ExtendAlign: a computational algorithm for delivering multiple global alignment results originated from local alignments.

Mariana Flores-Torres,[1] Laura Gómez-Romero,[2] Joshua I. Haase-Hernández,[3] Israel Aguilar-Ordóñez,[4] Hugo Tovar,[2] S. Eréndira Avendaño-Vázquez[1]* and C. Fabián Flores-Jasso[1,5]*

[1]Consorcio de Metabolismo de RNA. Instituto Nacional de Medicina Genómica, INMEGEN. Periférico Sur 4809, Ciudad de México 14610, México.

[2]Unidad de Servicios Bioinformáticos. Instituto Nacional de Medicina Genómica, INMEGEN. Periférico Sur 4809, Ciudad de México 14610, México.

[3]Subdirección de Bioinformática. Instituto Nacional de Medicina Genómica, INMEGEN. Periférico Sur 4809, Ciudad de México 14610, México.

[4]Subdirección de Investigación Básica. Instituto Nacional de Medicina Genómica, INMEGEN. Periférico Sur 4809, Ciudad de México 14610, México.

[5]Lead contact.

*Correspondence:

S.E.A-V. (seavendano@inmegen.gob.mx), C.F.F-J. (cfflores@inmegen.gob.mx)

ABSTRACT
MicroRNAs are conserved small RNAs that mediate gene silencing. Although the first plant and animal microRNAs reported in 1993 were considered an abundant class of RNAs, it is with the advent of high throughput sequencing that a tremendous amount of conserved and non-conserved microRNA families has been uncovered across several clades. The search for microRNAs has motivated the development of computational tools that incorporate a diverse set of parameters that aid in the recognition of novel candidate sequences, or the assignation into microRNA families by hierarchical clustering methods. In order to establish homology, multiple sequence alignment tools (MSATs) are used to determine the identity between highly similar sequences. However, as the similarity between two or more sequences declines, the MSATs that perform alignments locally have difficulties in reporting the correct number of matches and mismatches (m/mm) for entire short sequences, compromising the alignment and creating biased results for multiple novel microRNAs against dissimilar subject sequences, or distant genomes. Conversely, the refinements provided by MSATs that perform alignments globally that deal with dissimilar sequences, have a better performance in establishing identity among sequences of similar length, limiting the identification of multiple alignment hits for multiple microRNA sequence queries against distant genomes. To address this, we developed ExtendAlign, a computational tool that corrects any unreported m/mm originated from a local multiple alignment. ExtendAlign extends the alignment achieved by a local MSAT and provides an end-to-end report of true m/mm for every hit in each query sequence, reducing the aforementioned alignment bias. We tested the efficacy of ExtendAlign in calculating the number of true m/mm that resulted from a local multiple alignment among mature and precursor microRNAs, and also distant genomes. Our results showed that ExtendAlign increases significantly the number of m/mm originally missed by an MSAT in all alignments tested. Remarkably, ExtendAlign corrects the alignment hits of dissimilar sequences in the range of ~35–50% similarity —also known as the twilight zone; suggesting that it could be used also as a regular procedure if high accuracy is required after any local multiple alignment aimed to establish a phylogenetic origin among distant microRNAs, or any other short sequences. ExtendAlign is available for download at https://github.com/Flores-JassoLab.

INTRODUCTION

microRNAs are conserved small RNAs that mediate gene silencing by binding mRNAs by base complementarity at their 3′-UTRs (David P. Bartel, 2009). Since the discovery of the first microRNA in *C. elegans*, a myriad of them have been discovered in several other species; thus unveiling an important layer of gene expression regulation (Lee *et al.*, 1993; David P. Bartel, 2009; Bartel, 2018). For this reason, experimental and computational methods have been used to search for conserved and non-conserved microRNA families among related and unrelated species (Meunier J, Lemoine F, Soumillon M, Liechti A, Weier M, Guschanski K, Hu H, Khaitovich P, 2013; Rajendiran *et al.*, 2018; Friedländer *et al.*, 2014; Artzi *et al.*, 2008). The microRNA database (miRBase) is the repository of all microRNAs in all metazoa, viridiplantae, protoctista and viruses inspected so far. The first microRNAs deposited in early versions of miRBase were identified experimentally, in part due to their high abundance (Griffiths-Jones, 2004). However, as deep sequencing has become more available for most laboratories, the number of sequences reported as microRNAs has increased considerably, some detected even at very low expression levels. Currently, the vast majority of microRNAs deposited in miRBase have been predicted from small RNA deep sequencing experiments and computational analyses that seek for homologous microRNA sequences in related and unrelated genomes, and by *de novo* discovery (Kozomara and Griffiths-Jones, 2014). Because new microRNA candidates deposited in miRBase began to accumulate swiftly, a new set of rules for novel microRNA annotation based on structural features has emerged. These rules —comprising: conservation, abundance, overhanging 3′ ends, percent of reads sharing the same 5′ end, precursor folding free energy (< –0.2 kcal/mol/nt), and minimum percent of bases forming the stem in the hairpin structure— yielded a high confidence microRNA data set, that alongside the complete collection reported thus far, are meant to minimize the false-positive discovery rate observed after the explosion of recent sequencing studies (Kozomara and Griffiths-Jones, 2014).

One important criterion for microRNA annotation is conservation. Sequence alignment is usually the first step performed in bioinformatics to investigate the molecular phylogeny of an unknown sequence as it provides an identity measure between two or more compared sequences (Allmer, 2014; Rajendiran *et al.*, 2018; Guerra-Assunção and Enright, 2010; Griffiths-Jones, 2004; David P Bartel, 2009; Friedman *et al.*, 2009). In order to obtain

the sequence identity, many algorithms used for microRNA discovery and annotation rely on other computational tools to perform the sequence alignments as the first step to then determine if a microRNA is to be considered either homologous, or non-homologous (Allmer, 2014). From the general types of computational alignment methods currently available (local and global alignments), the tools that perform multiple sequence alignments *locally* are preferred for microRNA discovery since it is more straightforward to locate several regions sharing certain identity within a genome (Chowdhury and Garai, 2017), enabling to discover whether a microRNA has multiple copies (or hits) in a genome, or whether it belongs to a larger family cluster (Tanzer and Stadler, 2004; Pfeffer *et al.*, 2005; Smalheiser and Torvik, 2005; He and Hannon, 2004; Roush and Slack, 2008; Arteaga-Vazquez *et al.*, 2006; Meunier J, Lemoine F, Soumillon M, Liechti A, Weier M, Guschanski K, Hu H, Khaitovich P, 2013; Guerra-Assunção and Enright, 2010). In contrast, multiple sequence alignment tools (MSATs) that perform alignments *globally* with hierarchical clustering, like methods based on the Needleman-Wunsch algorithm, are employed to align typically sequences of similar length and whose relationships have already been inferred; consequently, are not suitable for identifying multiple hits of a series of shorts queries over large sequences or genomes as subjects (Needleman and Wunsch, 1970; Corpet, 1988; Edgar, 2004; Notredame *et al.*, 2000).

Albeit most MSATs that perform local alignments (like BLASTn, BWA, Bowtie, etc.) are robust in dealing with sequences above 100 nt long, there is one limitation in the alignment of sequence queries shorter than 30 nt long: as the size and identity between two sequences declines, some small stretches of sequences flanking the alignment achieved by the MSAT, that are shorter than the word size and also contain multiple mismatches or gaps, end up unreported. This provokes a match and mismatch (m/mm) report bias for some alignments (Li and Durbin, 2009; Langmead *et al.*, 2009; Polyanovsky *et al.*, 2011; Li and Homer, 2010).

