

# Navigating the unexplored seascape of pre-miRNA candidates in single-genome approaches

Nuno D. Mendes<sup>1,2,3,7</sup>, Steffen Heyne<sup>4</sup>, Ana T. Freitas<sup>2,\*</sup>, Marie-France Sagot<sup>1,3,\*</sup> and Rolf Backofen<sup>4,5,6,\*</sup>

<sup>1</sup> Équipe BAOBAB, Laboratoire de Biométrie et Biologie Évolutive (UMR 5558); CNRS; Univ. Lyon 1, 43 bd du 11 nov 1918, 69622, Villeurbanne Cedex, France

<sup>2</sup>INESC-ID/IST, 9 Rua Alves Redol, 1000-029 Lisbon, Portugal

<sup>3</sup>BAMBOO Team, INRIA Rhône-Alpes, 655 avenue de l'Europe, 38330 Montbonnot Saint-Martin, France

<sup>4</sup>Department of Computer Science, University of Freiburg, Georges-Köhler-Allee 106, 79110 Freiburg, Germany

<sup>5</sup>Center for Biological Signalling Studies (BIOSS), University of Freiburg, 79110 Freiburg, Germany

<sup>6</sup>Center for non-coding RNA in Technology and Health, Faculty of Life Sciences, University of Copenhagen, Grønnegårdsvej 3, 1879 Frederiksberg C, Denmark

<sup>7</sup>Instituto Gulbenkian de Ciência, Rua da Quinta Grande 6, 2780-156 Oeiras, Portugal

Associate Editor: Prof. Ivo Hofacker

## ABSTRACT

**Motivation:** The computational search for novel miRNA precursors often involves some sort of structural analysis with the aim of identifying which type of structures are prone to being recognised and processed by the cellular miRNA-maturation machinery. A natural way to tackle this problem is to perform clustering over the candidate structures along with known miRNA precursor structures. Mixed clusters allows then the identification of candidates that are similar to known precursors. Given the large number of pre-miRNA candidates that can be identified in single-genome approaches, even after applying several filters for precursor robustness and stability, a conventional structural clustering approach is unfeasible.

**Results:** We propose a method to represent candidate structures in a feature space which summarises key sequence/structure characteristics of each candidate. We demonstrate that proximity in this feature space is related to sequence/structure similarity, and we select candidates which have a high similarity to known precursors. Additional filtering steps are then applied to further reduce the number of candidates to those with greater transcriptional potential. Our method is compared to another single-genome method (TripletSVM) in two datasets, showing better performance in one and comparable performance in the other, for larger training sets. Additionally, we show that our approach allows for a better interpretation of the results.

**Availability and Implementation:** The MinDist method is implemented using Perl scripts and is freely available at

<http://www.cravela.org/?mindist=1>.

**Contact:** backofen@uni-freiburg.de

## 1 INTRODUCTION

MicroRNAs (miRNAs) constitute one of several classes of small RNAs found in plant and animal branches of Eukaryota. Since the

discovery of the first miRNAs in *C. elegans* (Lee *et al.*, 1993), an abundant number of these regulatory RNAs (ranging from 18-25 nt in length) have been discovered and their underlying mechanisms investigated (for an overview see e.g. Bartel (2009)). MiRNAs originate from the maturation of larger precursors of  $\sim 70$  nt called pre-miRNAs.

An important feature of pre-miRNAs that elicits their recognition by the miRNA-processing machinery is their secondary structure. Pre-miRNAs typically exhibit a stem-loop structure with few internal loops or asymmetric bulges but the variety of structures that are efficiently recognised has escaped any strict characterisation (Lindow and Gorodkin, 2007). Previously, we proposed a combination of measures that distinguishes true pre-miRNAs from the large number of stem-loops that can be found in metazoan genomes (Mendes *et al.*, 2010). However, the number of precursor candidates (in the order of a few hundred thousand) obtained above the optimal cutoff of the score, which combines measures of stability and robustness, is still impractically large to be subjected to experimental confirmation (see the CRAVELA framework website for further details: <http://www.cravela.org>). Despite the fact that all these candidates consist of or contain stem-loops, the details of their secondary structure have not been subjected to a thorough analysis.

