

## Feature Review

# *De novo* prediction of structured RNAs from genomic sequences

Jan Gorodkin<sup>1</sup>, Ivo L. Hofacker<sup>2</sup>, Elfar Torarinsson<sup>1</sup>, Zizhen Yao<sup>3</sup>, Jakob H. Havgaard<sup>1</sup> and Walter L. Ruzzo<sup>3,4</sup>

<sup>1</sup> Section for Genetics and Bioinformatics, IBHV and Center for Applied Bioinformatics, University of Copenhagen, Grønnegårdsvej 3, DK-1870 Frederiksberg C, Denmark

<sup>2</sup> Institut für theoretische Chemie, University of Vienna, Währingerstr. 17, A-1090 Vienna, Austria

<sup>3</sup> Fred Hutchinson Cancer Research Center, 1100 Fairview Ave. N. PO Box 19024, Seattle, Washington 98109, USA

<sup>4</sup> Departments of Computer Science and Engineering and Genome Sciences, University of Washington, Seattle, Washington 98195, USA

**Growing recognition of the numerous, diverse and important roles played by non-coding RNA in all organisms motivates better elucidation of these cellular components. Comparative genomics is a powerful tool for this task and is arguably preferable to any high-throughput experimental technology currently available, because evolutionary conservation highlights functionally important regions. Conserved secondary structure, rather than primary sequence, is the hallmark of many functionally important RNAs, because compensatory substitutions in base-paired regions preserve structure. Unfortunately, such substitutions also obscure sequence identity and confound alignment algorithms, which complicates analysis greatly. This paper surveys recent computational advances in this difficult arena, which have enabled genome-scale prediction of cross-species conserved RNA elements. These predictions suggest that a wealth of these elements indeed exist.**

## Introduction


Non-coding RNAs (ncRNAs) are functional transcripts that do not encode proteins. A handful of examples, such as transfer and ribosomal RNAs, have been well known since the dawn of molecular biology, and probably have existed since the dawn of life. These few examples have critical and deeply central roles, but ironically, elucidation of the full spectrum of ncRNA activity has received relatively little attention. Within the past 10–15 years, however, several striking discoveries, including RNA interference, microRNAs and riboswitches, have demonstrated that RNAs have unexpectedly diverse, sophisticated and important roles in all living organisms, which has sparked renewed interest in the “modern RNA world” [1]. To hint at the scope of the issue, only 1.2% of the human genome encodes protein [2], but recent data suggest that 90% of the genome is transcribed on one or both strands, to some extent, at some time, in some tissues [3]. The significance of this observation remains unclear and controversial, but a growing body of evidence points to the presence of many functionally

important non-coding transcripts [4]. Even if the bulk of the non-coding transcription is noise, this still provides a vast substrate upon which natural selection could have acted to generate a multitude of biologically important ncRNAs. Thus, even striking ncRNA-related discoveries, such as microRNAs, may only be the tip of the iceberg.

Interest in conducting computational screens for functional RNAs has risen with their increased prominence [5]. However, in contrast to protein coding genes, whose regular codon structure provides strong signals within nucleotide sequences, the signals for ncRNAs are subtler. For example, many ncRNAs are believed to be processed out of longer primary transcripts, including introns of protein-coding genes, and hence, lack features such as proximal promoters in addition to codon structure [1]. There is, however, one general characteristic shared by many (but not all) known ncRNAs: they fold into complex shapes that are crucial to function and are thus conserved. Although prediction of RNA 3D structure is less well developed than protein structure prediction, the prediction of RNA secondary structure (the set of intramolecular, largely Watson–Crick, base pairs that define the fundamental units from which the tertiary structure is formed) is reasonably well understood and computationally tractable. Furthermore, secondary structure also tends to be conserved, whereas primary sequences often evolve more rapidly (Figure 1), and are of limited use except when searching for close homologs. These facts make secondary structure the key signal to be exploited in ncRNA prediction. Although secondary structure (henceforth, simply structure) prediction is tractable, it is not trivial (e.g. involving interactions between nucleotides at variable and sometimes large distances in the primary sequence), which makes the problem challenging intellectually and computationally [5].

