest scoring alignment can be found by a simple dynamic programming algorithm (Smith and Waterman, 1981), its quadratic running time is too slow for large databases. Most sequence alignment tools therefore follow the seed-and-extend paradigm, well-known from BLAST (Altschul et al., 1990). The first step is to build an index of one of the sequences, and then use it to efficiently identify short substrings that occur in both sequences. Each such short exact match, called *hit*, forms a *seed* of a potential alignment. In the extension phase, sequences surrounding each hit are inspected, and the hit is extended to a full alignment, or discarded, if no high scoring alignment can be found. The *sensitivity* of such an approach depends largely on how likely it is for a real alignment to contain a hit. Without the hit, the extension phase is never triggered, and the alignment cannot be identified. On the other hand, the *running time* depends mainly on how many false hits will trigger unnecessary extensions.

The original BLAST (Altschul et al., 1990) used 11 consecutive matches as a seed. Consecutive matches of a fixed

length are very easy to index by standard hashing techniques. Several other more complex seeding strategies are used in modern alignment tools (Brown, 2008).

In this work, we consider read mapping, which is a specific alignment task, where the query sequence is a read produced by a DNA sequencing machine, and the database consists of the known reference genome of the same or very closely related organism to the one we are sequencing. The goal is to find the region of the reference genome from which a particular sequencing read was obtained. This task is relatively simple for DNA sequencing technologies with low error rates, because we are looking for almost exact matches between the reads and the reference genome; the goal is to map a large number of reads as quickly as possible. However, in this paper we consider reads produced by MinION sequencer (Ip et al., 2015), which can have up to 30% error rate, and thus pose a challenge for read mappers. The main advantages of the MinION sequencing device are its small size and portability, and its ability to sequence very long reads (reads as long as 100 kbp were reported).

In this paper, we explore the use of simple seeding strategies for mapping MinION reads to the reference sequence. Instead of a single read sequence as a query, we propose to use an ensemble of sequences sampled from a hidden Markov model used for base calling. With such an ensemble of sequences representing alternative predictions of the true read sequence, we were able to achieve very high sensitivity and a very small number of false hits using a relatively simple alignment seeding strategy.

To sequence DNA, MinION uses measurements of electric current as a single-stranded fragment of DNA passes through a nanopore. The electric current depends mostly on the context of k bases of DNA passing through the pore at the time of the measurement. As the DNA fragment moves through the pore, this context changes and measurements change accordingly.

The raw measurements are processed by MinKnow software from Oxford Nanopore. MinKnow first splits raw measurements into *events*, where each event would ideally correspond to a single-base shift of the DNA through the pore. Each event is characterized by the mean and the variance of the corresponding raw measurements. This sequence of events is then uploaded to a cloud-based service Metrichor for base calling.

The exact details of the algorithms behind MinKnow and Metrichor are not disclosed by Oxford Nanopore.

However, the whole process is naturally modelled by a hidden Markov model (HMM) (Durbin et al., 1998, Chapter 3) with hidden states corresponding to k -mers of the underlying DNA sequence and observations representing the mean value of the current in each event. In fact, data provided by Oxford Nanopore include parameters of such a model. Open-source base caller Nanocall based on this idea was recently implemented by David et al. (2016), with accuracy similar to Metrichor. To decode a sequence of observations, Nanocall uses the standard Viterbi algorithm (Viterbi, 1967) for finding the most probable state path. Another open-source base caller DeepNano is based on recurrent neural networks (Boža et al., 2016).

Base calls produced by the Viterbi algorithm contain many errors (David et al., 2016; Boža et al., 2016); a typical error rate would be around 30%, dominated mostly by insertions and deletions. Currently, two general-purpose aligners are used in the community to map MinION reads: BWA-MEM (Li and Durbin, 2010), and LAST (Kielbasa et al., 2011). Both of these tools follow the seed-and-extend paradigm, but they use variable seed lengths, indexed with FM-index (Ferragina and Manzini, 2000) or suffix arrays (Manber and Myers, 1993). By extending seeds to the point of only a few occurrences, one can avoid most costly false positives. The adjustments for MinION reads in case of BWA-MEM and LAST include lowering the minimum length of a seed to be considered as a valid hit, and changes that make the extension phase less stringent. GraphMap tool (Sovic et al., 2016), specifically targeting MinION data, uses seeds allowing insertions and deletions in the context of a complex multi-step algorithm that goes well beyond a simple seed-and-extend framework.