The most immediate approach to analysing the variety of pre-miRNAs in our candidate set is to seek the identification of structural families amongst the precursor candidates. Although miRNAs have been grouped into families according to their sequence similarity in the miRBase database (Griffiths-Jones *et al.*, 2003), this approach does not give enough insight as to the structural features that are important for the recognition by the miRNA-processing machinery. Hence, the grouping has to be performed according to sequence *and* structure. Various algorithmic approaches have been introduced to determine structural similarities and to derive consensus structure patterns for structural RNAs with low sequence identity (Siebert and

\*to whom correspondence should be addressed



## 2.3 PCA of vectorial representations

The vectorial representation of the stem-loop structures used in this work captures information about sequence/structure features but, in general, the dimensions of these feature vectors are not independent. Furthermore, all vectors will always have zero values in some dimension as some combinations of left/right-hand paired and unpaired nucleotides are not possible in actual RNA structures. In order to reduce the number of dimensions and to ensure we represent our structures in a space with independent dimensions, we readily eliminate dimensions with zero variance. In practice, for a sufficiently large dataset, this will only eliminate dimensions for which all vectors have a value of zero. Over the remaining dimensions, we perform a principal component analysis (PCA) thus obtaining a space of linearly independent coordinates. Additionally, each dimension of the vectorial representation is scaled to unit variance before performing the PCA.

## 2.4 Evaluation

**2.4.1 Randomisation procedure** A randomisation procedure is used to compare our approach to the results obtained for samples of the datasets using conventional structural clustering in terms of the proportion of correctly assigned cluster members. This procedure allows us to estimate the likelihood that our values were obtained by chance or as a result of the way our candidates are spatially distributed in the principal components space. To obtain the proportion of correctly assigned cluster members in the randomised version of the samples, we keep each candidate in the same position of the principal components space but we shuffle their identities, *i.e.*, we randomly select two candidates and we swap their coordinates, repeating the process until all candidates have had their coordinates swapped. After having performed the random swapping of candidates, we calculate the centroids of each cluster in the shuffled space and the resulting proportion of correctly assigned cluster members. A similar procedure is used to obtain random samples of the median distance of known precursors to their centroid.

**2.4.2 Sampling the datasets for performance evaluation** To assess the performance of both our method and TripletSVM, four groups of samples of different sizes were prepared for each dataset. Each sample group was divided in training sets and testing sets, both with the same number of positive and negative examples. Each sample group is made of 1000 samples. The positive examples in the training set of each sample are a random subset of a pool of known miRNAs, with a single representative per miRNA family (either 5%, 10%, 20% or 50% of the pool) and the remaining pre-miRNAs in the pool make up the positive examples of the corresponding testing set. The negative examples in both the training and testing sets of each sample are random subsets of the candidates having the same size of the corresponding positive examples. Our method uses only the positive examples in the training set as a reference from which to compute the distance to the elements in the testing set, whereas TripletSVM, for each sample, is trained using both the positive and negative examples of the training set and is evaluated against the testing set.

## 3 DISCUSSION

### 3.1 Vector representation reflects structural clustering

To assess the adequacy of our approach with respect to its ability to identify regions of structural similarity in a way that resembles conventional sequence/structure clustering, we adopt the following procedure. We use LOCARNA to perform hierarchical structural clustering over 100 samples of 1000 randomly chosen stem-loops drawn from the *D. melanogaster* and *A. gambiae* datasets, always including the entire set of known miRNAs for each organism, and we determine the optimal partition into clusters applying a tree node evaluation rule for various significance levels called *k*-levels (the

<i>k</i> -level	<i>A. gambiae</i>			<i>D. melanogaster</i>		
	Corr. assign.	<i>p</i> -value	Avg. cl. size	Corr. assign.	<i>p</i> -value	Avg. cl. size
0.00	83.60%	8.58e-87	3.05	82.10%	3.45e-125	3.29
0.10	82.50%	1.69e-87	3.33	81.24%	1.92e-120	3.54
0.20	81.21%	8.68e-84	3.70	80.01%	4.31e-113	3.89
0.30	79.30%	2.37e-76	4.27	78.24%	3.31e-108	4.44
0.40	77.09%	9.22e-65	5.31	76.08%	1.64e-96	5.44
0.50	74.12%	2.24e-55	7.61	72.80%	1.43e-84	7.54
0.60	71.23%	2.76e-42	12.09	69.45%	1.01e-61	11.44
0.70	68.70%	1.05e-31	17.24	67.41%	1.74e-54	15.32
0.80	68.14%	6.01e-26	19.52	66.07%	9.62e-44	17.77
0.90	67.37%	2.57e-22	21.08	65.72%	4.01e-37	20.09

**Table 1.** Evaluation of vectorial representations. For each *k*-level, the table shows the percentage of correct assignments in the datasets of *A. gambiae* and *D. melanogaster*, the *p*-value of Welch's two-sample *t*-test comparing the observed correct assignments with a randomised version of each dataset shuffling the correspondence between candidates and their vectorial representation, and the average number of cluster members.

details of this procedure are described in the Methods section). For low values of *k*, the procedure produces clusters with highly similar structures. An increasing value of *k* allows for more dissimilar structures to be included in the same cluster, therefore producing a lower number of clusters with an increasing number of structures.