Most RNA molecules will fold into some secondary structure, but given the importance of secondary structure in functional ncRNAs, it is reasonable to ask whether they have more stable structures than random genomic sequences. A negative answer to this question was provided about a decade ago. That is, the folding energy of RNA genes is not distinguished easily from that of appropriately chosen

Corresponding authors: Gorodkin, J. (gorodkin@genome.ku.dk); Ruzzo, W.L. (ruzzo@cs.washington.edu).



**Figure 1.** Compensating base changes. Changes in base pairing might preserve structure, but not the primary sequence. In addition to the usual Watson-Crick base pairs, less stable G-U pairs (sometimes called “wobble pairs”) are often seen in RNAs, and are evolutionarily important because they allow single base substitutions that are not structurally disruptive. This might allow sequences to accumulate substitutions much more rapidly than would be the case if both nucleotides in a base pair needed to be changed more or less simultaneously. Adapted with permission from Ref. [80].

background sequences, such as randomly shuffled versions of the RNA gene itself [6,7]. Therefore, a simple screening approach that seeks unusually stable genomic segments, will generally not work. MicroRNAs, whose precursors form unusually stable stem-loops, are a notable exception [8,9].

Although searching for possible structured elements in one nucleotide sequence is generally ineffective, as discussed above, searching in multiple (orthologous) sequences can be highly effective, since evolutionary conservation highlights functionally important regions of all kinds. Importantly, such searches leverage the rapidly increasing body of comparative genomic sequence data. A key issue, however, is that the evolutionary signature of an RNA gene is quite different from that of a protein-coding gene. In particular, as noted earlier, the nucleotide sequence of an ncRNA might evolve relatively rapidly, which makes identification and alignment of orthologous sequences difficult or impossible, especially when using tools that focus only on sequences. However, patterns of compensating base changes, for example, an A-U base pair in a human RNA sequence that corresponds to a C-G pair in mice, can provide evidence to support the existence of a conserved RNA structure (without requiring conserved sequence; Figure 1) and insight into the structure itself. Indeed, short of X-ray crystallography, this type of comparative analysis, done carefully by human experts [10], has been the gold standard for RNA secondary structure prediction for more than 40 years. In a nutshell, this highlights the key challenges in computational prediction of ncRNAs: to find orthologous regions, expose the common structure therein, and do so rapidly and accurately.

The comparative approach is important for an additional reason. Over the next few years, we expect that emerging technology such as high-throughput sequencing of RNA (RNAseq [11]) will reveal the transcriptomes of many organisms with unprecedented depth and precision. Yet, given the extensive breadth of genomic transcription now observed, at least in mammals, evidence of transcription can no longer be taken as proof of functional importance. Any observed transcript might just be noise or incompletely degraded detritus that arises from the expression of some nearby, functionally important RNA. Furthermore, experimental protocols will remain limited with regard to the diversity of species, cell types, states, growth and stress conditions that are probed. Consequently, lack of measured expression of a given genomic

segment is not proof of lack of function. By contrast, evolutionary conservation strongly suggests functional importance, whether or not expression has already been verified experimentally. This does not deny the value of experimental evidence, of course, nor the existence of functionally important species-specific ncRNAs, but merely argues that detection of evolutionarily conserved ncRNAs by comparative genomics is a powerful tool, and belongs in any effort to understand living systems.