This situation is not entirely solved by the sequence alignment tools that perform alignments globally —although alignments between dissimilar sequences are conducted in end-to-end manner (for instance, MSATs based on the Neddleman-Wunsch algorithm), they typically work in and paired-way manner, may not report the number of m/mm encountered by a hit query accurately, and alignments are usually found at one end on the subject when the two sequences compared have a large difference in size. This constrains

global alignments to sequences of similar length that seek only for the best alignment hit. As a consequence, if one is interested in finding multiple alignment hits for a short sequence in a genome, the best option is to perform a local alignment, with the caveat that some alignment hits will have unreported nucleotides in the query that may still have matched with the subject. Another drawback of global alignments is that, in order to minimize mismatches, gaps are introduced artificially (Chowdhury and Garai, 2017). Therefore, there is a need for computational tools whose alignment results have the strength of a multiple local alignment, and also the refinement provided similar to a global alignment in reporting accurately the number of m/mm for short queries.

Here we present ExtendAlign, a computational algorithm that address this problematic by extending the local alignments reported by an MSAT. ExtendAlign corrects the number of m/mm flanking a seed match in every hit provided by the multiple local alignments, and reports a rectified number of true m/mm in end-to-end fashion considering only the query sequence for all hits resulted from the MSAT without increasing the gap number in order to minimize mismatches. First, we evaluated the functionality of ExtendAlign by comparing its alignment results to local and global alignment tools using simulated data that included short sequence queries (22 nt) and sequence subjects with a wide range of sizes and dissimilarity. Then, we tested ExtendAlign to calculate the number of true m/mm after alignments of short sequences (human mature microRNAs) against slightly longer sequences (mouse pre-microRNAs), and against mouse, armadillo and platypus whole genomes. Our results show that ExtendAlign improves significantly the number of m/mm in all alignments tested, even at a sensitivity of a 7-mer seed size (BLASTn-short). Although ExtendAlign is intended to identify the exact number of m/mm in short sequences (18–25 nt), it executes well with longer sequences and is also capable of improving the m/mm counts for fragments longer than 60 nt. ExtendAlign can be used as a complement for MSATs that perform alignments locally when the highest accuracy for m/mm within a sequence is required, especially in studies where identity percent precision is key for determining homology between phylogenetically distant short sequences. ExtendAlign is available for download at https://github.com/Flores-JassoLab.

RESULTS

Motivation and Functionality of ExtendAlign

Common biases found after the alignment of short sequences are depicted in Figure 1A. For example, consider the default setting for BLAST (Basic Local Alignment Search Tool) and its version for aligning nucleotides (BLASTn) (Altschul *et al.*, 1990). When a short sequence query —i.e., 22 nt, is aligned to a subject with certain similarity, no hits will be reported with the default parameters (word size 11) (i). Adjusting a word size that performs better with short sequences —for example by using the default parameters for BLASTn-short (suitable for sequences shorter than 30 nt), allows the nucleotide alignments as long as the seed (word size, 7-mer) is not interrupted by a gap or mismatch (ii); in this example, BLASTn-short will report a hit with 13 matches out of 14 nucleotides aligned. If the parameters of BLASTn are further adjusted to achieve the highest sensitivity (hereinafter HSe-BLAST) —that is a 7-mer seed size minimizing the penalties for gaps and mismatches (see methods), HSe-BLAST finds an alignment with 15 matches and 2 mismatches for the same pair of sequences (iii). However, the true number of matches and mismatches for an end-to-end alignment after finding a 7-mer seed for this particular hit is 17 and 5, respectively (iv). Even at the highest sensitivity, there are nucleotides beyond the seed size (red, bold letters) still unreported by the MSAT, thus biasing the alignment results; the query in the example for HSe-BLAST is reported as a 17 nt long sequence with 15 matches —instead of 18, and 2 mismatches —instead of 4. This type of bias is particularly common for alignments with dissimilar sequences that contain multiple gaps or mismatches near the 5′ or 3′ ends on the query, and that also fall within a word size that equals or exceeds the unreported sequence length. Although this specific bias may be handled by global alignment tools, such as Needle (EMBOSS) (Suppl. Fig. 1), other intrinsic features of global alignment tools might cause a different type of bias in the reported results: the alignment achieved is reported as the subject length, and alignments will be reported as one hit only at one end if the subject is much larger than the query, making global alignments impractical when searching for multiple alignment hits to a genome.

One consequence of a global alignment is that, in order to find the best alignment, gaps may be introduced in order to minimize mismatches. We performed sequence alignments using a simulated database of randomly generated query sequences of 22 nt long and

aligned them against randomly generated longer subject sequences (49 to 170 nt long); every query was aligned to every subject sequence (see methods). Needle aligned 100% of each sequence for every hit achieved (5,950,500 hits), but also increased the total alignment size up to the longest subject (Suppl. Fig. 2A), which is also accompanied by a similar increase in the number of gaps (Suppl. Fig. 2B). This result implies that Needle: i) considers overhanging nucleotides in the query or subject as part of the alignment, and ii) introduces gaps interspersed along some sequences to achieve alignments. Conversely, HSe-BLAST increased marginally the number of gaps (up to 4), but does not achieve 100% nucleotide alignment coverage for all hits (6,717,736 hits) (Suppl. Fig. 2A and 2B). Importantly, the percentage of nucleotides in each hit that were not aligned by HSe-BLAST should display a complementary representation in the mismatch number (Suppl. Fig. 2C). For example, a 22 nt sequence query with 50% alignment contains 11 nt matches, and also 11 nt mismatches and/or gaps. However, neither the number of mismatches (up to 5, Suppl. Fig. 2C), nor the number of gaps (up to 4, Suppl. Fig. 2B) adds up to 11 nt. This indicates that there are unreported mismatches as a result of the local alignment.

ExtendAlign reduces these report biases by extending the alignment achieved by a local MSAT. If there exist unreported nucleotides flanking the local alignment, ExtendAlign corrects the number of m/mm reported by identifying first the number and identity of any remaining nucleotides in the query that did not align. Finally recalculates and reports the total number of m/mm, in an end-to-end fashion for the query length (Fig. 1B). ExtendAlign corrects the alignments of short-, or long-sequence subjects (for example, whole genomes) whose typical examples of undetected m/mm nucleotides in query searches locate near the 5′ or 3′ ends of queries (Fig. 1C): (i), undetected m/mm; (ii), overhanging query nucleotides are considered mismatches; (iii) and m/mm are extended up to the last query nucleotide at the 5′ or 3′ ends. Results are reported as a table containing the m/mm achieved before and after ExtendAlign for each query hit (Fig. 1C, right). As expected, in the outcome delivered by ExtendAlign there is no increase in the total alignment size, nor a decrease in the number of mismatches, and it always achieves 100% coverage of the sequence query for every hit reported (Suppl. Fig. 2A–C, in red)

To evaluate the performance of ExtendAlign we used human mature microRNAs as short-sequence queries. For short- and long-sequence subjects we used mouse precursor microRNAs and mouse whole genome. The rationale for this is that human microRNAs are used frequently as a reference in homology studies and it also contains the largest list of annotated microRNAs. Conversely, mouse genome contains the second largest list reported of precursor microRNAs, and most of its mature sequences are considered orthologs to human (miRBase release 22) (Kozomara and Griffiths-Jones, 2014). Since BLASTn delivers the highest hit number compared to other sequence alignment tools with similar word size and parameters (i.e., BLAT, Bowtie, Bowtie 2, BWA, Supplementary Table 1), we used BLASTn-short, and HSe-BLAST as local MSATs to initiate the alignments, to provide ExtendAlign with the largest datasets to test.