We then represent each structure from the samples using a vectorial representation summarising sequence/structural features, in an effort to capture the key elements distinguishing the various hairpins. These feature vectors contain, however, both interdependent dimensions and dimensions with different variance. To obtain a linearly independent set of dimensions, we perform a principal components analysis (PCA) over the vectorial representations mapping them to their principal components representation (feature space).

To determine whether our representation of the candidates in the principal components space reflects the structural clusters found by LOCARNA for the different *k*-levels, we calculate the proportion of correct assignments, which measures the ratio of cluster members that are closer to their assigned cluster centroid as opposed to a centroid of another cluster. The cluster centroid is calculated by determining the average position of the cluster members each dimension at a time. The distribution of this measure in our 100 samples is then compared to its distribution in a randomised version of our spatial representation of the candidates, where candidate positions are kept but candidate identities are shuffled.

The comparison with the randomised version of the spatial distribution of candidates allows us to address the fact that the significance of the absolute distances of each candidate to its cluster centroid in the feature space can only be determined comparatively. For instance, some clusters may have only one member, in which case it will always coincide with its cluster centroid and, more generally, it may happen that variations in distance of a candidate to its assigned cluster centroid for different *k*-levels, or even different vectorial representations, may partially reflect the overall density of the candidates rather than a better evaluation of structural similarity.

Table 1 shows that, for both datasets, a large proportion of cluster members are found closer to their cluster centroid than to the centroid of any other cluster. For the most heterogeneous clusters which are obtained at *k*-level 0.90 the proportion of correctly assigned cluster members is about two-thirds, and it rises above 80% for the structurally more homogeneous clusters obtained at *k*-level 0.00. The comparison with the randomised datasets shows that the results are statistically significant, *i.e.*, these results are well above

what one would hope to obtain by chance or simply due to the way candidates are spatially distributed.

### 3.2 Known precursors are clustered together

Using the same samples from the datasets presented in the previous section, we can observe that despite the fact that not all known precursors are grouped together in the same cluster by LocARNA at any  $k$ -level (data not shown), they are however significantly close and restricted to a limited portion of the features space. In fact, if we take the centroid of the known precursors and calculate the median distance of each known precursor to the centroid, we obtain a value which is much smaller than what would be expected by chance ( $p$ -value =  $8.00 \times 10^{-60}$  for *A. gambiae*,  $p$ -value =  $8.94 \times 10^{-71}$  for *D. melanogaster*, estimated using the randomisation procedure described in section 2.4.1).

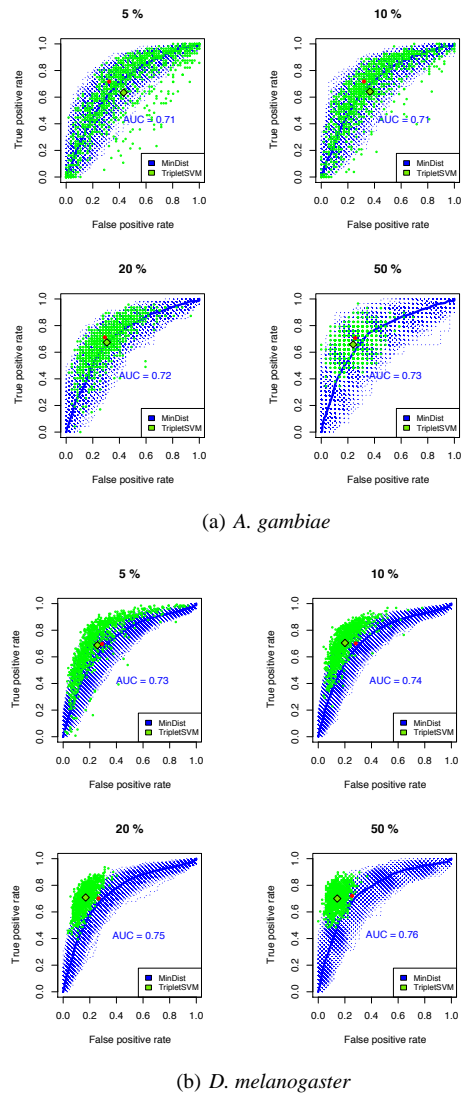
### 3.3 Distance to known precursors is good predictor

The results shown in the previous section suggest that known precursors tend to concentrate in a particular region of the feature space. This region, however, is also densely populated by other precursor candidates. Since the region where known precursors are found is inserted in an area of great density, it cannot be identified by a purely unsupervised approach. Therefore we take the coordinates of all known precursors and use them to identify the closest candidates. This method has the advantage of allowing for different pre-miRNA structural clusters to emerge around subsets of known precursors. The number of candidates that are included in the acceptance region is controlled by the maximum distance allowed to the closest pre-miRNA.

