

1 Atropos: specific, sensitive, and speedy 2 trimming of sequencing reads

3 John P Didion¹ and Francis S Collins¹

4 ¹National Human Genome Research Institute, National Institutes of Health, Bethesda,
5 MD

6 Corresponding author:
7 John P Didion¹

8 Email address: john.didion@nih.gov

9 ABSTRACT

10 **Summary.** A key step in the transformation of raw sequencing reads into biological insights is the
11 trimming of adapter sequences and low-quality bases. Read trimming has been shown to increase
12 the quality and reliability while decreasing the computational requirements of downstream analyses.
13 Many read trimming software tools are available; however, no tool simultaneously provides the accuracy,
14 computational efficiency, and feature set required to handle the types and volumes of data generated in
15 modern sequencing-based experiments. Here we introduce Atropos and show that it trims reads with
16 high sensitivity and specificity while maintaining leading-edge speed. Compared to other state-of-the-art
17 read trimming tools, Atropos achieves a XX% increase in trimming accuracy and a decrease in execution
18 time of at least XX% (using 16 parallel execution threads). Furthermore, Atropos maintains high accuracy
19 even when trimming simulated data with a high rate of error. The accuracy, high performance, and
20 broad feature set offered by Atropos makes it an appropriate choice for the pre-processing of most
21 current-generation sequencing data sets. **Availability.** Atropos is open source and free software written
22 in Python and available at <https://github.com/jdidion/atropos>.

23 1 INTRODUCTION

24 Sequencing on all current-generation sequencing technologies, including Illumina, SOLiD, PacBio, and
25 Nanopore, requires a library construction step that involves ligation of short adapter sequences to the ends
26 of the template DNA fragments. Depending on the sequencing platform and the fragment size distribution
27 of the sequencing library, an often substantial fraction of reads will consist of both template and adapter
28 sequences (Figure 1A). Additionally, the error rates of these sequencing technologies vary from ~0.1% on
29 Illumina to 5% or more on long-read sequencing platforms. Error rates tend to be enriched at the ends
30 of reads (where adapters are located), thus exacerbating the effects of adapter contamination. Adapter
31 contamination and sequencing errors can lead to increased rates of misaligned and unaligned reads, which
32 results in errors in downstream analysis including spurious variant calls (Del Fabbro et al., 2013; Sturm
33 et al., 2016). Certain sequencing protocols may introduce other artifacts in sequencing reads. For example,
34 some methylation sequencing (Methyl-Seq) protocols result in artificially methylated bases at the 3' ends
35 of reads that can result in inflated estimates of methylation levels (Bock, 2012).

36 Read trimming is an important step in the analysis pipeline to mitigate the effects of adapter contami-
37 nation, sequencing errors, and other artifacts. The development of tools for read trimming is an active
38 area of bioinformatics research, thus there are currently many options. In terms of adapter trimming,
39 these tools fall into two general categories (Figure 1B): 1) those that rely solely on matching the adapter
40 sequence (*adapter-match trimming*) using semi-global alignment (which is the only option available for
41 single-end reads); and 2) those that leverage the overlap between paired-end reads to identify adapter
42 starting positions (*insert-match trimming*) (Sturm et al., 2016). Cutadapt (Martin, 2011) is a mature and
43 feature-rich example of a tool that provides adapter-match trimming, while SeqPurge (Sturm et al., 2016)
44 is a recent example of a highly accurate insert-match trimmer designed specifically for paired-end data.
45 Additionally, hybrid tools are available that optimize their choice of read trimming method based on
46 the type of data. Skewer (Jiang et al., 2014) is a fast and accurate hybrid trimmer that works with both