ExtendAlign improves the alignment of unreported matches and mismatches

First, we initiated the alignment of 2,656 human mature microRNAs versus 1,234 mouse precursor microRNAs using BLASTn-short and its default parameters as MSAT; 2.66 $x10^4$ hits were found. The number of matches per hit achieved was plotted as a cumulative fraction histogram. Then we tested if ExtendAlign could increase the number of matches per hit on nucleotides that may had been unreported by the MSAT. Our results show that ExtendAlign significantly increased the number of matches originally unreported by BLASTn-short ($p < 0.0001$) (Fig. 2A). Next we tested the alignments made by the MSAT using the highest sensitivity parameters. As expected, when alignments were done with HSe-BLAST, the number of hits increased (4.27 $x10^4$ hits) compared to BLASTn-short, and also ExtendAlign detected unreported nucleotides that resulted in an increased number of matches even at the highest sensitivity MSAT alignment (Fig. 2B). Then, to test the performance of Extend Align with long-sequence subjects we aligned all human mature microRNAs against the mouse genome. Our results show that ExtendAlign also increased the number of matches found by BLASTn-short significantly ($p < 0.0001$) (Fig. 2C). When alignments were done with HSe-BLAST, we observed a marginal, but significant increase in the number of matches reported after the ExtendAlign correction ($p < 0.0001$, Kolmogorov-Smirnov $D = 0.005$) (Fig. 2D). The marginal difference achieved by ExtendAlign compared to HSe-BLAST cannot be attributable to the hit number, because it

is larger than BLASTn-short due to its higher sensitivity (2.17 x10^6 *vs*. 1.99 x10^6). One explanation for this result is that, due to its higher sensitivity, HSe-BLAST is able to find a larger number of perfect, or near perfect sequence alignments when the subject is a whole genome. To test this idea, we removed all hits whose alignments found perfect match to the mouse genome, because they are not candidates for ExtendAlign correction. When the data was plotted again, we found a less marginal ($D = 0.69$) and significant difference ($p < 0.0001$) for alignments that were susceptible for correction by ExtendAlign (Fig. 2D, inset). This result supports the purpose of ExtendAlign in correcting the total number of matches contained in an alignment for dissimilar sequences.

Next, we analyzed the number of mismatches for each of the alignments described above (Fig. 3). In contrast to the improvement achieved on the number of matches, the difference in the number of mismatches unreported by the MSAT is more drastic. When the human mature microRNAs were aligned *vs*. mouse precursor microRNAs, ExtendAlign improved significantly both, BLASTn-short ($p < 0.0001$) and HSe-BLAST ($p < 0.0001$) (Fig. 3A and 3B, respectively). The large difference encountered in the mismatches reported by ExtendAlign, compared to matches, is due to the large amount of nucleotides that go unreported by the MSAT: all hits containing query overhangs at their 5′ or 3′ ends with unreported mismatches vastly exceed the number hits corresponding to matches (Fig. 1C, examples ii and iii; and Suppl. Table 2). When alignments were done against the mouse genome, ExtendAlign also increased the number of matches found by BLASTn-short significantly ($p < 0.0001$) (Fig. 2C). However, similarly to what occurred for unreported matches (Fig. 3D), the number of mismatches corrected by ExtendAlign had a significant, but marginal difference with respect to HSe-BLAST alone ($p < 0.0001$, Kolmogorov-Smirnov $D = 0.016$). Likewise, by removing the hits that found a perfect match, a less marginal and significant difference is observed (Kolmogorov-Smirnov $D = 0.89$, $p < 0.0001$) (Fig. 2D, inset). Even at the highest sensitivity, HSe-BLAST missed more than 2% of all mismatches corrected by ExtendAlign (Suppl. Table 2); this percentage increased considerably when BLASTn-short was used for the alignments, and also when precursor microRNAs served as sequence subjects. Taken together these results, we conclude that ExtendAlign is a reliable algorithm that can be used to refine the initial alignment and minimize the biases caused by the unreported m/mm by the use of local MSATs alone.

In phylogenetic inference studies, genome sequencing has become important for establishing (or re-establishing) taxa in which is difficult to determine the genetic boundaries that comprise a taxonomic division; since a key determinant in homology searches is the identity at nucleotide level, having a precise number of m/mm between two contrasted sequences is of utmost importance. Some species are considered to be at a taxonomic border because, at genetic level, many of their genes contain dissimilar sequences and it is difficult to assign the correct taxa. Examples of this are the clades of placental mammals (Afrotheria, Xenarthra, Boreoeutheria), whose sequence relationships are at the border of phylogenetic classification (Tarver *et al.*, 2016). Since ExtendAlign performed robustly in correcting number of m/mm between more dissimilar sequence alignments, we used the list of human precursor microRNAs as sequence queries against distant genomes among mammals in order to find candidates with dissimilar hit matches for correction. First, we used the genome of *Dasypus novemcinctus* (armadillo) because these animals have been considered to be at the border of taxonomic classification. To challenge and compare the results provided by ExtendAlign we used HSe-BLAST as MSAT. Our results revealed a significant increase in the number of matches per hit detected by after the ExtendAlign correction (Fig. 4A, $p < 0.0001$). The number of mismatches also increased significantly ($p < 0.0001$), albeit at much larger extent (Fig. 4B; $D = 0.91$, mismatches *vs*. $D = 0.43$, matches).

Finally, we aligned the human precursor microRNAs against the genome of *Ornithorhynchus anatinus* (platypus). Platypus diverged from the mammalian lineages around 160 million years ago and share morphological and genomic features that represent a mix of mammals, reptiles, birds and fish (Warren *et al.*, 2008). Therefore, the platypus genome provides a unique genomic blend as an attractive source for dissimilar alignments that could serve as candidate hits for ExtendAlign. We observed a significant increase in the number of matches and mismatches per hit, compared with the alignments achieved by HSe-BLAST alone ($p < 0.0001$; Fig. 4C and 4D, respectively). These results suggest that phylogenetic analyses executed using MSATs only may contain underrepresented alignments for all those hit matches for sequences with unreported nucleotides; suggesting that ExtendAlign might be a useful tool to potentially improve cladistics and reconstruct evolutionary trees based on a refined alignment.

Finding a correct sequence identity is not only important for phylogenetic analyses, but also to predict secondary structures for nucleic acids and proteins. In order to do so, multiple sequence alignments are performed followed by multiple folding events based on optimized thermodynamic scores and covariance (Bremges *et al.*, 2010). However, a conserved secondary structure prediction is highly limited for alignments considered within the so-called "twilight zone" —the sequence alignment that ranges from ~35–50% similarity, because the degree of conservation cannot be estimated reliably (Bremges *et al.*, 2010; Rost, 1999; Wilm *et al.*, 2006). Notably, ~50% of all alignment hits achieved by HSe-BLAST in the armadillo and platypus genomes span across this range (Fig. 5A and 5B, HSe-BLAST; medians 43% and 41%, respectively). After ExtendAlign correction, the alignments within this zone improved above 50% of matches per hit (Fig. 5A and 5B, ExtendAlign; medians 44% and 45%, respectively; $p < 0.0001$). These results support the use of ExtendAlign as a regular procedure after alignments performed by local MSATs aimed to find an accurate identity percent, and reinforce the notion that unreported nucleotides, especially mismatches, may lead to inaccurate short sequence annotations.

DISCUSSION

There are several computational tools that serve as sequence aligners. Most of them perform robustly in identifying closed alignments with a high identity percent (>90%) for long sequences. However, as the size and similarity between two sequences declines, there is a bias originated by the word size set by the MSAT that yields unreported nucleotides at the query's ends that affects the number of total m/mm for certain types of alignments. This bias is more acute for short-sequence queries (<30 nt). Here, we present ExtendAlign —a computational tool that minimizes this bias and deliver the true number of m/mm, in end-to-end query hits contained in multiple sequence alignment. ExtendAlign was developed to incorporate the basic output of a MSAT and provide a list of all hits with their corrected m/mm for the total length of the query sequence. In this work we employed BLASTn because it provides the highest number of hits compared to other MSATs, but we envision it could be adapted to also accept the output from other MSATs. Since full-length, perfect match alignments are not prone for improvement, ExtendAlign is more robust in correcting the number of true m/mm in more dissimilar alignments; a useful feature for sequence alignments between more evolutionarily distant species. We showed that ExtendAlign refines the alignment of more dissimilar sequences, particularly those alignments situated within the ~30-50% of matching nucleotides, also called the twilight zone (Ponting, 2017). Examples of alignments that ExtendAlign may help to improve because of their similarity percent span across this range include the conservation and functionality of lncRNAs (which subject to weak functional constraint and rapid turnover during evolution) (Kapusta and Feschotte, 2014), or also a recent controversy rooted in the similarity found in conserved and non-conserved microRNAs from bovine origin in human and murine cell cultures due to RNA transfer caused by fetal bovine serum into cell lines (Wei *et al.*, 2016). We are currently using ExtendAlign to re-evaluate the extent at which this phenomenon might influence the published literature in the microRNA field (Flores-Torres and Flores-Jasso; unpublished data).