All of these tools consider Metrichor base calls, equivalent to the most probable state path in the HMM, as the query sequence. Our approach to the challenges imposed by MinION reads is to use an ensemble of sub-optimal sequences instead of a single DNA sequence. To this end, we have implemented a sampling algorithm that can generate samples from the posterior distribution of state paths in an HMM given the sequence of observations (see, e.g. Cawley and Pachter (2003)). We adapt common seeding strategies to such ensembles of sequences and show that on real data we can find a seed that is easy to index, 99.9% sensitive, and yields only a small number of false positives that would trigger extension phase unnecessarily.

Sampling of suboptimal MinION base calls was also considered by Szalay and Golovchenko (2015). They use sampling from the base calling HMM to arrive at the correct consensus sequence for an alignment of multiple reads. Due to the nature of errors in MinION reads and availability of a reasonable probabilistic model, we consider sampling strategies to be a promising alternative in many applications of MinION.

2 HMM for Sampling MinION Base Calls

Both Oxford Nanopore Metrichor base caller and recently released open-source base caller Nanocall (David et al., 2016) use hidden Markov models. Briefly, each hidden state of the HMM represents one k -mer passing through the pore, and the emission of the state is the value of the electric current. Actual measurements of the current provided by the device with high sampling rate are segmented by the MinKnow software into discrete *events*, each corresponding to the shift of the DNA sequence through the pore by a single base. The base callers then use the HMM to obtain the sequence of hidden states given the sequence of events from the MinION read.

Definition of the model. Our HMM follows the same general idea. The set of states of our HMM is composed of all 4^k possible k -mers (we denote state for a k -mer x by S_x) and the starting state S_0 . Different versions of MinION use different values of k ; in our experiments we have used a data set with $k = 5$, while the newer chemistry uses $k = 6$.

Emission of state S_x is represented as a continuous random variable. The probability of observing a measurement e for a k -mer x is given as

$$\Pr(e|x) \sim \mathcal{N}(\text{scale} \cdot \mu_x + \text{shift}, \sigma_x \cdot \text{var}), \quad (1)$$

where $\mathcal{N}(\mu, \sigma)$ is the normal distribution with mean μ and standard deviation σ . Parameters μ_x , σ_x (specific for each version of the chemistry and each k -mer x), and *scale*, *shift*, *var* (scaling parameters specific for each read) are provided by Oxford Nanopore and can be obtained from the FAST5 file containing each read. Starting state S_0 is silent.

Under ideal conditions, each event corresponds to a shift by a single base in the DNA sequence. This corresponds to four outgoing transitions from each state S_x to state S_y , where x and y overlap by exactly $k - 1$ bases (i.e., S_{ACTG} has transitions to states S_{ACTGA} , S_{ACTGC} , S_{ACTGG} , and S_{ACTGT}). This organization closely resembles de Bruijn graphs commonly used in sequence assembly (Pevzner et al., 2001). All four transitions have an equal probability. From the starting state S_0 , we have a transition to each possible S_x with equal probability $1/4^k$.

Segmentation of raw measurements into events is known to be error prone. In particular, two events with similar measurements can be fused together, or a single event can be split artificially into multiple events. Thus the assumption that each event corresponds to a single-nucleotide shift is unrealistic. To account for this fact, we have introduced additional transitions to our model.

First, we have added a self-transition (so called *split transition*) to each state, which models splitting of a single true event into multiple predicted events. Second, we have also added so called *skip transitions* between all pairs of states S_x and S_y , which correspond to shifts of the k -mer by up to k bases instead of one.

The transition probabilities for split and skip transitions are not provided by MinION. We have estimated these parameters directly from the data, as outlined in Section 3. Alternatively, we could employ a more elaborate error model for Oxford Nanopore event segmentation process.