The larger the permitted distance, the greater the chance of selecting a region that includes all interesting candidates, but at the expense of enlarging the number of false positives. The Youden index (Youden, 1950),  $J$ , defined as  $\max_c \{TPR(c) - FPR(c)\}$ , i.e., the maximum difference between the true positive rate (TPR) and the false positive rate (FPR) over all cutoff values,  $c$ , is a standard method to select the best compromise in such a trade-off. The optimal cutoff value,  $c^*$ , is the cutoff for which  $J = TPR(c^*) - FPR(c^*)$ .

To estimate the optimal cutoff, we consider subsets of known precursors as reference and calculate the true/false positive rate with respect to the remaining known precursors and other candidates (see detailed description in the Methods section). Fig. 2 shows the ROC curves for *A. gambiae* and *D. melanogaster* when using samples of 5%, 10%, 20%, and 50% of known precursors as reference and computing the trade-off between the true/false positive rates with respect to the remaining precursors and an equal number of sampled candidates. Each figure shows the ROC curves of 1000 such samples as well as the average curve, computed as the average performance over all samples across the full range of cutoff values. Additionally, the figures also show the average performance of our method, computed as the average TPR and FPR across all samples for the optimal cutoff on each sample (note that this may be significantly different from the optimal cutoff calculated on the average ROC curve).

The optimal cutoff in each of these ROC curves can be interpreted as the best choice of maximum distance allowed between a structure and the closest precursor so that the former may be included in the acceptance region. We have observed that there is a log-linear



**Fig. 2.** ROC curves for the minimum distance (MinDist) to pre-miRNAs method and the performance of TripleTSVM over 4000 samples equally divided into 4 groups. Each group uses 5%, 10%, 20% or 50% of the known precursors of (a) *A. gambiae* and (b) *D. melanogaster* to set up the positive examples of the training set. The positive examples of the testing set are made up by the remaining precursors. In both sets, the negative examples are samples of the set of candidates. ROC curves for each individual sample are shown in dashed lines and the average curve across the range of cutoff values is shown in a solid line. The red dot represents the average performance of the MinDist method over all samples considering the optimal cutoff for each sample. The green dots represent the performance of TripleTSVM on each sample, whereas the green diamond refers to its average performance

relation between the value of the average optimal cutoff and the percentage of known precursors that is used as reference ( $R^2 = 0.998$ , for *A. gambiae*, and  $R^2 = 0.989$ , for *D. melanogaster*). Since the best choice of cutoff cannot be directly determined for the entire set of known precursors, we estimate it by extrapolating the log-linear model. The estimated optimal cutoff can be interpreted as the best choice of maximum distance to include additional (yet unknown) precursors with the least number of false positives.



% known	<i>A. gambiae</i>						<i>D. melanogaster</i>					
	MinDist		TripletSVM				MinDist		TripletSVM			
	Sens.	Spec.	$F_1$	Sens.	Spec.	$F_1$	Sens.	Spec.	$F_1$	Sens.	Spec.	$F_1$
5 %	0.72	0.68	0.71	0.63	0.57	0.61	0.70	0.71	0.70	0.69	0.74	0.71
10 %	0.72	0.68	0.71	0.64	0.73	0.67	0.70	0.72	0.71	0.70	0.80	0.74
20 %	0.71	0.71	0.71	0.67	0.69	0.68	0.70	0.73	0.71	0.71	0.83	0.76
50 %	0.71	0.74	0.72	0.66	0.76	0.69	0.72	0.75	0.73	0.70	0.86	0.76
80 %	0.70	0.80	0.74	0.64	0.78	0.69	0.75	0.75	0.75	0.69	0.86	0.75
90 %	0.74	0.82	0.77	0.63	0.78	0.68	0.77	0.78	0.77	0.68	0.87	0.75
95 %	0.83	0.79	0.81	0.64	0.78	0.69	0.79	0.81	0.80	0.68	0.87	0.75

**Table 2.** Sensitivity (TPR), Specificity (1 - FPR) and the  $F_1$  measure ( $2 \frac{TP/(TP+FP) * TPR}{TP/(TP+FP) + TPR}$ ) of TripletSVM and MinDist computed as the average performance across all samples for training sets whose positive examples consist of a fraction of known pre-miRNAs in

#### *A. gambiae* and *D. melanogaster*

Using the estimated optimal cutoffs, the selected regions include 23.5% (77 366) and 23.5% (67 619) of all candidates from the *A. gambiae* and *D. melanogaster* datasets, respectively.