Although ExtendAlign was initially developed to satisfy the need for multiple local alignments in end-to-end manner of short-sequence queries, it executes robustly with longer sequences, like precursor microRNAs, which are frequent targets for the discovery of novel microRNAs. Moreover, sequence alignments whose queries/subjects are biased because of one or more internal nucleotides, for example SNPs, could also be corrected

efficiently after ExtendAlign analysis. Also, we envision the use of ExtendAlign to pinpoint the alignment results in all hits provided by MSATs in pathogenic or complex samples (*i.e*. host-pathogen metagenomics) where analyses commonly employ the "best N-hit match"; ExtendAlign can be used to reevaluate an alignment data set containing all hits and then extracting the hits with an identity percent of interest (Shah *et al.*, 2018). The sequencing era has given us the opportunity to analyze several species at genomic level; the prediction of conserved and non-conserved microRNAs among species has become a field of great interest in the last decade. Thus, novel tools like ExtendAlign can help in reducing false-positives commonly found in homology searches for small RNAs and also to increase the refinement of alignments between dissimilar sequences (Brudno *et al.*, 2003; Polyanovsky *et al.*, 2011).

METHODS

Datasets

All mature and precursor microRNAs used as queries and subjects for ExtendAlign analysis were downloaded from miRBase (release 22). Mouse (*Mus musculus*, mm10p6), armadillo (*Dasypus novemcinctus*, v.3.0), and platypus (*Ornithorhynchus anatinus*, v.5.0.1) genomes used as long-sequence subjects were downloaded from NCBI.

Simulations were performed by building 11,901 random 22 nt long query sequences following an uniform probability distribution for each nucleotide. A no variable seed of 7 nucleotides was fixed for every sequence. The start position for this constant region was chosen at random. To generate the subjects, randomly generated nucleotides were added based on the 11,901 query sequences. These extra nucleotides flanked a fixed 7mer position in each query (the seed) to yield a total size that ranged from 49 to 170 nt long. The start position for this constant region was chosen at random. One or up to fifteen nucleotide changes were introduced into the region corresponding to each query sequence. One insertion or deletion (indel) was introduced per every nine mismatches. The length of the indels followed an exponential distribution with lambda=2. The files containing the simulated database were analyzed using HSe-BLAST, ExtendAlign, as described below, and Needle (EMBOSS) with default parameters for command line application (gapextension, 0.5; gapopen, 10).

BLASTn Alignments

The parameters of command line BLASTn were adjusted to achieve the highest sensitivity (HSe-BLAST) as follows: word size, 7; reward, 1; penalty, -1; gap open, 2; gap extend, 2; e-value, 10. In order to keep all query sequences even when scored as low complexity, the DUST filtering and soft masking options were disabled. For constructing the lookup table used by BLASTn, the BLAST extensions were performed without soft masking. BLASTn databases for subject sequences were constructed using the makeblastdb command-line tool including the parse_seqids option (in order to keep the original sequence identifiers) and the dbtype nucl (specific for input nucleotide sequences) parameter. BLASTn-short

was used with default parameters for the command line application. BLASTn output files were adopted to serve as ExtendAlign input.

In the simulation analysis, in order to keep all alignments regardless of their score, the "eValue" and "max_target_seqs" parameters were change to 5000 and 500 respectively.

ExtendAlign Implementation

The ExtendAlign algorithm was implemented in four stages:

1) It starts by performing a BLASTn alignment (either BLASTn-short or HSe-BLAST for short-sequence queries (sQ) and short-, long-sequence subjects (sS and lS, respectively).

2) From HSe-BLAST results, ExtendAlign determines if the reported MSAT´s query start/end positions ($_{MSAT}Q$) correspond to the full-length query ($_{FL}Q$). ExtendAlign identifies whether there exist unaligned nucleotides on the queries (extensions) for every hit reported by the BLASTn search according to its reported start/end positions.

3) ExtendAlign calculates first the length of extensions at each end: 5′ extension is calculated by simply subtracting 1 to the $_{MSAT}Q$ start end ($_{MSAT}Q$ –1) —for example, if 5′-$_{MSAT}Q$ = 4, there is a 3 nt 5′ extension; 3′ extension is calculated by subtracting the $_{FL}Q$ minus the $_{MSAT}Q$ end position —for example, if 3′-$_{MSAT}Q$ = 16 and $_{FL}Q$ = 21, there is a 5 nt 3′ extension. If extensions do not equal zero at either end, ExtendAlign fixes the $_{MSAT}Q$ coordinates and then uses Samtools to extract every sequence extension.

4) Finally, every nucleotide in the extension is aligned directly $_{MSAT}Q$ to $_{FL}Q$ to obtain the correct number of m/mm. As verification step, $m_{(total)} + mm_{(total)} + gap_{(total)} = {}_{FL}Q$, for every hit. In order to avoid mistakes while dealing with overhanging nucleotides, ExtendAlign inverts the extensions and performs the procedure described above. All extensions with true overhangs after ExtendAlign are considered mismatches. In sQ *vs*. sS, ExtendAlign allows BLASTn searches only for sense alignments; in sQ *vs*. lS, BLASTn searches are allowed for sense and antisense alignments. The output file of ExtendAlign contains the query name, subject name, $_{MSAT}Q$, $_{FL}Q$ and total match and mismatch values for MSAT and ExtendAlign.

Statistics

Exact data (no bins) were plotted as cumulative frequency distributions. Significance between datasets was obtained by using the Kolmogorov-Smirnov test for non-parametric data with a significance value of $p < 0.005$. Non-linear fit curve by least squares is shown for every dataset; coefficient of determination as measure of goodness of fit. Box plots are presented as min to max values shown by quartiles; significance was obtained by using the Wilcoxon test for paired, nonparametric data with a significance value of $p < 0.05$. Except for simulated data (plotted with R software), all plots were done using the GraphPad Prism software (v.7).

AVAILABILITY AND REQUIREMENTS

ExtendAlign and all datasets related to this work are available at:

https://github.com/Flores-JassoLab.

Working requirements: all input files must be FASTA format; BLASTn , v. 2.2.31+; Samtools, v. 1.5 (Li *et al.*, 2009); EMBOSS v. 6.6.0.0; Execline, v. 2.5.0.1, and 9base/plan9port for Unix.

AUTHOR CONTRIBUTIONS

M.FT., J.I.HH., and I.AO., constructed EA and set GitHub repository; M.FT., L.GR, H.T., S.E.AV. and C.F.FJ., validated EA; M.FT., S.E.AV. and C.F.FJ., envisioned the project and wrote the manuscript.

ACKNOWLEDGEMENTS

We thank all members of the Avendaño-Vázquez and Flores-Jasso laboratories for critical comments on the manuscript. We also thank Wayne Matten (NCBI-BLAST crew) for valuable help for implementing the command line HSe-BLAST. This work was supported in part by INMEGEN (08/2017/I-322) and SS/IMSS/ISSSTE-CONACyT (289862) to S.E.A-V; and INMEGEN (05/2017/I-321) to C.F.F-J.