Inference in the model. A traditional way of decoding HMMs is by finding the most probable sequence of states by the Viterbi algorithm (Viterbi, 1967), which is the approach taken both by Metrichor and Nanocall. The resulting sequence of states (which is, in fact, a sequence of k -mers corresponding to individual events) can be translated into the DNA sequence. In most cases, the neighbouring k -mers in the sequence should be shifted by one, and thus each state in the sequence introduces one additional base of the DNA sequence. However, split and skip transitions may introduce between 0 and k bases for each event. In these cases, the result is not necessarily unique: for example state sequence $S_{ACTTCTC}S_{CTCTCA}$ could correspond to one of the sequences ACTCTCA, ACTCTCTCA, ACTCTCTCTCA, or even ACTCTCCTCTCA. We have decided to adopt the shortest possible interpretation, as is done in Nanocall.

Since the base calls produced by the Viterbi algorithm contain many errors, we have decided to explore the use of an ensemble of alternative sequences instead of a single base call sequence. To this end, we have implemented the stochastic traceback algorithm (see, e.g. Cawley and Pachter (2003)) to generate samples from the posterior distribution of state paths given the sequence of observations.

Briefly, for a given sequence of observations $e_1e_2\dots e_n$, the algorithm starts by computing *forward probabilities*, where $F(i, s)$ is the probability of generating first i observations $e_1e_2\dots e_i$, and ending in state s (Rabiner, 1989). The last state s_n of the path is sampled proportionally to the probabilities $F(n, s_n)$. When state s_n is fixed, we can sample state s_{n-1} proportionally to $F(n-1, s_{n-1}) \cdot t_{s_{n-1}, s_n}$, where $t_{s, s'}$ is the transition probability from state s to state s' . This can be continued, until we sample the complete path $s_1\dots s_n$. The running time of the algorithm is $O(nm^2)$, where n is the length of the sequence and m is the number of states. Forward probabilities need to be computed only once if multiple samples are required for the same read.

Figure 1 illustrates typical differences between individual samples. Note that the samples are almost identical in some regions, but these high-confidence regions are interspersed by regions with very high uncertainty. This is a typical pattern for MinION data.

3 Experiments

Our goal is to consider various seeding strategies for seed-and-extend algorithms. As discussed in Section 1, a typical seed-and-extend algorithm first uses an index structure to locate hits between the query sequence and the target. For example, the most basic BLAST strategy looks for exact matches of length 11 (Altschul et al., 1990). It then

tries to extend each cluster of hits to a full alignment. The extension phase usually involves dynamic programming and is therefore time consuming.

The seed-and-extend algorithms cannot locate alignments that do not contain a hit of the seed between the query and the target sequence. We call these alignments *false negatives*. Note that even a single hit is often sufficient to locate the whole alignment, depending on the particular extension strategy. On the other hand, spurious hits between random locations can also trigger extension. We call such spurious hits *false positives*. Unsuccessful extension of spurious hits often dominates the running time of the alignment algorithm, and consequently, we need to minimize the number of false positives.

In this section, we concentrate on the evaluation of a trade-off between false positives and false negatives for a variety of seeding strategies. We start from simple strategies, e.g. k consecutive exact matches between the reference sequence and the Viterbi base call, and we also propose several novel seeding strategies using the sampling framework and multiple hit requirement that greatly improve the trade-off.

Data preprocessing and model training. To evaluate seeding strategies, we use a data set composed of real sequencing reads from the model bacteria *E. coli* (strain MG1655) with accession number ERR968968 produced by the Cold Spring Harbor Laboratory by using MinION sequencer with SQK-MAP005 kit. For simplicity, we have only considered template reads (complement reads from the reverse strand use different model parameters).

Some of the MinION reads are of very low quality and consequently it is impossible to map them to the reference sequence even with the most sensitive tools available. To filter out these low quality reads, we have mapped Metrichor base calls to the reference genome by BWA-MEM (Li and Durbin, 2010) with `-x ont2d` parameters optimized for mapping Oxford Nanopore reads. The reads that did not map to the reference at all were discarded. We also discarded reads where Metrichor predicted skips in the event sequence longer than two. From the original 27,073 reads, we were left with 25,162 reads.