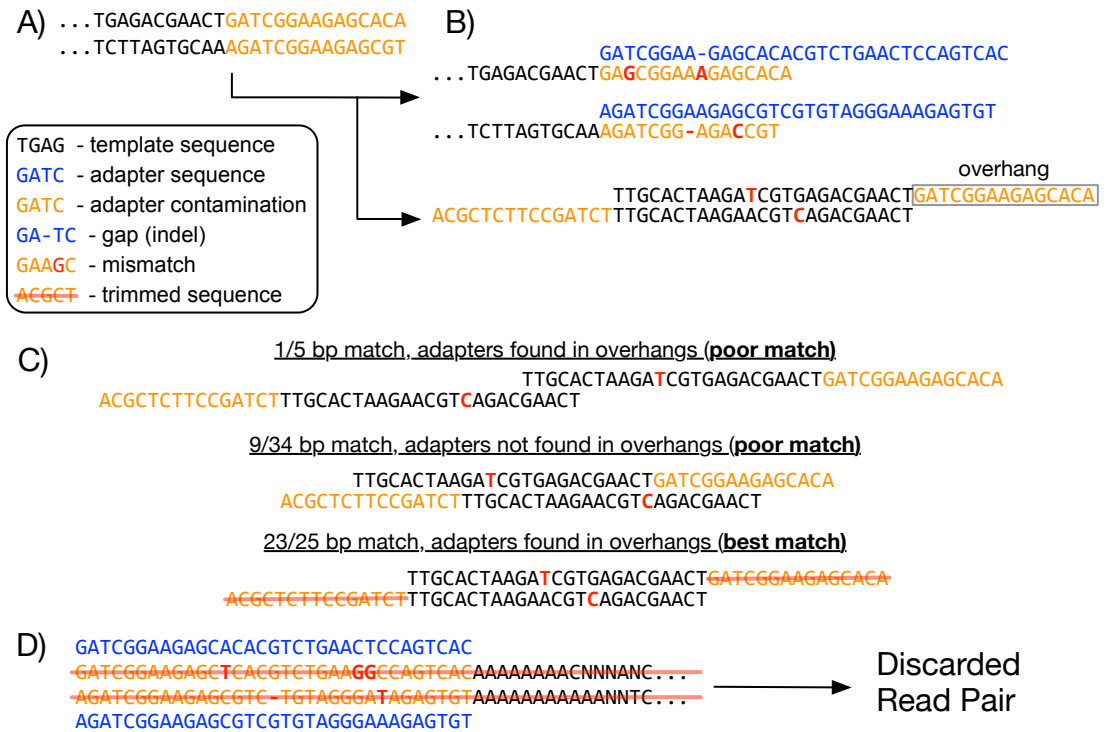


Figure 1. Adapter detection and trimming. A) When a fragment (insert) is shorter than the read length, the read sequence will contain partial to full-length adapter sequences (orange). B) Adapters contamination may be detected by adapter-match (top) or insert-match (bottom). C) Atropos uses semi-global alignment to find the best insert match (if any) for which the overhangs contain adapter contamination. D) If insert-match fails – for example when the adapter appears at or near the beginning of the read – adapter-match is optionally performed. Reads that are too short after trimming are discarded.

47 single-end and paired-end data. However, choosing a read-trimming tool currently requires a trade-off
48 between feature set, efficiency, and accuracy. Furthermore, even state-of-the-art tools still have a relatively
49 high rate of over-trimming (removing usable template bases from reads) and/or under-trimming (leaving
50 low-quality and adapter-derived bases in the read sequence) (Sturm et al., 2016).

51 We sought to develop a read trimming tool that would combine the best aspects of currently available
52 software to provide high speed and accuracy while also offering a rich feature set. To accomplish this aim,
53 we used Cutadapt as a starting point, as it provides the broadest feature set of currently available tools and
54 is published under the MIT license, which allows modification and improvement of the code. We focused
55 on making three specific improvements to Cutadapt: 1) improve the accuracy of paired-end read trimming
56 by implementing an insert-match algorithm; 2) improve the performance by adding multiprocessing
57 support (as it is currently only able to use a single processor); and 3) add important additional features
58 such as automated trimming of Methyl-Seq reads and automated detection of adapter sequences in reads
59 where the experimental protocols are not known to the analyst. Because these modifications required
60 substantial changes to the Cutadapt code base, and because there are software tools that depend on the
61 current implementation of Cutadapt, we chose to "fork" the Cutadapt code base and develop our software,
62 Atropos, as a separate tool. Here, we show that we have accomplished our three aims. In addition to
63 extending the already rich set of features provided by the original Cutadapt tool, Atropos demonstrates
64 paired-end read trimming accuracy that is superior to other state-of-the-art tools, and it is among the
65 fastest read trimming tools when a moderate number of parallel threads are used (4 or 8). Furthermore,
66 Atropos achieves a performance increase that is roughly linear with the number of threads used, making it
67 the fastest tool when 16 or more threads are available.

68 2 IMPLEMENTATION

69 Atropos is developed in Python and is available to install from GitHub or via the pip package manager
70 (see Data Availability).

71 2.1 Insert Match Algorithm

72 For each read pair, the insert-match algorithm finds all possible alignments between the first read and
73 the reverse complement of the second read that satisfy user-specified minimum overlap and maximum
74 mismatch thresholds (Figure 1C). The candidate alignments are tested in order of decreasing length
75 until one is found in which the overhanging sequences on either end match the user-specified adapter
76 sequences.