REFERENCES

Allmer,J. (2014) miRNomics: MicroRNA Biology and Computational Analysis. **1107**, 157–175.

Altschul,S.F. *et al.* (1990) Basic local alignment search tool. *J. Mol. Biol.*

Arteaga-Vazquez,M. *et al.* (2006) A Family of MicroRNAs Present in Plants and Animals. *PLANT CELL ONLINE*.

Artzi,S. *et al.* (2008) miRNAminer: A tool for homologous microRNA gene search. *BMC Bioinformatics*, **9**, 1–7.

Bartel,D.P. (2018) Metazoan MicroRNAs. *Cell*.

Bartel,D.P. (2009) MicroRNA Target Recognition and Regulatory Functions. *Cell*, **136**, 215–233.

Bartel,D.P. (2009) MicroRNAs: Target Recognition and Regulatory Functions. *Cell*.

Bremges,A. *et al.* (2010) Fine-tuning structural RNA alignments in the twilight zone. *BMC Bioinformatics*.

Brudno,M. *et al.* (2003) Glocal alignment: Finding rearrangements during alignment. *Bioinformatics*, **19**.

Chen,Y. *et al.* (2015) High speed BLASTN: An accelerated MegaBLAST search tool. *Nucleic Acids Res.*, **43**, 7762–7768.

Chowdhury,B. and Garai,G. (2017) A review on multiple sequence alignment from the perspective of genetic algorithm. *Genomics*.

Corpet,F. (1988) Multiple sequence alignment with hierarchical clustering. *Nucleic Acids Res.*

Edgar,R.C. (2004) MUSCLE: multiple sequence alignment with high accuracy and high throughput. *Nucleic Acids Res*.

Friedländer,M.R. *et al.* (2014) Evidence for the biogenesis of more than 1 , 000 novel human microRNAs. 1–17.

Friedman,R.C. *et al.* (2009) Most mammalian mRNAs are conserved targets of microRNAs. *Genome Res.*, **19**, 92–105.

Griffiths-Jones,S. (2004) The microRNA Registry. *Nucleic Acids Res.*

Guerra-Assunção,J.A. and Enright,A.J. (2010) MapMi: Automated mapping of microRNA loci. *BMC Bioinformatics*.

He,L. and Hannon,G.J. (2004) MicroRNAs: Small RNAs with a big role in gene regulation. *Nat. Rev. Genet.*

Kapusta,A. and Feschotte,C. (2014) Volatile evolution of long noncoding RNA repertoires: Mechanisms and biological implications. *Trends Genet.*, **30**, 439–452.

Kozomara,A. and Griffiths-Jones,S. (2014) MiRBase: Annotating high confidence microRNAs using deep sequencing data. *Nucleic Acids Res.*, **42**, 68–73.

Langmead,B. *et al.* (2009) Ultrafast and memory-efficient alignment of short DNA sequences to the human genome. *Genome Biol.*, **10**.

Lee,R.C. *et al.* (1993) The C. elegans heterochronic gene lin-4 encodes small RNAs with antisense complementarity to lin-14. *Cell*.

Li,H. *et al.* (2009) The Sequence Alignment/Map format and SAMtools. *Bioinformatics*, **25**, 2078–2079.

Li,H. and Durbin,R. (2009) Fast and accurate short read alignment with Burrows-Wheeler transform. *Bioinformatics*, **25**, 1754–1760.

Li,H. and Homer,N. (2010) A survey of sequence alignment algorithms for next-generation sequencing. *Brief. Bioinform.*, **11**, 473–483.

Meunier J, Lemoine F, Soumillon M, Liechti A, Weier M, Guschanski K, Hu H, Khaitovich P,K.H. (2013) Birth and expression evolution of mammalian microRNA genes. *Genome Res.*, **23**, 34–45.

Needleman,S.B. and Wunsch,C.D. (1970) A general method applicable to the search for similarities in the amino acid sequence of two proteins. *J. Mol. Biol.*

Notredame,C. *et al.* (2000) T-coffee: A novel method for fast and accurate multiple sequence alignment. *J. Mol. Biol.*

Pfeffer,S. *et al.* (2005) Identification of microRNAs of the herpesvirus family. *Nat. Methods*.

Polyanovsky,V.O. *et al.* (2011) Comparative analysis of the quality of a global algorithm

and a local algorithm for alignment of two sequences. *Algorithms Mol. Biol.*, **6**, 1–12.

Ponting,C.P. (2017) Biological function in the twilight zone of sequence conservation. *BMC Biol.*, **15**, 1–9.

Rajendiran,A. *et al.* (2018) Computational Approaches and Related Tools to Identify MicroRNAs in a Species: A Bird's Eye View. *Interdiscip. Sci. Comput. Life Sci.*, **10**, 616–635.

Rost,B. (1999) Twilight zone of protein sequence alignments. *Protein Eng.*

Roush,S. and Slack,F.J. (2008) The let-7 family of microRNAs. *Trends Cell Biol.*

Smalheiser,N.R. and Torvik,V.I. (2005) Mammalian microRNAs derived from genomic repeats. *Trends Genet.*

Tanzer,A. and Stadler,P.F. (2004) Molecular evolution of a microRNA cluster. *J. Mol. Biol.*

Tarver,J.E. *et al.* (2016) The interrelationships of placental mammals and the limits of phylogenetic inference. *Genome Biol. Evol.*, **8**, 330–344.

Warren,W.C. *et al.* (2008) Genome analysis of the platypus reveals unique signatures of evolution. *Nature*, **453**, 175–183.

Wei,Z. *et al.* (2016) Fetal Bovine Serum RNA Interferes with the Cell Culture derived Extracellular RNA. *Sci. Rep.*

Wilm,A. *et al.* (2006) An enhanced RNA alignment benchmark for sequence alignment programs. *Algorithms Mol. Biol.*

FIGURE LEGENDS

Figure 1. (*A*) Common biases caused by unreported matches and mismatches and ExtendAlign correction examples. Unreported nucleotides; red, bold letters. (*B*), ExtendAlign (EA) takes the output file of an MSAT after the alignment of long-, or short-sequence subjects and short-sequence queries. (*C*) i–iii, examples of unreported nucleotides comprising matches or mismatches that originate alignment biases. Nucleotides flanking the seed-word size after MSAT alignment (red, bold) are extended and the true number of m/mm is reported by ExtendAlign according to the query length. Typical examples of nucleotides unreported by commonly used MSATs locate near the 5′ or 3′ ends on the queries aligned to long- (*C*, i), or short-sequence (*C*, ii and iii) subjects. The nucleotides that overhang on queries either at their 5′ or 3′ ends (underlined, grey; example ii) are interpreted as mismatches by ExtendAlign; overhanging nucleotides on subjects are not reported.

Figure 2. ExtendAlign corrects the number of matches per hit reported by MSAT in short-sequence queries. (*A* and *B*) Human mature microRNAs *vs*. mouse precursor microRNAs. (*A*) BLASTn-short *vs.* ExtendAlign (EA). (*B*) High Sensitivity BLASTn (HSe-BLAST) *vs.* ExtendAlign. (*C* and *D)* Human mature microRNAs *vs*. mouse genome. (*C*) BLASTn-short. (*D*) HSe-BLAST. BLASTn-short (grey), HSe-BLAST (grey) and EA (red) are plotted as cumulative fraction histogram of the number of matches found in each alignment. Significance values were calculated using the Kolmogorov-Smirnov [K-S] test for unpaired nonparametric data; D = K-S distance.

Figure 3. ExtendAlign corrects the number of mismatches per hit unreported by MSATs in short-sequence queries. (*A* and *B*) Human mature microRNAs *vs*. mouse precursor microRNAs. (*A*) BLASTn-short *vs.* ExtendAlign (EA). (*B*) High Sensitivity BLASTn (HSe-BLAST). (*C* and *D*) Human mature microRNAs *vs*. mouse genome. (*C*) BLASTn-short. (*D*) HSe-BLAST. The number of mismatches per hit achieved by BLASTn-short (grey), HSe-BLAST (grey) and EA (red) are plotted as cumulative fraction histogram. Significance values were calculated using the Kolmogorov-Smirnov [K-S] test for unpaired nonparametric data; D = K-S distance.