Computational search for conserved RNA structure does have certain important limitations. We highlight two of them here, because they color much of what follows. First, these searches are expensive computationally, principally because of the nature of the underlying RNA folding algorithms that need to be applied to these multiple sequences [12,13]. Even single-sequence folding algorithms have run times that grow as the cube of the sequence length. Applied naively, however fast screening of a sequence of 1 kb is, it will be 1000 times slower for 10 kb and 1 million times slower for 100 kb. Hence, all successful programs in this arena are engineered carefully to control run time, which entails some, hopefully modest, loss of accuracy on long genomic sequences. For example, one simple, widely used strategy is the “sliding window” approach, wherein the genomic sequence is cut into multiple, overlapping, fixed-length segments (windows) that are processed separately. This obviously limits the cubic run time penalty to the length of the window, but unfortunately, also limits the lengths of discoverable structures and risks arbitrarily truncating them. Even using substantially more sophisticated techniques, genome-scale ncRNA analyses often consume tens to hundreds of computer years. These high computational costs are one reason why ncRNA gene finding is still in its infancy.

The second significant limitation of these general searches for conserved RNA secondary structures is more conceptual. It is natural to want to think of each element discovered as an ncRNA gene, but the truth is more complex because the approaches described here might generate only a partial picture for each ncRNA. For example, technical limitations related to window boundaries or splicing could result in partial or fragmentary predictions. More intrinsically, some ncRNAs lack conserved secondary structures, or may have only patches of conserved structure embedded in longer, largely unstructured transcripts. Additionally, conserved, functionally important RNA structures, such as the selenocysteine insertion sequence (SECIS element), are known to exist in mRNAs, usually in their untranslated regions. Therefore, identification of RNA structures in genomic data should trigger post-processing steps and follow-up experiments to characterize transcript boundaries (and function) more precisely. For reasons of simplicity, however, we will refer to individual conserved structures as ncRNA genes.

This review focuses explicitly on computational prediction of ncRNA elements by comparative genomics, that is, the discovery of conserved structured elements in multiple genomic sequences. Other methods for *de novo* prediction have succeeded in some contexts (e.g. exploiting organism-specific differences in mono- or dinucleotide frequencies of ncRNAs versus background [14–16]), but the comparative

approach appears to be the most broadly applicable. We will say little about the related ncRNA homology search problem (finding new instances of a particular RNA family given one or more examples) but this equally important task comes with its own set of issues [17], especially the difficulty of finding homologs outside the phylogenetic range of known examples.

### From RNA folding to gene finding

Even though RNA structure cannot be detected reliably by merely folding single sequences, the principles obtained from folding single sequences are fundamental and often constitute an implicit part of more elaborate methods. For example, to date, no large-scale RNA structure screens have accounted for so-called pseudoknots, because the underlying RNA folding algorithms do not do that. Without pseudoknots, RNA secondary structure can be represented by nested parentheses. For example, the hairpin structure of the sequence AAAAUUCGGCAAUUUU with base pairs exactly between As and Us can be written as ((((((...)))))). Although most large RNA structures do contain pseudoknots, they make up a relatively small proportion of the overall structure. Algorithms that account for these more general cases are dramatically more complex (and slower) for seemingly modest gains in accuracy.

Folding of sequences typically is performed through energy minimization, whereby different structural elements are associated with experimentally measured energy parameters. Well-known programs for this task are mfold [18] and RNAfold [19]. Rather than minimizing the free energy, folding can also be carried out in a probabilistic framework, wherein the structure with the highest probability is sought. This is achieved through so-called stochastic context-free grammars (SCFGs; for example see Refs. [20,21]), which typically are parameterized using probabilities extracted from curated rRNA and tRNA alignments [22,23]. The probabilistic and energetic frameworks are connected intimately, because one expects energetically destabilizing substitutions to be evolutionarily disfavored. Indeed, recent work has shown that statistics derived from good structural alignments can be used to improve upon experimentally measured energy parameters [24]. Energy-based and SCFG-based methods are now part of the core inventory of ncRNA gene prediction methods. These principles for folding single sequences are employed for multiple sequences as well, often together with schemes that simultaneously search for compensating base changes.