77 To determine whether an adapter sequence is present in the overhang, our algorithm uses the same
78 probabilistic approach as in the SeqPurge algorithm (Sturm et al., 2016). Briefly, starting at the end of the
79 insert overlap, a pairwise comparison is made between the adapter and the read at each possible offset.
80 The probability of a random match at k bases out of the n bases being compared is computed using the
81 binomial distribution:

$$P = \sum_{i=k}^n \frac{n!}{i!(n-i)!} p^i (1-p)^{n-i} \quad (1)$$

82 The offset having the lowest random-match probability is taken to be the location of the adapter
83 sequence, and all bases from that position to the 3' end of the read are removed. If an adapter is only
84 found in one of the two reads, then the same offset is used to trim both reads, under the assumption that
85 the location of the adapter sequence must be symmetric across the read pair. If no adapter is found, then
86 an adapter-match approach is optionally attempted in each read separately (Figure 1D).

87 2.2 Parallel processing

88 The performance improvements in Atropos relative to Cutadapt and other read trimming tools are based in
89 two observations: 1) each read (or read pair) is trimmed separately, and thus trimming can be parallelized
90 across multiple processor cores, and 2) a significant fraction of the execution time is spent decompressing
91 input files and re-compressing results. Compression of sequencing data is increasingly becoming necessary
92 due to the large volumes of data generated in sequencing experiments.

93 To address the first bottleneck, we implemented a parallel processing pipeline based on the Python
94 `multiprocessing` module. Briefly, a single thread is dedicated to a "reader" process that loads reads
95 (or read pairs) from input file(s), with support for a variety of data formats and automatic decompression
96 of compressed data. Reads are loaded in batches, and each batch is added to an in-memory queue. A user-
97 specified number of "worker" threads (which is constrained by the number of processing cores available
98 on the user's system) extract batches from the queue and perform trimming and filtering operations
99 on the reads in the same manner as Cutadapt. Atropos addresses the second bottleneck by offering a
100 choice of three modes for writing the results to disk. The first two modes involve adding the results to a
101 second queue, from which a dedicated "writer" process extracts batches and performs the serialized write
102 operation. These modes differ in how the trimmed reads are compressed – in "worker compression" mode,
103 each worker is responsible for compressing the results using the Python `gzip` module prior to placing
104 the results on the queue, whereas in "writer compression" mode, the writer process performs compression
105 using the much faster system-level `gzip` program. The choice between these two modes is selected
106 automatically based on the number of worker threads used, with worker compression becoming faster
107 than writer compression when at least 8 threads are available. The third output mode, called "parallel
108 writing," does not use a dedicated writer process (and thus an additional worker process can be used in
109 its place). Instead, each worker process writes its results to a separate file. This can dramatically reduce
110 the execution time of the program (~100% reduction in our experiments; see Results) and is generally
111 compatible with downstream analysis since most modern mapping and assembly tools accept multiple
112 input files. An additional speed-up is gained by recognizing that the reader process often finishes loading
113 data well before the worker processes finish processing it, thus an additional worker thread is started as
114 soon as the reader process completes.

115 2.3 Adapter detection

116 Often, details of sequencing library construction are not fully communicated from the individual or
 117 facility that generated the library to the individual(s) performing data analysis. Manual determination of
 118 sequencing adapters and other potential library contaminants can be a tedious and error-prone task. Thus,
 119 we implemented in Atropos a new 'detect' command that automatically identifies adapters/contaminants
 120 from a sample of read sequences. First, a profile is built of k -mers (where k is a fixed number of
 121 consecutive nucleotides, defaulting to $k = 12$) within N read sequences (where N defaults to 10,000).
 122 When at least 8 consecutive A bases are detected, those bases along with all subsequent bases (in the 3'
 123 direction) are first trimmed, as that pattern is a strong indicator that the sequencer scanned past the end of
 124 the read (i.e. the length of the fragment + adapter is less than the read length; Figure 1D). Additionally,
 125 low-complexity reads are excluded, where complexity $X(S)$ is defined as follows. Let $C(i, S)$ be the
 126 number of elements of a nucleotide sequence $S = s_1, \dots, s_n$, that are nucleotide $i \in A, C, G, T$.

$$X(S) = -\sum \frac{C(i, S) * \log(C(i, S))}{\log(2)} \quad (2)$$

127 Sequences with $X(S) < 1.0$ are defined as low-complexity. All remaining k -mers are counted, and
 128 each k -mer is linked to all of the sequences from which it originated. This process continues iteratively for
 129 increasing values of k , with only those read sequences linked to high-abundance k -mers in the previous
 130 iteration being used to build the k -mer profile in the next iteration. k -mer K is considered high-abundance
 131 when:

$$|K| > \frac{N * (l - k + 1) * O}{4^k} \quad (3)$$

132 where l is the read length and $O = 100$ by default. Finally, high-abundance k -mers of all lengths
 133 are merged to eliminate shorter sequences that are fully contained in longer sequences. Optionally, the
 134 high-abundance k -mers are matched to a list of known adapters/contaminants. Atropos reports to the user
 135 an ordered list of up to 20 of the most likely contaminants.