From these reads, we have randomly selected a training set (693 reads) and a testing set (307 reads). The training set was used to estimate the transition probabilities in our HMM. In particular, we set the probability of each transition to be proportional to the number of times the transition was observed in the training data set. We added pseudocount of 1 to avoid zero transition probabilities for rare transitions.

Preparing testing data. For each sequence in the testing set, we have produced a Viterbi base call and 250 samples from the posterior distribution as outlined in Section 2. Figure 2 shows comparison of sequence identities of individual base calls to the reference genome. Note that our Viterbi base calls are not too different from Metrichor base

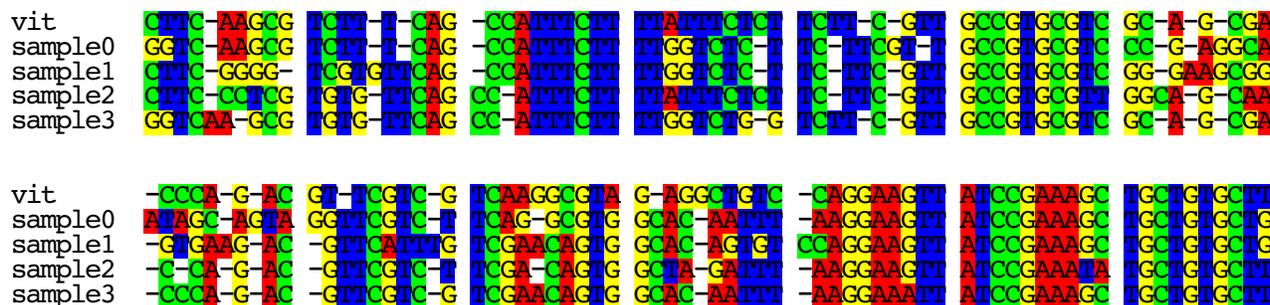


Figure 1: Example of base calling samples from a MinION read. The first line corresponds to the Viterbi base call, other lines correspond to four samples from the posterior distribution defined by the HMM. Base calls were aligned according to events in the sequence of observations.

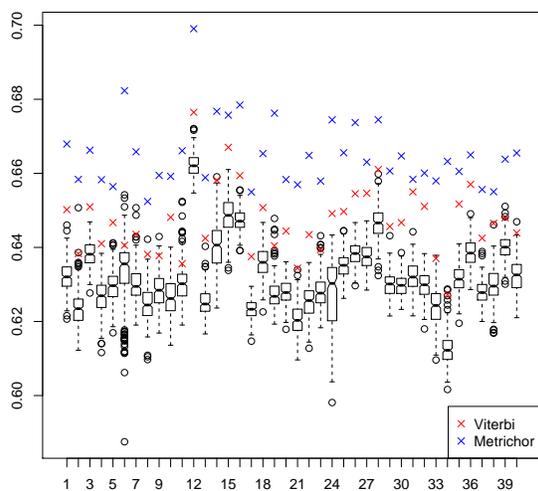


Figure 2: Comparison of sequence identity to the reference genome for Metrichor base calls, our Viterbi base calls, and posterior samples for 40 randomly selected reads. All sequences were aligned to the reference by BWA-MEM. Sequence identity of Metrichor, Viterbi, and a box plot distribution of sequence identities of 250 samples are shown on the y-axis.

calls; slight decrease in the quality of calls is expected due to simplicity of the model we have used (the decrease in the sequence identity is similar to that observed by David et al. (2016)). Sampled sequences have in general lower sequence identity than the Viterbi base call, as can be expected, since they correspond to sub-optimal paths through the model. However, the difference from the Viterbi base call identities is not very high.

Experimental setup and evaluation. To evaluate application of various seeding strategies in MinION data, we have split the reads into *windows*, each corresponding to 500 events. These windows represent our query sequences. For each window we have an alignment of all samples and

the Viterbi base call. This alignment was based on event boundaries: we have padded all sequences generated for one event in individual samples by gap symbols so that they have the same length (see also Figure 1).