### *In silico* screening for RNA structures

Prediction of RNA structure in genomic sequences is related closely to the existing methods for RNA structure prediction. As indicated above, multiple sequences are needed to predict RNA structures reliably. Several strategies exist for RNA structure prediction based on multiple sequences [25], which can be categorized loosely as “align-first,” “fold-first” and “joint.” As the name suggests, align-first strategies start by aligning all sequences using standard multiple sequence alignment tools, followed by inference of their presumed common structure. That

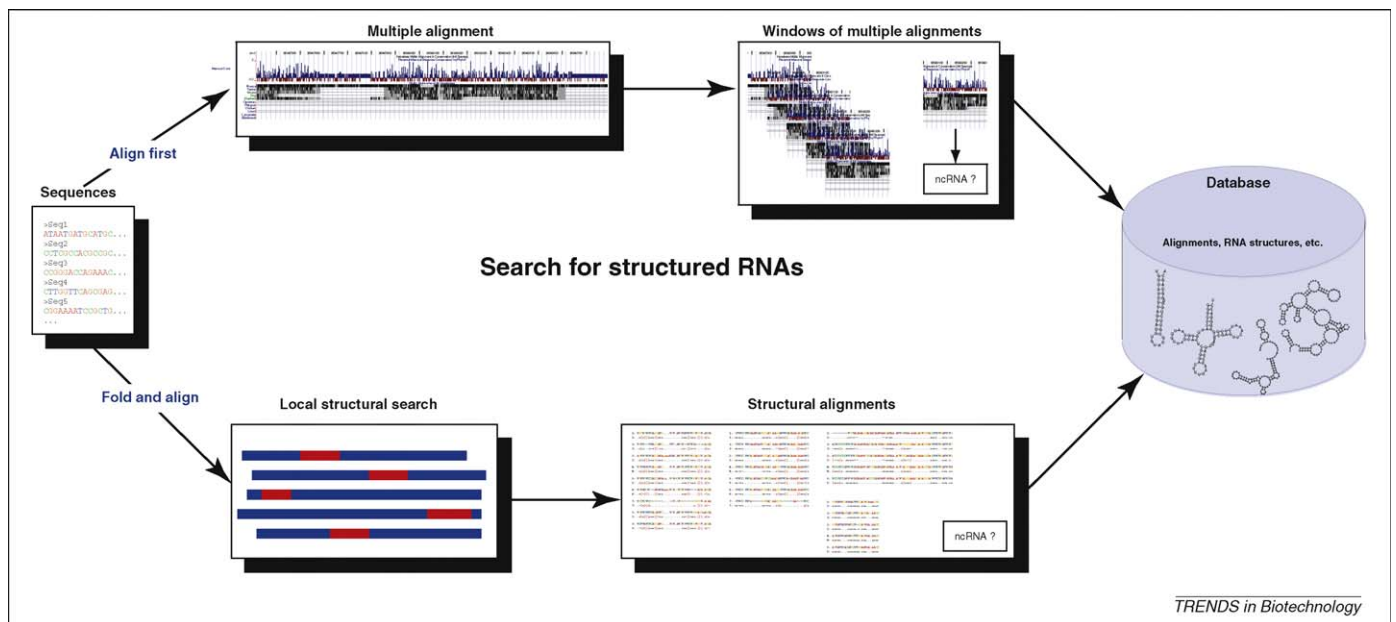
inference typically proceeds by some combination of folding energy prediction and detection of compensatory base changes, for example, as quantified by high mutual information [20,22] between pairs of columns in the alignment. These methods work best when sequence conservation is sufficiently low that compensatory changes are not rare, but high enough that most are bracketed by well-conserved patches, which constrain the changes to be correctly aligned (and hence visible as a high mutual information score between paired columns). Fold-first strategies fold individual sequences separately, and then align the structures to each other. These strategies can be expected to excel when sequence conservation is low, but structure is clear-cut, stable and well-conserved. Joint strategies, such as approaches based on Sankoff's algorithm [26], simultaneously align and fold their input sequences. In a sense, the align-first and fold-first strategies are aimed at opposite evolutionary extremes. The joint approach, with an appropriate joint model for sequence and structure evolution, subsumes the other two, as shown in Ref. [27]. In practice, however, high computational costs and difficulties in appropriately specifying such an evolutionary model have hindered the joint methods, which has left viable niches for all three approaches. Current approaches for *de novo* genomic screening mostly utilize methods based on the fold-first and joint strategies, as exemplified in Figure 2, in which prediction is based on finding structures within relatively short sequences locally, rather than globally. Meanwhile, fold-first methods mainly have been applied in homology searches [28,29].