Figure 4. ExtendAlign improves the number of matches and mismatches for undetected nucleotides in pre-microRNA as queries vs armadillo and platypus genome. (*A* and *B*) Comparison of HSe-BLAST and ExtendAlign (EA), human precursor microRNAs (query), *vs*. armadillo genome (subject). (*C* and *D*) Comparison between HSe-BLAST and ExtendAlign (EA), human precursor microRNAs (query), *vs*. platypus genome (subject). (*A* and *C*) Detected number of matches; (*B* and *D*) Detected number of mismatches. The number of matches and mismatches per hit achieved by HSe-BLAST (grey) and ExtendAlign (red) are plotted as cumulative fraction histogram. Significance values were calculated using the Kolmogorov-Smirnov [K-S] test for unpaired nonparametric data; *D* = K-S distance.

Figure 5. ExtendAlign improves the alignment percent of dissimilar sequences that span across the twilight zone. (*A* and *B*) Comparison of the identity percent for all matches per hit made by HSe-BLAST and after ExtendAlign correction. (*A*) Human precursor microRNAs vs. armadillo genome. (*B*) Human precursor microRNAs vs. platypus genome. Significance was obtained by Wilcoxon test for paired, non-parametric data.

SUPPLEMENTAL INFORMATION

Supplemental Figure Legends

Supplementary Figure 1. Examples of results obtained with the global alignment tool Needle (EMBOSS). Overhanging nucleotides (underlined letters) beyond the query size are reported as gaps and the alignment length is reported as the subject, (i). If there are overhanging nucleotides on both the query and subject are summed and reported as the total alignment length (ii).

Supplementary Figure 2. ExtendAlign achieves 100% query alignment without increasing the alignment size or introducing gaps. 11,901 randomly generated query and subject sequences (22mers and 49–170, respectively) were aligned using HSe-BLAST, Needle, and ExtendAlign. Percent of query aligned (*A*), Number of gaps per sequence pair (*B*), and Number of mismatches per query (*C*) were plotted as a function of total alignment size for each alignment hit; the frequencies for each intersection were plotted in *z*. Hits per alignment: HSe-BLAST (grey) 6,717,736 hits; Needle (cyan) 5,950,500 hits; ExtendAlign (red) 6,717,736 hits.

Flores-Torres et al., Figure 1

A

(i) 5′-UGGAUUCUUUAGAGUGCAUCUGCGUC-3′
5′-GAUGCUGUAGAGGGCAUCUGCU-3′

BLASTn (word size 11):
No hits
Query size 22 nt

(ii) 5′-UGGAUUCUUUAGAGUGCAUCUGCGUC-3′
||||| ||||||||
5′-**GAUGCUG**UAGAGGGCAUCUGC**U**-3′

BLASTn-short (word size 7):
13 matches, 1 mismatch
Query size reported 14 nt
Unreported nt: 8

(iii) 5′-UGGAUUCUUUAGAGUGCAUCUGCGUC-3′
|| ||||| ||||||||
5′-**GAUG**CUGUAGAGGGCAUCUGC**U**-3′

HSe-BLAST (word size 7):
15 matches, 2 mismatches
Query size reported 17 nt
Unreported nt: 5

(iv) 5′-UGGAUUCUUUAGAGUGCAUCUGCGUC-3′
||| || ||||| ||||||||
5′-GAUGCUGUAGAGGGCAUCUGCU-3′

Total query matches: 18
Total query mismatches: 4

B

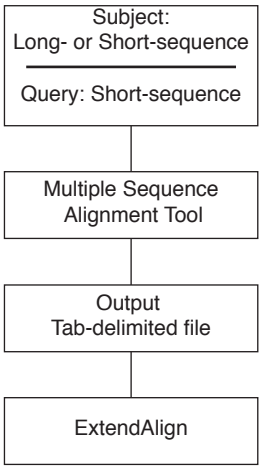

C

	MSAT	MSAT + EA	
(i)	14	17	Matches
	2	4	Mismatches
	16	21	Query length
(ii)	7	9	Matches
	0	7	Mismatches
	7	16	Query length
(iii)	19	21	Matches
	0	5	Mismatches
	19	26	Query length