136 3 RESULTS

137 3.1 Benchmarks

Data Set	Error Rate	Read Length	Total Read Pairs	Reads w/ Adapters	Adapter Bases
Simulated 1	0.20%	125	781,923	325,982	12,447,262
Simulated 2	0.60%	125	780,899	325,105	12,416,861
Simulated 3	1.20%	125	782,237	325,860	12,464,235
XXX Amplicon	Unknown	???	1,000,000	???	???
GM12878 WGBS	Unknown	125	1,000,000	16,999*	1,020,017*

Table 1. Description of data sets. *For the real data sets, total adapter content is unknown; we provide the number of full-length adapters and adapter bases present.

138 We evaluated both the speed and the accuracy of Atropos relative to other state-of-the-art read
 139 trimming tools using both simulated and real-world data (Table 1). As trimming of single-end reads
 140 is unchanged from the original Cutadapt method and is also decreasing in relevance as most current
 141 experiments used paired-end data, we focused our benchmarking on trimming of paired-end reads. Sturm
 142 et al. demonstrate that Skewer (Jiang et al., 2014) and SeqPurge (Sturm et al., 2016) stand out as having
 143 superior performance in paired-end read trimming; thus, we chose to benchmark Atropos against these
 144 tools. We also evaluated the accuracy of our new insert-match algorithm against the original Cutadapt
 145 adapter-match algorithm.

146 To simulate paired-end read data, we use the ART simulator (Huang et al., 2012) that was modified
 147 by Jiang et al. to add adapter sequences to the ends of simulated fragments. ART simulates reads based
 148 on empirically derived quality profiles specific to each sequencing platform. A quality profile consists

149 of distributions of quality scores for each nucleotide at each read position, expressed as read counts
 150 aggregated from multiple sequencing experiments, where quality scores are in Phred scale ($-10\log_{10}(e)$,
 151 where e is the probability that the base call is erroneous). We developed an R script to artificially inflate
 152 the error rates in an ART profile to a user-defined level. For each row in the profile with quality score bins
 153 $e_1..e_n$ and corresponding read counts $r_1..r_n$, the overall error rate can be computed as:

$$E = \frac{\sum_{i=1}^n e_i r_i}{\sum_{i=1}^n r_i} \quad (4)$$

154 We use the R function *optim* with the variable metric ("BFGS") algorithm to optimize a function
 155 that adds an equal number of counts C to each Phred-score bin in the distribution and then compares the
 156 overall error rate to the user-desired error rate E' :

$$f(C, E') = \frac{\sum_{i=1}^n e_i (r_i + C)}{\sum_{i=1}^n (r_i + C)} - E' \quad (5)$$

157 We simulated ~800k 125 bp paired-end reads using the Illumina 2500 profile at error rates that were
 158 low/typical (~0.2%, the unmodified profile), intermediate (~0.6%), and high (~1.2%).

159 We evaluated the accuracy of the tools on the simulated data by comparing each trimmed read pair
 160 to the known template sequence. We counted the frequency of following outcomes: the fragment does
 161 not contain adapters but is trimmed anyway ("wrongly trimmed"), the fragment is under-trimmed, or the
 162 fragment is over-trimmed. We also counted the total number of under- and over-trimmed bases.

163 We also benchmarked the tools on a two real-world data sets of 1M read pairs each: an amplicon
 164 sequencing data set generated in our lab (REF: Maria's paper) and 1M randomly sampled read pairs from
 165 a whole-genome bisulfite sequencing (WGBS) library generated from the GM12878 cell line as part of
 166 the ENCODE project (ENCODE Project Consortium, 2012). Since the genomic origins of the templates
 167 are not known *a priori*, we instead compared the read trimming tools based on improvement in the results
 168 of mapping the untrimmed versus trimmed reads. We used BWA MEM (Li, 2013) to map the amplicon
 169 reads to human genome build 37 (GRCh37), and we used bwa-meth (Pedersen et al., 2014) to map the
 170 WGBS reads to the bisulfite-converted GRCh37. We also compared the results of only adapter trimming
 171 to the results of adapter trimming plus additional quality trimming, using a minimum quality threshold of
 172 20. The supplement describes in detail how bisulfite reads are mapped and how we used the mapping
 173 information to determine the effects of read trimming.