The Viterbi base call sequence from each window was aligned to the reference genome by LAST software (Kielbasa et al., 2011) with parameters $-q 1 -a 1 -b 1 -T 1$. We have kept only those windows that aligned to a unique place in the genome, and the alignment covered the entire length of the window. After this step, we were left with 3192 windows out of 4948. A randomly chosen subset of 143 windows was used as a validation set for exploring various seeding strategies, and the remaining 3049 windows were used for final testing. The region of the reference covered by the alignment to the window is considered to be the only true alignment of the window to the reference sequence.

When testing a seeding strategy, we locate hits of a particular seed in both the query window and the reference sequence. We represent each hit by the coordinates of the left endpoints of the hit in the query window and in the reference. The seed hit is considered to be *valid*, if the endpoint in the reference is within the region covered by the alignment of this particular window and on the correct strand of DNA. The entire window of a read is considered to be a *true positive (TP)*, if it contains at least one valid hit; we assume that the extension algorithm would be able to recover the alignment within this window starting from this hit. The *sensitivity (S_n)* of the seeding strategy is the number of true positives divided by the total number of windows.

Many seed hits are invalid, and they contribute to the false positive rate. Often we see clusters of hits with very similar coordinates in both reference and the query window. Presumably the extension algorithm would be called only once for each such cluster. Therefore, we compute the number of *false positives* by greedily selecting one hit from each cluster so that each hit in the cluster differs in both coordinates by at most 10 from the selected hit.

Simple seeding strategies. The simplest seeding strategy is to consider k consecutive exact matches as a hit. The traditional approach would create an index of all k -mers in the reference genome and then scan all k -mers in the query windows. Each cluster of matching k -mers would trigger the extension phase.

For example, if we consider the Viterbi base calls and use 13 exact consecutive matches as a seed, we will be able to map 98.8% windows to the correct region of the reference (see Fig. 3), but we will also incur a substantial number of false positives (more than 240,000 or about 80 per each query window).

Our strategy of using sampling instead of Viterbi base calls works as follows. In the basic version of our approach (see $t = 1$ in Fig. 3), we consider k -mers from n different samples from the given window. Each k -mer can potentially form the seed triggering the extension phase of the alignment. To match the sensitivity of the Viterbi base calls for $k = 13$, we need only $n = 3$ samples. One advantage of our approach is that we can increase the sensitivity by increasing the number of samples n . For example with 8 samples we reach 99.9% sensitivity and with 14 samples 100% sensitivity. The cost for this very high sensitivity is a high false positive rate; even with 3 samples we have about $2.6\times$ more false positives than the Viterbi.

To improve the false positive rate, we use a simple pre-filtering step: at each position in the window we consider only those k -mers that appear in at least t different samples. Assuming that the true sequence has a high posterior probability in the model, we expect that it will occur in many samples, whereas most other variants would occur rarely and thus be filtered out. Indeed, for $t = 4$ we can reach the sensitivity of the Viterbi algorithm with about 20% reduction in the false positive rate, using $n = 25$ samples. Performance for other values of n and t can be found in Figure 3.

Multiple seed hits to trigger extension. We have also considered a more complex seeding strategy, where we first find matching 10-mers, and then we join them into *chains* of length 3. This technique has previously proved to be very effective for regular alignment tasks (Altschul et al., 1997).

Matching 10-mers in the chain are required to have increasing coordinates in both the read and the reference sequence, and the distance between starts of adjacent matching k -mers in the chain must be at least 10 and at most 50 in both sequences. However, the distances of the two 10-mer matches in the two sequences may differ, accommodating insertions and deletions in the intervening regions. The entire chain is then again represented by its leftmost point in both the read and the reference for the purpose of determining if it is valid.

As we see in Figure 4, this seeding strategy is too stringent for the Viterbi base calls, achieving only 71.3% sensitivity. On the other hand, false positives are extremely rare, totaling only 136 in all 3049 windows.

When using samples, different 10-mer matches in the chain may come from different samples. The chaining of weaker hits helps to accommodate regions with high uncertainty present in the MinION data. Some settings of our strategy can achieve the same sensitivity as the Viterbi algorithm with even lower false positives, but more importantly, by considering more samples, we can increase sensitivity while keeping the false positives quite low (Figure 4). For example, for $t = 2$ and $n = 43$ our strategy can reach 99.9% sensitivity with only 6407 false positives.