Besides ignoring pseudoknots, another crucial limitation of current genomic screening methods is their inability to cope with structural variation. That is, they assume that the individual RNA sequences share a conserved structure with little variation. In particular, structural inserts, which result in large length differences, are handled poorly by current programs. Such structural inserts are, however, frequently observed in evolutionarily old ncRNAs, such as RNaseP [30]. Even the state-of-the-art SCFG-based homology search program Infernal deals only indirectly with structural inserts through its ability to match a local substructure [31]. Recent programs for comparison of RNA structure, such as LocARNA [32], point out some ways to address the problem. Additionally, the presence of two nearby structural alignments suggests an intervening structural insertion/deletion in some of the sequences, a situation that potentially could be detected in a post-processing step. However, none of these approaches fully overcomes this limitation of the current *de novo* screening methods.

In the following discussion, we provide an overview of the main methodologies applied for *de novo* RNA structure discovery. These are also listed in Table 1.

### *Exploiting sequence-based multiple alignments for ncRNA screens*

A natural approach for *de novo* structure-based search is to exploit existing sequence based alignments, that is, the align-first strategy outlined above. QRNA was an early tool of this kind [33]. It screened pairwise BLASTn [34] alignments using three probabilistic models: a pair-SCFG for



**Figure 2.** Strategies for ncRNA screening. The upper path illustrates RNA structure prediction using existing sequence alignments that are divided into overlapping windows (an align first strategy). By contrast, in the lower path, labeled “fold and align,” sequence and structure alignments are performed directly from unaligned sequence data (a joint strategy), which searches simultaneously for conserved structure and sequence, and results in structural alignments. To date, alternative fold-first strategies have not been applied to genome-scale screening.

detecting RNA structures, and two pair-HMMs (hidden Markov models) for detecting coding and background sequences, respectively. QRNA screens on sequences from *Escherichia coli* and *Saccharomyces cerevisiae* lead to the identification and subsequent experimental verification of several novel ncRNAs [35,36].

Given the apparent importance of a stable structure, it is reasonable to speculate that the predicted folding energy of a given sequence would be a useful clue for the detection of putative ncRNAs. This appears to be true, but has proved difficult to verify, and required careful statistical analysis

and the consideration of both mono- and dinucleotide frequencies, because the latter affect stacking energies [6,7,37]. The magnitude of the energy difference between random sequences and functional ncRNAs, although generally not sufficient for reliably detecting them based on the genome of a single species, becomes more pronounced in multiple alignments. AlifoldZ introduced this approach [38]. It compared the folding energy that resulted from a consensus structure prediction on the native alignment with those of shuffled alignments, that is, those that arise from randomly shuffling the columns of the native alignment. More

**Table 1. A Summary of the main screens described here**

| Method <sup>a</sup>            | Organisms <sup>b</sup> | #Cand <sup>c</sup> | Experimental verification               | Size <sup>d</sup> | Cons <sup>e</sup> | Constraints   |
|--------------------------------|------------------------|--------------------|---|-------------------|-------------------|---|
| QRNA [33,35] (pair-SCFG)       | Bacteria               | 275                | 11 of 49 randomly sampled, northern     | 200               | Seq               | Sliding windows; sequence based alignments                  |
| RNAz [39,40] (Energy, SVM)     | Vertebrates            | 46,000             | 3 of top 5, northern                    | 200               | Seq               | Sliding windows; sequence based alignments                  |
|                                | ENCODE                 | 7,093              | 43 of 175, RT-PCR <sup>f</sup>          |                   |                   |   |
| EvoFold [43] (phylo-SCFG)      | Vertebrates            | 48,500             | (subsequent studies)                    | 750 (<50)         | Seq               | Sliding windows; sequence based alignments                  |
|                                | ENCODE                 | 9,953              | 43 of 175, RT-PCR <sup>f</sup>          |                   |                   |   |
| FOLDALIGN [52,59,61] (Sankoff) | Human-mouse            | 1,300              | 32 of 36, RT-PCR; 4 of top 12, northern | 200               | Str               | Motif size and type in local pairwise structural alignments |
| Dynalign [53,54] (Sankoff)     | Bacteria               | 1,200              | none reported                           | 150               | Seq               | Sliding windows; HMM based alignment constraints            |
| CMfinder [62,64,83] (SCFG, EM) | Bacteria               | 1,466              | (subsequent studies)                    | None (<200)       | Str               | Motif size and type in local multiple structural alignments |
|                                | ENCODE                 | 6,587              | 10 of top 11, RT-PCR; 1 of 10, northern |                   |                   |   |