174 Although sequence analysis is sometimes performed using a desktop computer, analysis of the volumes
 175 of data currently being generated increasingly requires the use of high-throughput computing facilities
 176 ("clusters"). The hardware architecture of a cluster is often different from that of a desktop computer.
 177 Most importantly, storage in a cluster is typically centralized and accessed by the compute nodes via
 178 high-speed networking. Such an architecture inevitably adds latency to file reading and writing operations
 179 ("I/O"). Cluster nodes also typically have more processing cores and memory available than desktop
 180 computers. This means that the performance of software with intensive I/O usage (such as read trimming)
 181 is likely to be quite different on a desktop versus a cluster. To examine the impact of these architectural
 182 differences, we ran all benchmarks on both a desktop computer (a Mac Pro) having a 3.7 GHz quad-core
 183 Intel Xeon E5 processor and 32 GB RAM, and on a cluster node having 4 X 3.2 GHz eight-core Intel Xeon
 184 E5 processors and 128 Gb memory, and with all data being read from and written to network-accessible
 185 storage.

186 Finally, we provide scripts in our GitHub repository that can be used to re-run our analysis, and that
 187 also enable other tool developers to benchmark their software against Atropos.

188 3.2 Simulated Data

189 3.2.1 Performance

190 On a desktop computer with 4 processing cores, we found that Skewer had the fastest overall execution
 191 time and SeqPurge the slowest (2). Atropos running in parallel-write mode and using the adapter-match
 192 algorithm was nearly as fast as Skewer.

193 As expected, execution times on a cluster node using 4 threads were approximately twice those
 194 observed on a desktop computer (3). Part of this difference may be explained by a slight difference in

Program	Execution Time (sec.)	
	Min	Max
Atropos (adapter + writer compr.)	41.07	55.51
Atropos (adapter + parallel write)	23.18	30.75
Atropos (insert + writer compr.)	41.75	56.22
Atropos (insert + parallel write)	32.98	41.33
SeqPurge	96.67	114.24
Skewer	20.75	29.40

Table 2. Execution time for programs running on desktop with 4 threads. Each program was executed multiple times, and Atropos was run with all combinations of alignment algorithm (insert-match or adapter-match) and output mode (writer-compression or parallel-write). The minimum and maximum execution times for each program are shown, with the lowest overall execution time in bold.

Program	Execution Time (Min Max sec.)					
	4 Threads	8 Threads		16 Threads		
Atropos (adapter + worker compr.)	99.39	493.85	56.64	99.05	30.38	64.05
Atropos (adapter + writer compr.)	75.87	346.80	77.63	109.05	72.37	102.00
Atropos (adapter + parallel write)	46.35	121.18	27.63	34.39	21.27	25.40
Atropos (insert + worker compr.)	115.28	292.61	65.36	100.84	34.72	53.82
Atropos (insert + writer compr.)	77.27	119.45	77.53	108.97	72.01	109.27
Atropos (insert + parallel write)	65.30	137.21	36.65	47.14	21.57	28.11
SeqPurge	372.54	517.83	192.85	226.01	94.88	122.39
Skewer	40.12	122.07	40.22	58.78	40.80	100.71
		Memory Usage (Min Max Gb)				
Atropos	1.2	2.0	0.011	0.012	0.008	0.008
SeqPurge	0.47	2.0	0.011	0.012	0.009	0.011
Skewer	0.51	1.8	0.011	0.012	0.012	0.014

Table 3. Top: Execution time for programs running on a cluster with 4, 8, or 16 threads. Each program was executed multiple times, and Atropos was run with all combinations of alignment algorithm (insert-match or adapter-match) and output mode (worker-compression, writer-compression or parallel-write). The minimum and maximum execution times for each program are shown, and notable results are highlighted in bold. Bottom: The minimum and maximum memory usage for each program across all executions.

processor speeds (3.7 GHz on the desktop versus 3.2 GHz on the cluster node); however, we expect that most of the difference is due to the increased latency involved in network-based I/O on the cluster.

Interestingly, when increasing the number of threads from 4 to 8 and 16, the execution time of Skewer remains unchanged, while Atropos and SeqPurge achieve roughly linear decreases in execution time. With 8 and 16 threads, Atropos in parallel-write mode is the fastest of the tools, and with 16 threads Atropos in worker-compression mode is also faster than SeqPurge and Skewer.

On the other hand, Atropos uses significantly more memory than SeqPurge or Skewer (3). We expect this is partially due to overhead of automatic memory management in Python compared to C++ (in which SeqPurge and Skewer are implemented), but in larger part results from Atropos' use of in-memory queues to communicate between parallel processes. We note that Atropos provides parameters to limit memory usage (although typically at the expense of performance).