### 3.4 Comparison to other methods

TripletSVM (Xue *et al.*, 2005) is a classifier based on a support vector machine (SVM) that purports to determine whether a given stem-loop is a pre-miRNA. The features considered by this SVM are quite similar to those of the TRIPLET\_S vectorial representation that is described in the Supplementary Materials. It is also, to our knowledge, the only single-genome method whose source code is made available and which includes the necessary routines to re-train the model with new data. Many other single-genome methods exist (Mendes *et al.*, 2009), but their time complexity makes their use in the classification of hundreds of thousands of candidates unfeasible. TripletSVM was trained using positive examples from samples of known precursors and negative examples from samples from the candidate set (a detailed description is given in the Methods section). A graphical representation of the performance of TripletSVM in each of the 4000 samples (evenly distributed between training sets using 5%, 10%, 20%, and 50% of the annotated pre-miRNAs), as well as the average performance in each group of samples, is shown in Fig. 2. Table 2 shows the sensitivity, specificity and the  $F_1$  measure for TripletSVM as well as our method across training sets including varying proportions (from 5 to 95%) of known precursors.

The average performance of our method in *A. gambiae* is superior to that of TripletSVM, and slightly worse in *D. melanogaster*, except for sample groups containing a greater number of known precursors. The slightly better performance of TripletSVM in the *D. melanogaster* dataset is the result of a tendency for having comparatively higher specificity but similar sensitivity. This is probably due to the fact that MinDist is sensitive to the inclusion of heterochromatic sequences in this dataset, which introduces greater variability in terms of sequence/structure features. As a consequence, the variation between the features of pre-miRNAs and those of other stem-loops with more regular features becomes less apparent. It is worth noting, however, that despite achieving a slightly better average performance in *D. melanogaster* for sample groups containing a lower number of known precursors, the actual performance of TripletSVM in these samples varies greatly from one sample to the other, alternating between very good and very poor performances, particularly for the 5% sample group, as is shown in Fig. 2 (b), which is a major disadvantage when exploring recently sequenced genomes, for which few clear homologs with

which to seed the search for new pre-miRNAs are generally available.

TripletSVM also bears the inconvenience of requiring negative examples which are inevitably chosen under the assumption – however plausible and defensible – that miRNA precursors are rare with respect to the overall number of candidates, but one cannot generally guarantee that hairpins which would normally be processed by the miRNA-maturation pathway are not being included in the negative training set. Our approach, despite assuming all candidates to be false positives for the purposes of performance evaluation, does not use this information to shape the acceptance region and, since it does not try to identify the optimal margin between positive and negative examples, it is also less likely to suffer from overtraining. Additionally, by reflecting sequence/structure similarity in a way comparable to conventional structural clustering, our method offers a better interpretation of the decision rule that is made when selecting candidates.

### 3.5 Transcriptional potential assessment further restricts the number of candidates

The number of candidates obtained after our structural analysis (77 366 for *A. gambiae*, and 67 619 for *D. melanogaster*), albeit considerably lower than the original candidate set (328 829 for *A. gambiae*, and 287 469 for *D. melanogaster*), is still quite numerous. A plausible interpretation of these results is that, despite their structural similarity to known precursors, the majority of these candidates are not pre-miRNAs due to other factors. Chiefly among these is the fact that most remaining candidates are probably not efficiently transcribed or are playing different biological roles. This illustrates the often ignored distinction between having an adequate secondary structure and actually being transcribed and processed.

A straightforward way to address the need to assess the transcriptional potential of the candidates is the observation that many fall within regions that have been annotated. Genomic locations with no annotation or which have been annotated as introns may contain miRNA precursors, but candidates which overlap regions annotated as exons, transposons or other non-coding RNAs are less likely to contain pre-miRNAs. If we filter out non-viable candidates by this criterion, our candidate set is reduced to 44 210 for *A. gambiae* and 40 582 for *D. melanogaster*.

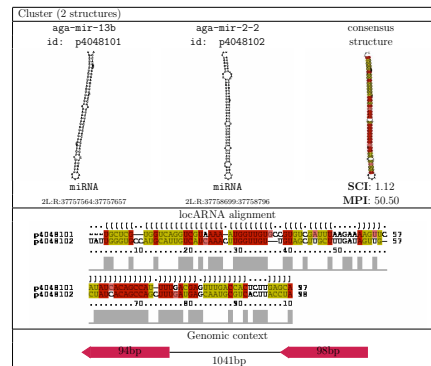
If we additionally restrict our search to putative miRNA cluster members, which are very common genomic organisations of miRNA precursors in metazoans, we can lower the number of candidates by considering only those which are found in the vicinity of known pre-miRNAs. The price to pay for this reduction is that we risk missing yet unidentified miRNAs which happen to occur in genomic locations far from previously identified precursors, or which are not part of a miRNA cluster. By selecting candidates with viable annotation and at a genomic distance not greater than 50 kb (as suggested in Baskerville and Bartel (2005)) from pre-miRNAs, we reduce our candidate set to 439 for *A. gambiae* and 1 604 for *D. melanogaster*.