4 Conclusions and Future Work

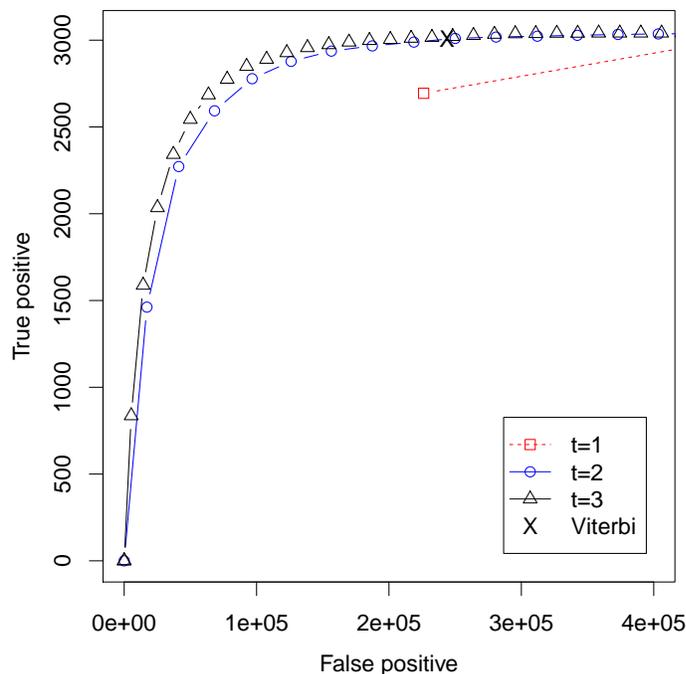
In this paper, we have examined the problem of mapping MinION reads to the reference sequence. The error rate of MinION is high, with many insertions and deletions. Consequently the standard alignment techniques do not achieve sufficient mapping sensitivity.

Instead of representing the read by a single base called sequence, we have proposed to use an ensemble of sequences generated from the posterior distribution defined by the HMM capturing the properties of the MinION sequencing. We have adapted the standard k -mer based techniques for alignment seeding to ensembles of sequences and identified a seed (three 10-mer hits spaced by at most 50 bases) that achieved 99.9% sensitivity with an extremely small number of false positives. With such a low false positive rate, we could investigate more precise (and slower) algorithms for the extension phase as the next step towards sensitive alignment of MinION reads.

An obvious extension of our approach would be to consider spaced seeds (Ma et al., 2002). Our experiments suggest that a typical MinION read consists of short regions of high-confidence sequence (often under 15 bases) interspersed with regions of high uncertainty with many insertions and deletions. The spaced seeds would have to target mainly the high-confidence regions; however, these regions seem to be too short to admit complex seeds of any significant weight. One possibility would be to build a probabilistic model capturing high-confidence and high-uncertainty regions and transitions between them, and attempt to design optimized spaced seeds, for example by techniques suggested by Brejova et al. (2004).

Another option would be to use seeds that also allow indels at do not care positions. These types of seeds were successfully used by Sovic et al. (2016) for MinION read mapping, but the overall algorithm was much more complicated than a simple seed-and-extend. Moreover, these types of seeds are much more difficult to index than continuous or spaced seeds and we believe, that our sampling approach together with multiple chained seed hits provides an elegant answer to the problem.

In this work, we have only considered template reads from MinION. However, MinION attempts to read both strands of the same DNA molecule and these two readouts can be combined in postprocessing to a single sequence



Method	Sn	FP
Viterbi	0.988	243796
$t = 1, n = 2$	0.974	433890
$t = 1, n = 3$	0.990	625782
$t = 1, n = 8$	0.999	1456049
$t = 2, n = 7$	0.963	156472
$t = 2, n = 11$	0.990	280912
$t = 2, n = 26$	0.999	727678
$t = 3, n = 12$	0.960	122830
$t = 3, n = 18$	0.988	216743
$t = 3, n = 38$	0.999	529692
$t = 4, n = 17$	0.957	110197
$t = 4, n = 25$	0.989	192648
$t = 4, n = 53$	0.999	488198
$t = 5, n = 22$	0.956	103666
$t = 5, n = 34$	0.988	196604
$t = 5, n = 72$	0.999	497944