3.2.2 Accuracy

We found that the four adapter matching algorithms had different biases toward under- and over-trimming (Table 4). Across the three sequencing error rates, Skewer had the lowest frequency of wrongly trimming reads while SeqPurge had the highest. The Atropos adapter-match algorithm and Skewer had similarly low frequencies of over-trimming, while the Atropos insert-match algorithm and SeqPurge had similarly low frequencies of over-trimming. Overall, the Atropos insert-match algorithm and SeqPurge demonstrated the lowest error rates (0.01%).

Program	Reads				Bases		
	Wrongly Trimmed	Over-trimmed	Under-trimmed	Total Error	Over-trimmed	Under-trimmed	Total Error
Error rate 0.2%							
Atropos							
adapter-match	51 (0.01%)	1 (0.00%)	28,991 (9.72%)	3.71%	490	102,133	0.054%
insert-match	134 (0.03%)	26 (0.01%)	27 (0.01%)	0.01%	1,036	66	0.001%
SeqPurge	2,639 (0.55%)	37 (0.01%)	29 (0.01%)	0.01%	7,708	29	0.004%
Skewer	10 (0.00%)	7 (0.00%)	283 (0.09%)	0.04%	21	22,029	0.012%
Error rate 0.6%							
Atropos							
adapter-match	72 (0.01%)	6 (0.00%)	28,843 (9.68%)	3.69%	733	101,839	0.054%
insert-match	168 (0.03%)	14 (0.00%)	31 (0.01%)	0.01%	1,434	62	0.001%
SeqPurge	2,595 (0.54%)	26 (0.01%)	38 (0.01%)	0.01%	6,701	41	0.004%
Skewer	6 (0.00%)	3 (0.00%)	413 (0.14%)	0.05%	9	36,685	0.019%
Error rate 1.2%							
Atropos							
adapter-match	76 (0.02%)	5 (0.00%)	30,152 (10.12%)	3.86%	721	117,027	0.062%
insert-match	175 (0.04%)	12 (0.00%)	46 (0.02%)	0.01%	1,440	110	0.001%
SeqPurge	2,652 (0.55%)	24 (0.01%)	34 (0.01%)	0.01%	7,628	36	0.004%
Skewer	7 (0.00%)	5 (0.00%)	644 (0.22%)	0.08%	12	67,066	0.035%

Table 4. Trimming accuracy on simulated data with three different base-call error rates.

213 In terms of numbers of over- and under-trimmed bases, the Atropos insert-match algorithm and
 214 SeqPurge clearly had the best performance (Table 4) at all sequencing error rates. The two algorithms
 215 had similarly low numbers of under-trimmed bases, but the Atropos insert-match algorithm had a lower
 216 number of over-trimmed bases, giving it the lowest overall error rate (0.001%). On the other hand, Skewer
 217 and the Atropos adapter-match algorithm left substantial numbers of under-trimmed bases, resulting in
 218 about 10-fold higher overall error rates.

219 Additionally, we found that all tools discarded very similar numbers of reads (~1.8%) that were below
 220 the minimum length threshold of 25 bp after trimming. These were reads with short insert sizes, which
 221 have a high rate of spurious mapping, and thus it is common practice to discard them.

222 3.3 Real Data

223 We first tested Atropos' adapter detection module on the real data sets. Using the first 10,000 reads in each
 224 pair of FASTQ files, Atropos correctly detected that sequences of the adapters used in constructing each
 225 library. In the case of the WGBS data set, the adapters were found in the list of known contaminants (read
 226 1: "TruSeq Adapter, Index 7" and read 2: "TruSeq Universal Adapter"). On the other hand, a custom
 227 index was used in the read 1 adapter of the amplicon library. Atropos correctly identifies the custom index
 228 (AACGTGAT), and also correctly identifies the TruSeq Universal Adapter as the read 2 adapter.

229 3.3.1 Performance

230 We performed adapter trimming on the real data set in a cluster environments using 16 parallel cores.
 231 Again, we found Atropos (in worker-compression mode) to have the best performance both with and
 232 without additional quality trimming (Table 5). We also performed read mapping on the cluster with 16
 233 cores and observed that the trimmed reads produced by each combination of tool and quality threshold
 234 resulted in substantially different execution times (Table 5). The mapping of reads trimmed by Atropos
 235 using the adapter-match algorithm took twice as long as those trimmed by any other tool, while the reads
 236 trimmed by Atropos using the insert-match algorithm and no additional quality trimming took the least
 237 amount of time to map. At a quality trimming threshold of 20, the trimmed reads generated by SeqPurge
 238 mapped slightly faster than those generated by Atropos.