<sup>a</sup>Name and core methodology

<sup>b</sup>Genomes screened; “ENCODE” means ENCODE regions of 17-way vertebrate alignments

<sup>c</sup>Number of predicted RNA structures

<sup>d</sup>Approximate maximum size of motifs; number in parenthesis is typical size of motifs

<sup>e</sup>Type of conservation exploited to produce alignments (Seq means sequence-level; Str means structure-level)

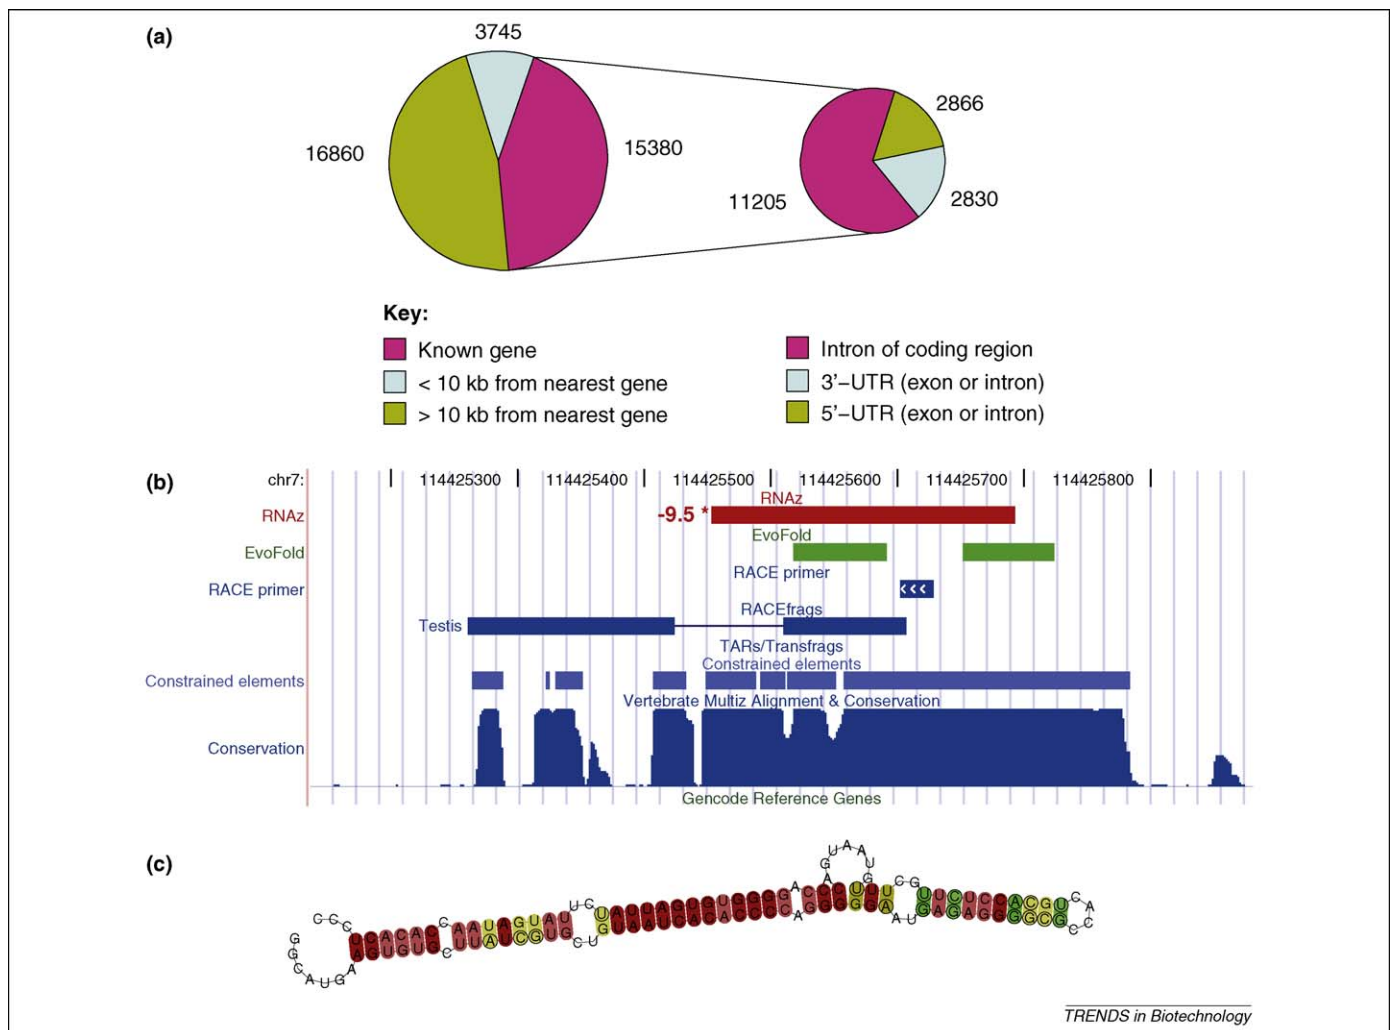
<sup>f</sup>This number is joint for RNAz and EvoFold