Flores-Torres et al., Figure 2

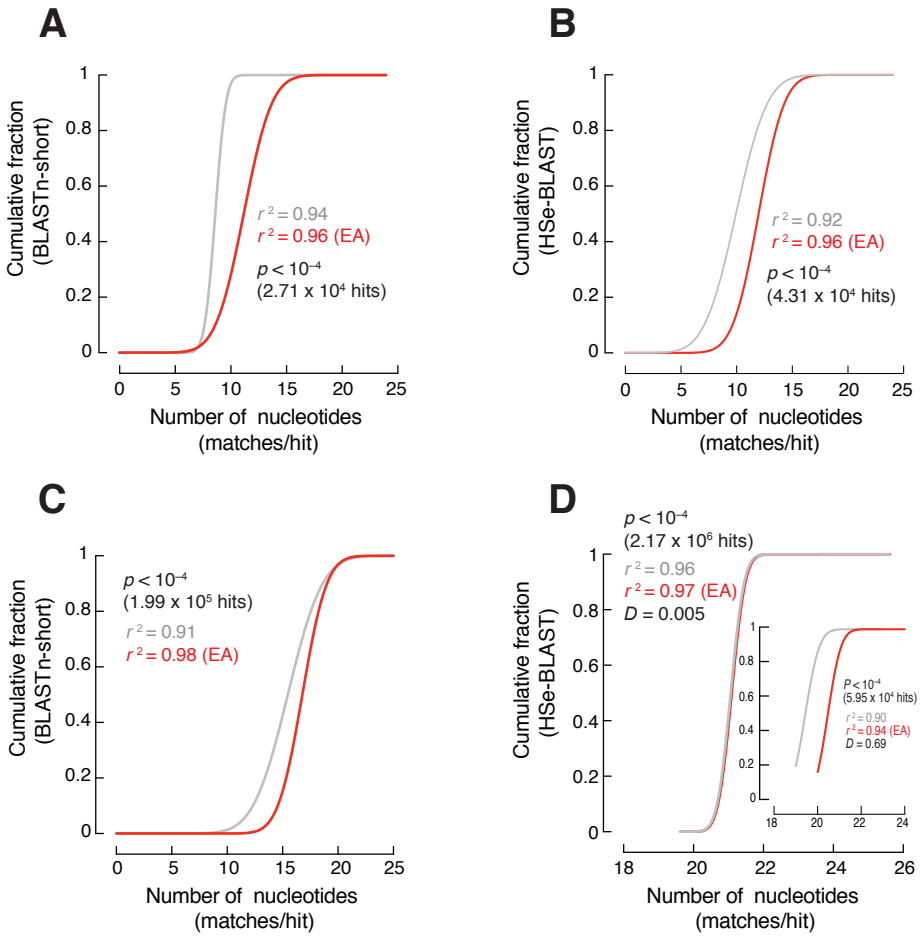

Flores-Torres et al., Figure 3

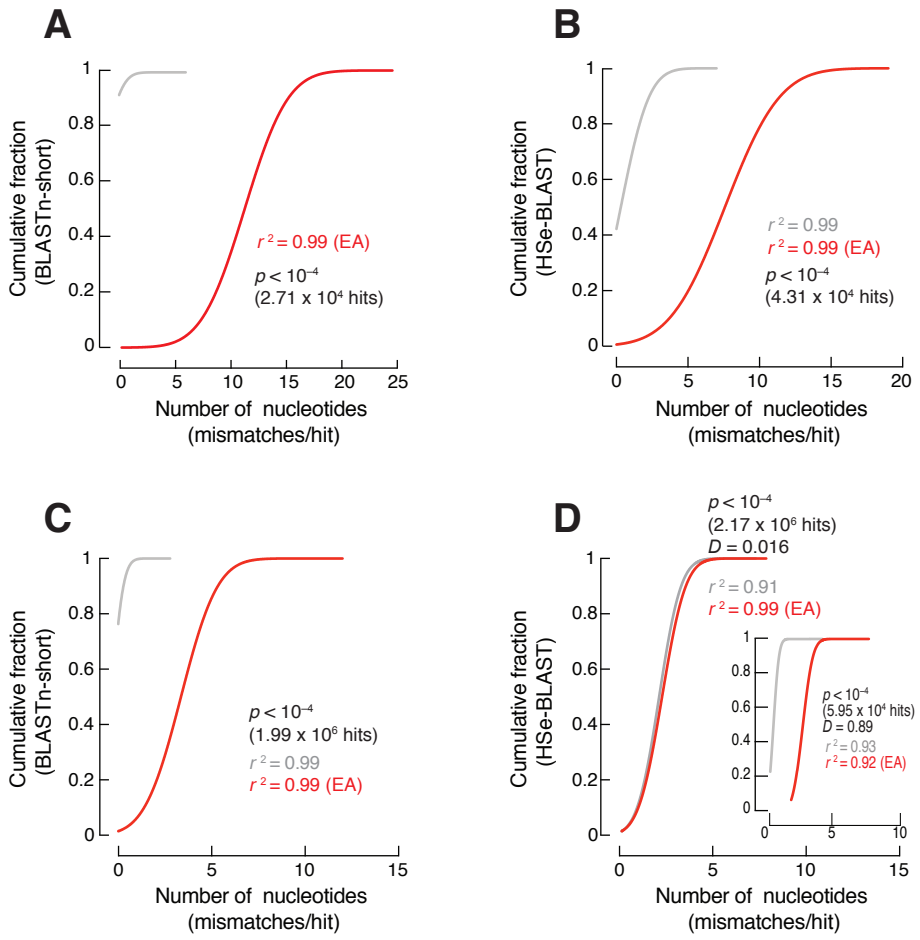

Flores-Torres et al., Figure 4

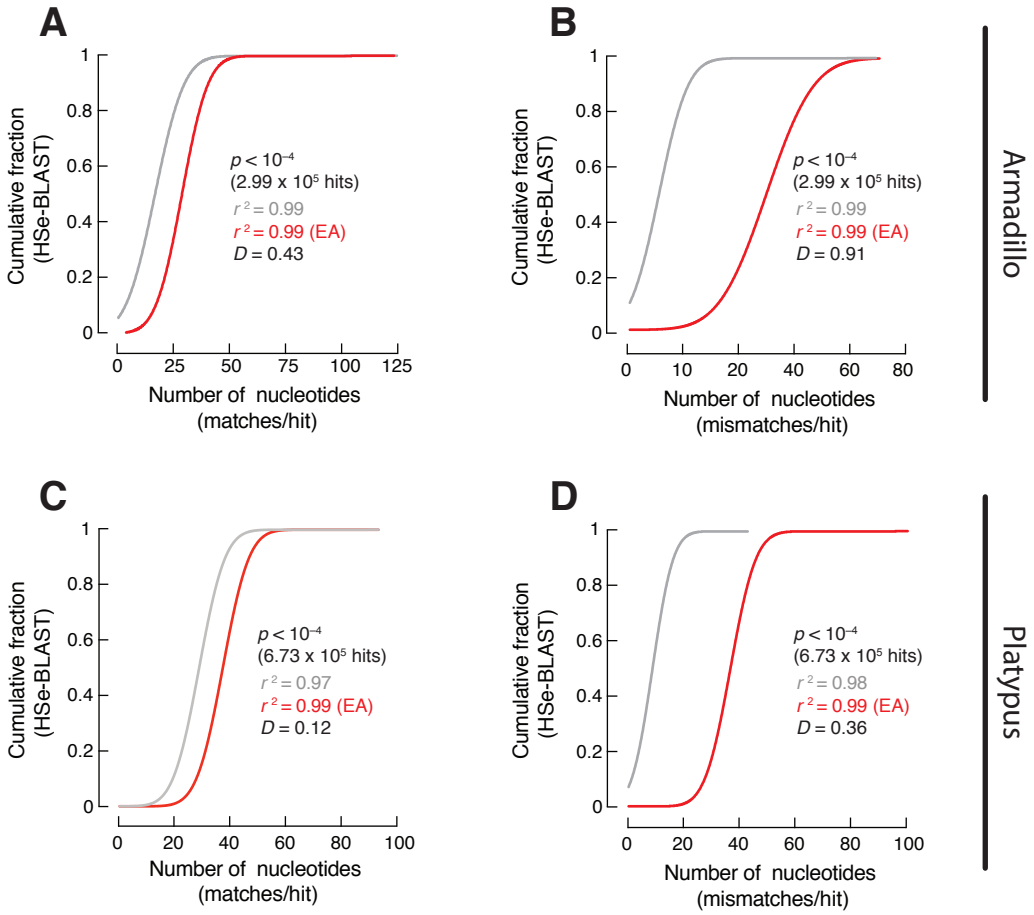

Flores-Torres et al., Figure 5

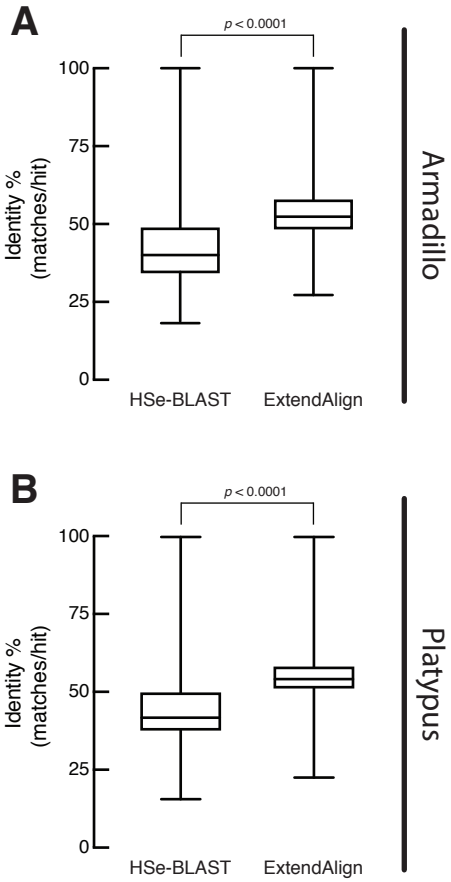

Flores-Torres et al., Supplementary Figure 1

Global Alignment
(Needle—EMBOSS)