Figure 3: **Performance of our approach compared to the Viterbi base calls for seeding with a single 13-mer.** The x -axis of the plot is the total number of false positives on the whole testing set; the y -axis is the number of true positives out of 3049 windows in total. Performance of the Viterbi base calls is shown by the black X . Three lines show our approach for different values of threshold t for filtering k -mers. Each point on the line represents performance for a particular number of samples $n = 1, 2, \dots$. The table shows sensitivity and the number of false positives for selected values of n and t . In particular for each t , we show the smallest value of n achieving sensitivity at least 95%, the sensitivity of the Viterbi algorithm, and sensitivity at least 99.9%.

called 2D read. Since 2D reads are much more accurate (typical error rate is about 15%), most researchers use only 2D reads in further analysis. Accurate 1D read mapping would allow researchers to consider all data; usually there is about $4 - 5 \times$ more 1D reads than 2D reads (Ip et al., 2015). Moreover, recent application of MinION to monitoring disease outbreaks (Quick et al., 2016) require that reads are analyzed on-the-fly as they are produced, and at this stage 2D reads are not available.

In this paper, we have evaluated seeding strategies for mapping reads to the reference genome. In future work, we would like to investigate the seed-and-extend approaches for read-to-read alignment. With our sampling approach, we would not need to commit to a single interpretation of either of the sequences, potentially increasing the sensitivity of detecting overlaps between reads in a given data set. Sensitive read-to-read alignment is essential for *de novo* genome assembly.

Acknowledgements. This research was funded by APVV grant APVV-14-0253 and VEGA grants 1/0719/14 (TV) and 1/0684/16 (BB).

References

Altschul, S. F., Gish, W., Miller, W., Myers, E. W., and Lipman, D. J. (1990). Basic local alignment search tool.

Journal of Molecular Biology, 215(3):403–410.

Altschul, S. F., Madden, T. L., Schaffer, A. A., Zhang, J., Zhang, Z., Miller, W., and Lipman, D. J. (1997). Gapped BLAST and PSI-BLAST: a new generation of protein database search programs. *Nucleic Acids Research*, 25(17):3389–3392.

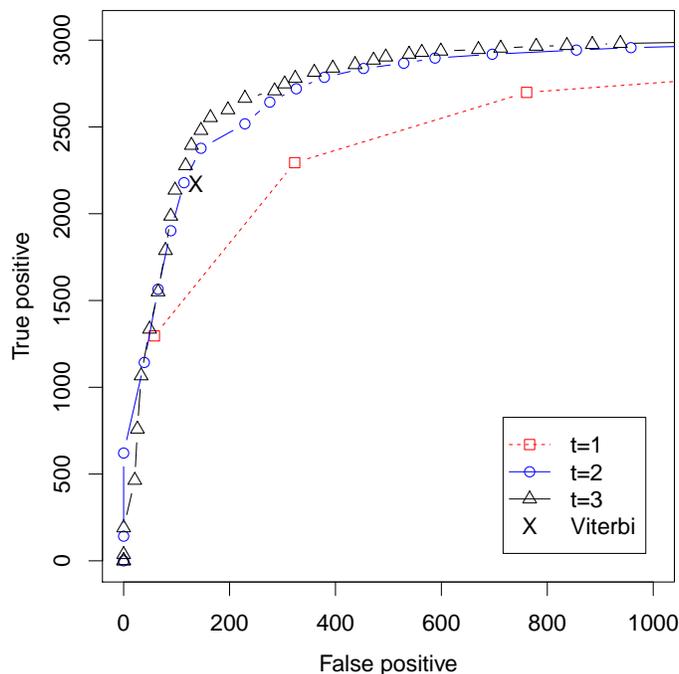
Boža, V., Brejová, B., and Vinař, T. (2016). DeepNano: Deep recurrent neural networks for base calling in MinION nanopore reads. Technical Report arXiv:1603.09195, arXiv.org.

Brejova, B., Brown, D. G., and Vinar, T. (2004). Optimal spaced seeds for homologous coding regions. *Journal of Bioinformatics and Computational Biology*, 1(4):595–610.