After having determined, for each dataset, the reduced list of candidates, a more detailed analysis was performed for two different approaches, both using LOCARNA as a method to perform structural hierarchical clustering of our candidates along with the annotated precursors. Unlike before, we do not use a parameterised partition rule to enumerate our clusters at different  $k$ -levels. Instead, we use

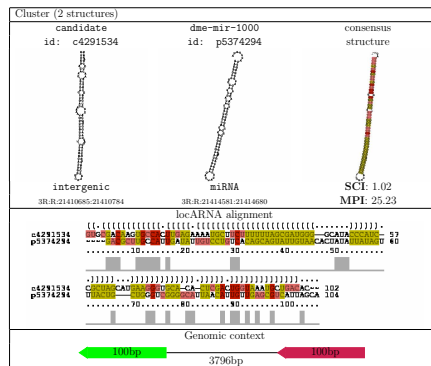
criteria aimed at identifying miRNA genomic clusters. The first approach consists in identifying, starting from the leaves of the similarity tree, the smallest structural clusters that include at least one known precursor, and within these clusters we enumerate all subsets of stem-loops that are located in close genomic proximity to each of the precursors in the structural cluster. This way we can identify the candidates which are both structurally similar and occurring close to each given precursor in the genome. The second approach drops the requirement that a known precursor be present and simply identifies leaves in the LocARNA similarity tree, extracting the top-scoring clusters in terms of SCI (Washietl *et al.*, 2005) (Structure Conservation Index). Additionally, we enumerate all subsets of stem-loops that are in close genomic proximity to each other, regardless of whether a known pre-miRNA is present. This way we try to identify candidates which are both similar and in the vicinity of one another. Subsets of stem-loops identified in the second approach which happen to be included in the output of the first approach are discarded.

The first approach identifies 108 and 422 candidates in the extracted clusters from *A. gambiae* and *D. melanogaster*, respectively, of which 5 and 11 were in the genomic vicinity of the precursors included in their respective clusters. A total of 9 and 23 potential genomic clusters of pre-miRNAs (not to be confused with structural clusters) corresponding to the relevant subsets of stem-loops of each structural cluster were identified using this approach for *A. gambiae* and *D. melanogaster*. Respectively, 4 and 13 were composed entirely of known pre-miRNAs at a median distance of 3 327 and 237.5 bp while 5 and 10 contained precursor candidates at a median distance of 12 424 and 22 820 bp (see Supplementary materials for the list of most promising cluster candidates). The genomic clusters exclusively made up of precursors, as the one shown in Fig. 3(a), attest the ability of our method to identify structurally homogenous pre-miRNA genomic clusters while the clusters which include new candidates, as seen in Fig 3(b), may indicate new instances of this type of genomic organisation and plausible miRNA precursors.

The second approach, which purports to identify potential genomic clusters where all members are non-annotated is, naturally, limited to those candidates which happen to be included in the initial set and are therefore close to known precursors in the genome, but which are structurally more similar to each other than to any pre-miRNA. This approach identified 481 and 1618 candidates in the extracted structural clusters for *A. gambiae* and *D. melanogaster*, respectively, of which 81 and 147 were not more than 50 kb away from another stem-loop in the same structural cluster. The potential genomic clusters identified for this approach total 7 and 65 at a median distance of 17 880 and 12 797 bp, respectively for *A. gambiae* and *D. melanogaster*. Interestingly, there are several instances of identical or highly similar candidates in both datasets that make up genomic clusters. Some of these clusters (see Supplementary materials) have up to three identical candidate members and are evenly spaced across the genome. They are either the result of a duplication event (more specifically, triplication) or they correspond to instances of long repeats, transposons or other repeat sequences. In any case, the fact that they were selected based on their structural similarity to known pre-miRNAs hints at the closeness of these structures and supports the much discussed possibility that repetitive sequences may serve



(a) Genomic cluster of known pre-miRNAs in *A. gambiae*



(b) Putative genomic cluster in *D. melanogaster*

**Fig. 3.** Genomic clusters of pre-miRNAs. Shown are the secondary structure of both stem-loops, the consensus structure along with the SCI (structure conservation index) and the MPI (mean pairwise identity), the LocARNA alignment and a representation of their genomic loci.

as a dormant repository of stem-loops which could be co-opted as pre-miRNAs (Smalheiser and Torvik, 2005).