Program	Trimming Time (sec.)		Mapping Time (min.)	
	Q0	Q20	Q0	Q20
WGBS Data				
Untrimmed reads			45:45	
Atropos	56.46	55.58	37:32	30:37
SeqPurge	70.37	80.33	39:10	31:55
Skewer	69.26	72.21	36:06	32:15

Table 5. Execution and read mapping times for programs trimming real data on a cluster with 16 threads. Each program was executed with no additional quality trimming (Q0) and with quality trimming at a minimum base quality of 20 (Q20). Atropos was run with the insert insert-match algorithm in worker-compression mode. SeqPurge and Atropos were run with error correction enabled (Skewer performs error correction by default).

239 3.3.2 Accuracy

240 We chose to assess read trimming accuracy in practical terms – as the number of trimmed reads mapped
241 at at given quality (MAPQ) cutoff, relative to the number of untrimmed reads mapped at that cutoff.
242 We found that trimming by Atropos resulted in the greatest increase in number of mapped reads at all
243 quality cutoffs (Figure 2). We found that trimming with SeqPurge resulted in similar, but smaller, gains in
244 mapping quality, while trimming with Skewer resulted in the smallest gains in quality.

245 Additionally, we found that additional quality trimming in addition to adapter trimming has a substan-
246 tial negative effect on read mapping, at least for bisulfite reads mapped using bwa-meth. However, quality
247 trimming by Skewer had the least negative effect on mapping quality of the three programs, and quality
248 trimming by SeqPurge had the greatest negative effect on mapping quality.

249 4 CONCLUSIONS

250 Our results demonstrate that adapter trimming tools are approaching optimal accuracy, at least for the
251 (currently) most common type of data – paired-end short reads with 3' adapters. On synthetic data with
252 varying error rates, Atropos (using our new insert-match algorithm) and SeqPurge both demonstrated
253 nearly negligible overall error rates of 0.01% at the read level, and Atropos has the lowest base-level error
254 rate of 0.001%.

255 On real data, we also find that Atropos has superior performance, leading to the greatest increase in
256 read mapping quality. We also find unequivocally that stringent quality trimming has a negative effect on
257 WGBS data. For reads trimmed with a quality threshold of 20, all mapping statistics are worse than those
258 for untrimmed reads.

259 In terms of performance, Atropos is faster than SeqPurge, and is the overall fastest tool when at least
260 16 processing cores are available.

261 Thus, overall, Atropos offers the best combination of accuracy and performance of the tools that
262 we evaluated. Furthermore, Atropos has the richest feature set of the three tools, including Methyl-
263 Seq-specific trimming options, automated adapter detection, and support for data generated by many
264 sequencing methods (ABI SOLiD, Illumina NextSeq, mate-pair libraries, and single-end sequencing).
265 Finally, Atropos is easily installable on any system with Python 3.3+ (using the command 'pip install
266 atropos').

267 5 DATA AVAILABILITY

- 268 • Atropos can be installed from the Python Package Index (pypi) using the *pip* tool: 'pip install
269 atropos'.
- 270 • The Atropos source code, including all scripts needed to execute the analyses in this paper, are
271 available at <https://github.com/jdidion/atropos>. The portions of Atropos developed
272 by JPD are a work of the US government, and thus all copyright is waived under a CC0 1.0 Universal
273 Public Domain Dedication ([https://creativecommons.org/publicdomain/zero/
274 1.0/](https://creativecommons.org/publicdomain/zero/1.0/)).