Brown, D. G. (2008). A survey of seeding for sequence alignment. In *Bioinformatics algorithms: techniques and applications*, pages 126–152. John Wiley & Sons, Inc.

Cawley, S. L. and Pachter, L. (2003). HMM sampling and applications to gene finding and alternative splicing. *Bioinformatics*, 19(S2):ii36–41.

David, M., Dursi, L. J., Yao, D., Boutros, P. C., and Simpson, J. T. (2016). Nanocall: An open source basecaller



Method	Sn	FP
Viterbi	0.713	136
$t = 1, n = 2$	0.752	323
$t = 1, n = 5$	0.967	2412
$t = 1, n = 13$	0.999	18249
$t = 2, n = 7$	0.715	114
$t = 2, n = 15$	0.950	588
$t = 2, n = 43$	0.999	6407
$t = 3, n = 13$	0.747	117
$t = 3, n = 26$	0.952	495
$t = 3, n = 80$	0.999	6180
$t = 4, n = 18$	0.714	89
$t = 4, n = 39$	0.950	517
$t = 4, n = 139$	0.999	8877
$t = 5, n = 24$	0.721	139
$t = 5, n = 51$	0.954	483
$t = 5, n = 165$	0.999	7003

Figure 4: **Performance of our approach compared to the Viterbi base calls for seeding with a chain of three 10-mers**, each in distance at most 50 from the previous one. The plot and the table have the same form as in Figure 3.

for Oxford Nanopore sequencing data. Technical Report bioRxiv:046086, Cold Spring Harbor Laboratory.

Durbin, R., Eddy, S. R., Krogh, A., and Mitchison, G. (1998). *Biological Sequence Analysis: Probabilistic Models of Proteins and Nucleic Acids*. Cambridge University Press.

Ferragina, P. and Manzini, G. (2000). Opportunistic data structures with applications. In *Foundations of Computer Science (FOCS)*, pages 390–398. IEEE.

Ip, C. L. C., Loose, M., Tyson, J. R., de Cesare, M., Brown, B. L., Jain, M., Leggett, R. M., Eccles, D. A., Zalunin, V., Urban, J. M., Piazza, P., Bowden, R. J., Paten, B., Mwaigwisya, S., Batty, E. M., Simpson, J. T., Snutch, T. P., Birney, E., Buck, D., Goodwin, S., Jansen, H. J., O’Grady, J., and Olsen, H. E. (2015). MinION Analysis and Reference Consortium: Phase 1 data release and analysis. *F1000Research*, 4:1075.

Kielbasa, S. M., Wan, R., Sato, K., Horton, P., and Frith, M. C. (2011). Adaptive seeds tame genomic sequence comparison. *Genome Research*, 21(3):487–493.

Li, H. and Durbin, R. (2010). Fast and accurate long-read alignment with Burrows-Wheeler transform. *Bioinformatics*, 26(5):589–595.

Ma, B., Tromp, J., and Li, M. (2002). PatternHunter: faster and more sensitive homology search. *Bioinformatics*, 18(3):440–445.

Manber, U. and Myers, G. (1993). Suffix arrays: a new method for on-line string searches. *SIAM Journal on Computing*, 22(5):935–948.

Pevzner, P. A., Tang, H., and Waterman, M. S. (2001). An Eulerian path approach to DNA fragment assembly. *Proceedings of the National Academy of Sciences of the USA*, 98(17):9748–9753.

Quick, J. et al. (2016). Real-time, portable genome sequencing for Ebola surveillance. *Nature*, 530(7589):228–232.

Rabiner, L. R. (1989). A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286.

Smith, T. F. and Waterman, M. S. (1981). Identification of common molecular subsequences. *Journal of Molecular Biology*, 147(1):195–197.

Sovic, I., Sikic, M., Wilm, A., Fenlon, S. N., Chen, S., and Nagarajan, N. (2016). Fast and sensitive mapping of nanopore sequencing reads with GraphMap. *Nature Communications*, 7:11307.

Szalay, T. and Golovchenko, J. A. (2015). De novo sequencing and variant calling with nanopores using PoreSeq. *Nature Biotechnology*, 33(10):1087–1091.

Viterbi, A. J. (1967). Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Transactions on Information Theory*, 13(2):260–269.