## 4 CONCLUSION

We have presented a method to assess the sequence/structures similarity of a large dataset of hairpins in search for novel pre-miRNAs and we have placed these candidates in a multidimensional space in a way that reflects their structural characteristics. The portion of the multidimensional space selected around the known pre-miRNAs purports to include most structures which have the potential of being efficiently recognised by the cellular miRNA-processing machinery.

The fact that this region is very dense in terms of the number of precursor candidates it contains tells us that a large number of genome locations have the potential to generate stable and robust structures which present sequence/structure similarities to known pre-miRNAs. The use of annotation information helps reducing the number of selected candidates but after this filtering step, which is nevertheless dependent on the quality and breadth of the available annotation data, they remain in the tens of thousands. Therefore, there is either an exceedingly large number of pre-miRNAs in these datasets or, more plausibly, most of these candidates are not efficiently transcribed but could otherwise be recognised as miRNA precursors.

The initial set of candidates extracted from the genomes of *A. gambiae* and *D. melanogaster* and described in Mendes *et al.* (2010) consisted of 2 245 014 and 1 316 305 candidates, respectively. The total number of candidates, after applying stability and robustness measures described in earlier work and the structural analysis along with the annotation filtering detailed in this paper, is reduced by two orders of magnitude to 44 210 and 40 582.

Unlike many machine learning approaches to the identification of miRNA precursors that use features of the sequence and secondary structure to provide a classifier, our approach does not need to postulate a set of negative examples. In fact, we contend that if the purpose is to characterise the structures which have the potential of being recognised by the enzymes involved in miRNA maturation, one needs to reduce one's dependence on the positive set as well, since it will most likely not be representative – it suffices to observe that the set of recognisable structures is surely larger than the set of all the pre-miRNAs contained in the genome and that these two sets are subject to different evolutionary constraints. In our work, information about known precursors is used merely to pinpoint a region of interest in our multidimensional representation of sequence/structure features. Admittedly, this approach is not guaranteed to identify the entire portion of the feature space where structures recognisable by the miRNA processing machinery are located, since such a subspace is surely much larger than the examples that could ever be instantiated in any genome. However, our method outperforms a machine learning approach based on a support vector machine in one dataset and has comparable performance in the other, for larger training sets. For sample groups with a greater number of positive examples it outperforms the machine learning method in both datasets.

One can further limit the candidates to those occurring in the genomic vicinity of known precursors and which could, therefore, be part of miRNA clusters together with pre-documented pre-miRNAs. This approach produces a greatly reduced set of candidates (439 for *A. gambiae*, and 1 604 for *D. melanogaster*), even using a very liberal definition of miRNA cluster. This significant reduction of the number of candidates, albeit at the expense of the ability to identify novel miRNAs located elsewhere in the genome, elicits both the possibility of experimental validation and further detailed computational analyses. To this effect, it was possible to identify several plausible miRNA clusters with structurally similar stem-loops by performing conventional structural clustering over this reduced set, along with an analysis of their genomic disposition.

## ACKNOWLEDGEMENT

The contributions of NM and ATF were supported by Fundação para a Ciência e a Tecnologia (INESC-ID multiannual funding) through the PIDDAC Program funds and by PTDC program [PTDC/AGR-GPL/098179/2008]. The contributions of NM and MFS were funded by the Agence Nationale de la Recherche [BRASERO project BLAN06-3 138859, MIRI project BLAN08-1 335497] as well as by the ERC AdG SISYPHE held by MFS. This work was also funded by the German Research Foundation [DFG grants BA 2168/3-1, BA 2168/4-1 to RB], the Excellence Initiative of the German Federal and State Governments [grant EXC 294 to RB], and SH was supported by the German Federal Ministry of Education and

Research [BMBF grant 0313921 to RB].