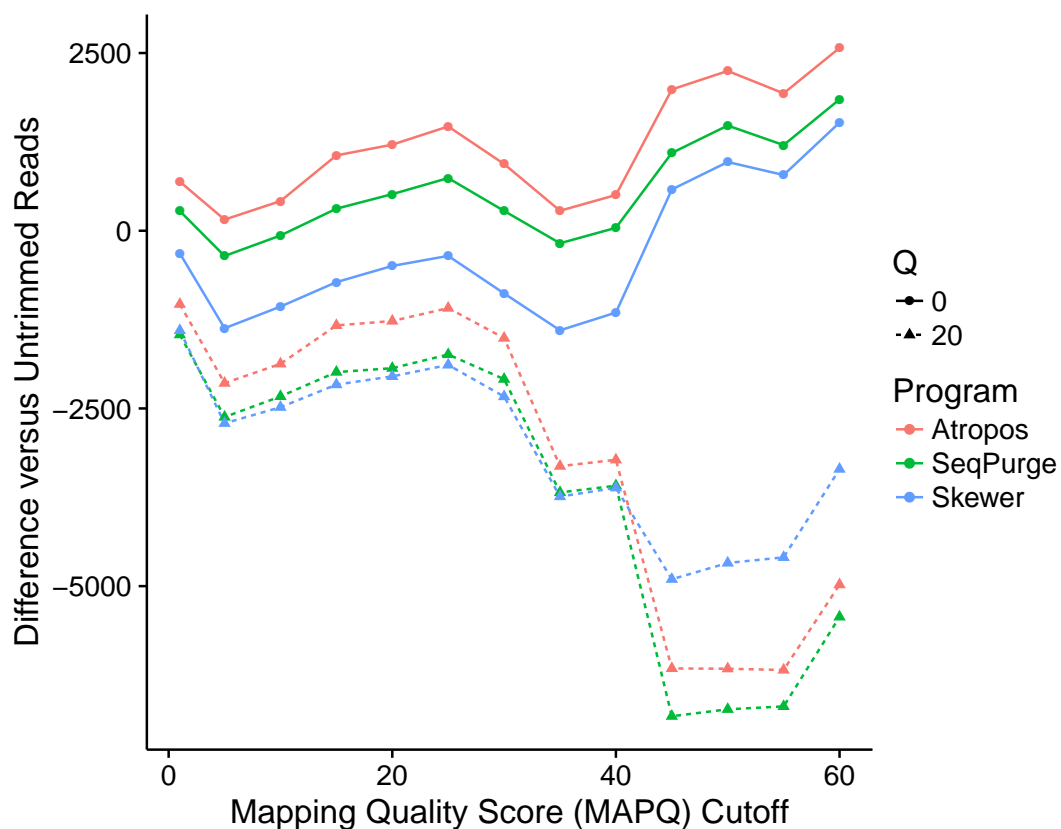


Figure 2. Atropos trimming best improves mapping of real sequencing reads. Difference between the number of trimmed and untrimmed reads mapped with quality (MAPQ) greater than or equal to each cutoff (1..60 at intervals of 5). We mapped both untrimmed and trimmed reads to the genome. Reads were trimmed with all three programs, both without additional quality trimming (Q=0) and with quality trimming at a threshold of Q=20.

- 275 • The simulated data sets are also available in the Atropos GitHub repository.
- 276 • The real amplicon data set is available at XXX.
- 277 • The real GM12878 WGBS data (accession ENCLB794YYH) is available from the ENCODE project
- 278 website: <https://www.encodeproject.org/experiments/ENCSR890UQ0/>.

279 6 ACKNOWLEDGEMENTS

280 We would like to thank Marcel Martin and colleagues for publishing Cutadapt under an open source
 281 license that is compatible with modifications and improvements. JPD and FSC are funded by the NIH
 282 intramural program.

283 REFERENCES

- 284 Bock, C. (2012). Analysing and interpreting DNA methylation data. *Nature Reviews Genetics*, 13(10):705–
 285 719.
- 286 Del Fabbro, C., Scalabrin, S., Morgante, M., and Giorgi, F. M. (2013). An extensive evaluation of read
 287 trimming effects on Illumina NGS data analysis. *PLoS One*, 8(12):e85024–None.
- 288 ENCODE Project Consortium (2012). An integrated encyclopedia of DNA elements in the human genome.
 289 *Nature*, 489(7414):57–74.
- 290 Huang, W., Li, L., Myers, J. R., and Marth, G. T. (2012). ART: a next-generation sequencing read
 291 simulator. *Bioinformatics*, 28(4):593–594.

- 292 Jiang, H., Lei, R., Ding, S.-W., and Zhu, S. (2014). Skewer: a fast and accurate adapter trimmer for
293 next-generation sequencing paired-end reads. *BMC Bioinformatics*, 15:182–None.
- 294 Li, H. (2013). Aligning sequence reads, clone sequences and assembly contigs with BWA-MEM.
295 *arXiv:1303.3997 [q-bio]*. arXiv: 1303.3997.
- 296 Martin, M. (2011). Cutadapt removes adapter sequences from high-throughput sequencing reads. *EMB-*
297 *net.journal*, 17(1):pp. 10–12.
- 298 Pedersen, B. S., Eyring, K., De, S., Yang, I. V., and Schwartz, D. A. (2014). Fast and accurate alignment
299 of long bisulfite-seq reads. *arXiv:1401.1129 [q-bio]*. arXiv: 1401.1129.
- 300 Sturm, M., Schroeder, C., and Bauer, P. (2016). SeqPurge: highly-sensitive adapter trimming for
301 paired-end NGS data. *BMC Bioinformatics*, 17:208.