(i)

```
5′-UGGAUUCUUUAGAGUGCAUCUGCGUC-3′
   ||| || ||||| |||||||||
5′-GAUGCUGUAGAGGGCAUCUGCU-3′
```

Aligment length: 26
Query length: 22
Total query mismatch/gap: 4

(ii)

```
5′-AGGAGTCGGGAGCGGAGAGGGT-3′
                || |||||||
        5′-CGTAGAGGGTCATGCT-3′
```

Aligment length: 28
Query length: 16
Total mismatch/gap: 18

Flores-Torres et al., Supplementary Figure 2

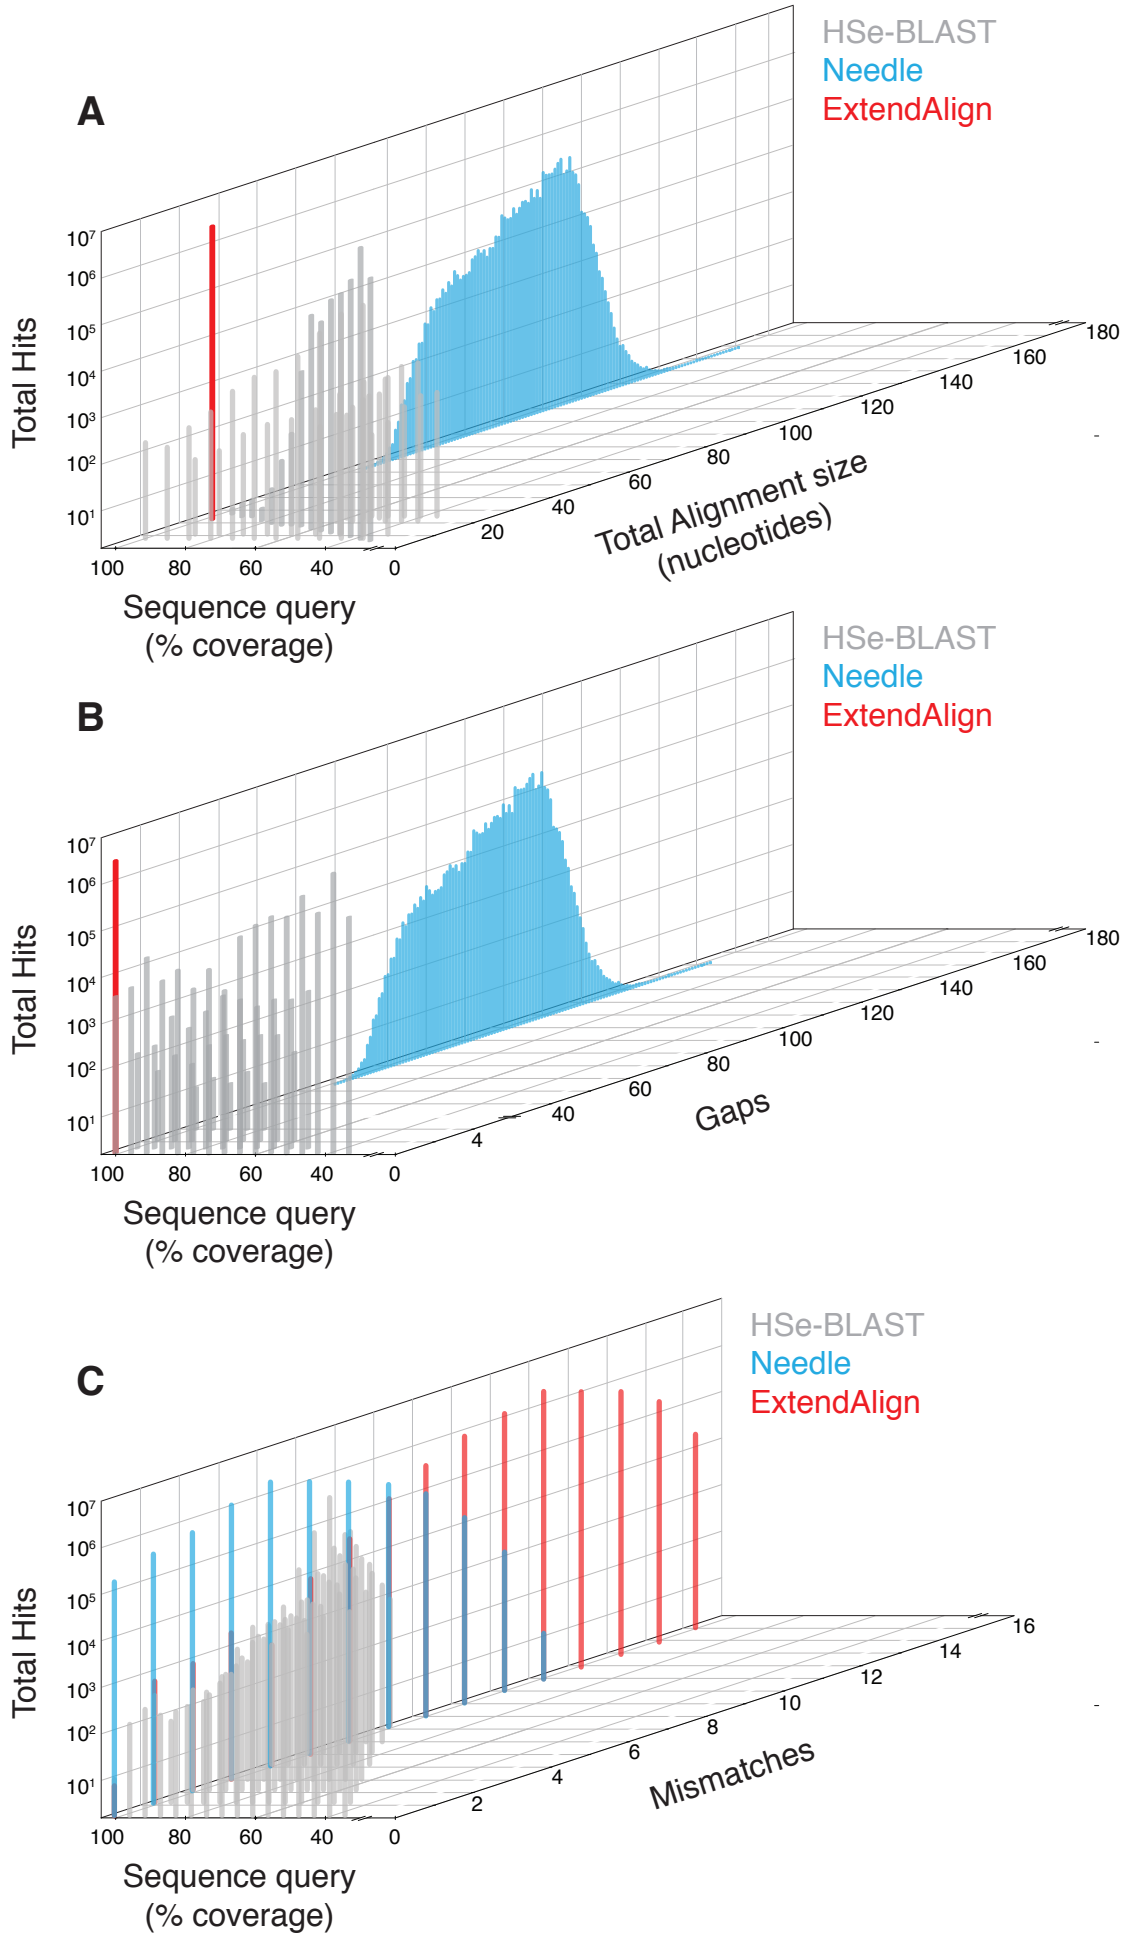

Supplementary Table 1. Total hits by local MSAT

Alignment tool	No. Query sequences	Total Hits	Sequences aligned	Word size
BLAT	2656	1067	626	8
Bowtie	2656	1176	692	5
Bowtie 2	2656	1589	1589	5
BWA	2656	1940	1435	7
BLASTn-short	2656	27108	2628	7
HSe-BLAST	2656	43179	2655	7

Supplementary Table 2. Mismatches corrected by ExtendAlign

	MSAT	MSAT + EA	No. nt-c / f.c.
BLASTn-short (m-miRNA vs pre-miRNA)	1,857	244,263	242406 / 131
HSe-BLAST (m-miRNA vs pre-miRNA)	48952	389,650	340698 / 8
BLASTn-short (m-miRNA vs genome)	47,225	807,014	759789 / 17
HSe-BLAST (m-miRNA vs genome)	4,238,650	4,362,321	123671 / 1

* Human mature microRNA = m-miRNA
Mouse precursor microRNA = pre-miRNA
Mouse genome = genome
Sequence Alignment Tool = SAT
ExtendAlign = EA
No. nt-c/f.c. = number of mismatches corrected / fold-change