## REFERENCES

- Bartel, D. P. (2009). MicroRNAs: target recognition and regulatory functions. *Cell*, **136**(2), 215–33.
- Baskerville, S. and Bartel, D. P. (2005). Microarray profiling of microRNAs reveals frequent coexpression with neighboring miRNAs and host genes. *RNA*, **11**(3), 241–247.
- Bompfunewerer, A. F., Backofen, R., Bernhart, S. H., Hertel, J., Hofacker, I. L., Stadler, P. F., and Will, S. (2008). Variations on RNA folding and alignment: lessons from Benasque. *Journal of Mathematical Biology*, **56**(1-2), 129–144.
- Bradley, R. K., Pachter, L., and Holmes, I. (2008). Specific alignment of structured RNA: stochastic grammars and sequence annealing. *Bioinformatics*, **24**(23), 2677–83.
- Duda, R. O., Hart, P. E., and Stork, D. G. (2001). *Pattern Classification*. John Wiley & Sons, INC.
- Gorodkin, J., Heyer, L. J., and Stormo, G. D. (1997). Finding the most significant common sequence and structure motifs in a set of RNA sequences. *Nucleic Acids Research*, **25**(18), 3724–32.
- Griffiths-Jones, S., Bateman, A., Marshall, M., Khanna, A., and Eddy, S. R. (2003). Rfam: an RNA family database. *Nucleic Acids Research*, **31**(1), 439–41.
- Havgaard, J. H., Lyngso, R. B., Stormo, G. D., and Gorodkin, J. (2005). Pairwise local structural alignment of RNA sequences with sequence similarity less than 40%. *Bioinformatics*, **21**(9), 1815–24.
- Heyne, S., Will, S., Beckstette, M., and Backofen, R. (2009). Lightweight comparison of RNAs based on exact sequence-structure matches. *Bioinformatics*, **25**(16), 2095–2102.
- Höchsmann, M., Töller, T., Giegerich, R., and Kurtz, S. (2003). Local similarity in RNA secondary structures. In *Proceedings of Computational Systems Bioinformatics (CSB 2003)*, volume 2, pages 159–168. IEEE Computer Society.
- Hofacker, I. L., Fontana, W., Stadler, P. F., Bonhoeffer, S., Tacker, M., and Schuster, P. (1994). Fast folding and comparison of RNA secondary structures. *Monatshfte Chemie*, **125**, 167–188.
- Hofacker, I. L., Fekete, M., and Stadler, P. F. (2002). Secondary structure prediction for aligned RNA sequences. *Journal of Molecular Biology*, **319**(5), 1059–66.
- Hofacker, I. L., Bernhart, S. H., and Stadler, P. F. (2004). Alignment of RNA base pairing probability matrices. *Bioinformatics*, **20**(14), 2222–7.
- Kaczkowski, B., Torarinsson, E., Reiche, K., Havgaard, J. K., Stadler, P. F., and Gorodkin, J. (2009). Structural profiles of human miRNA families from pairwise clustering. *Bioinformatics*, **25**(3), 291–4.
- Lee, R. C., Feinbaum, R. L., and Ambros, V. (1993). The *C. elegans* heterochronic gene *lin-4* encodes small RNAs with antisense complementarity to *lin-14*. *Cell*, **75**(5), 843–54.
- Lindow, M. and Gorodkin, J. (2007). Principles and limitations of computational microRNA gene and target finding. *DNA Cell Biology*, **26**(5), 339–51.
- Mathews, D. H. and Turner, D. H. (2002). Dynalign: an algorithm for finding the secondary structure common to two RNA sequences. *Journal of Molecular Biology*, **317**(2), 191–203.
- Mendes, N. D., Freitas, A. T., and Sagot, M.-F. (2009). Current tools for the identification of miRNA genes and their targets. *Nucleic*

- 
- Acids Res*, **37**(8), 2419–33.
- Mendes, N. D., Freitas, A. T., Vasconcelos, A. T., and Sagot, M.-F. (2010). Combination of measures distinguishes pre-miRNAs from other stem-loops in the genome of the newly sequenced *Anopheles darlingi*. *BMC Genomics*, **11**, 529.
- Sankoff, D. (1985). Simultaneous solution of the RNA folding, alignment and protosequence problems. *SIAM J. Appl. Math.*, **45**(5), 810–825.
- Siebert, S. and Backofen, R. (2005). MARNA: multiple alignment and consensus structure prediction of RNAs based on sequence structure comparisons. *Bioinformatics*, **21**(16), 3352–9.
- Smalheiser, N. R. and Torvik, V. I. (2005). Mammalian microRNAs derived from genomic repeats. *Trends Genetics*, **21**(6), 322–6.
- Washielt, S., Hofacker, I. L., and Stadler, P. F. (2005). Fast and reliable prediction of noncoding RNAs. *Proceedings of National Academy of Sciences USA*, **102**(7), 2454–9.
- Will, S., Reiche, K., Hofacker, I. L., Stadler, P. F., and Backofen, R. (2007). Inferring non-coding RNA families and classes by means of genome-scale structure-based clustering. *PLOS Computational Biology*, **3**(4), e65.
- Xue, C., Li, F., He, T., Liu, G.-P., Li, Y., and Zhang, X. (2005). Classification of real and pseudo microRNA precursors using local structure-sequence features and support vector machine. *BMC Bioinformatics*, **6**, 310.
- Youden, W. J. (1950). Index for rating diagnostic tests. *Cancer*, **3**(1), 32–35.