recently, the RNAz framework [39,40] has improved speed and sensitivity by introducing a fast regression to estimate the stability of a native RNA compared to shuffled sequences, together with a new measure for structural conservation and a support vector machine (SVM) classifier to predict whether an alignment contains an ncRNA.

None of the methods discussed so far directly exploit phylogenetic information, for example, the tree structure and branch lengths that relate individual sequences in the alignment. Pfold [41,42] and EvoFold [43] do this. These methods compare two alternative statistical models: a background model for unstructured sequence versus an SCFG model of RNA sequence and structure, both of which are evaluated with respect to the given phylogeny. A section of the alignment is predicted to be an ncRNA if it is more likely under the RNA model than under the background model. In particular, a pattern of substitutions in putatively base-paired columns that favor structure-preserving compensatory changes tends to support the RNA model over the background model.

QRNA, RNAz and EvoFold all use sliding window strategies. The main advantage of this approach is speed; its major disadvantage is its potential for mis-prediction near the window boundaries, including the possibility of missing RNAs that are longer than the overlap of the windows. These drawbacks might be ameliorated by adjusting window lengths dynamically, but to the best of our knowledge, this has not been attempted.

The University of California Santa Cruz (UCSC) genome browser provides whole-genome MULTIZ alignments of vertebrates [44]. Initial screens of their eight-way alignments (human, chimp, dog, mouse, rat, chicken, zebrafish and fugu) using RNAz and EvoFold have resulted in an unexpectedly large number of ncRNA candidates: 36 000 [40] and 48 000 [43], respectively. In both cases, candidates were widely distributed throughout the genome, as illustrated in Figure 3 for the RNAz screen. One of the EvoFold candidates, the ncRNA gene HAR1F (human accelerated region 1F), has been the subject of an in-depth follow-up study [45] that has demonstrated its specific expression



**Figure 3.** RNAz screen. **(a)** Genomic locations of RNAz predictions from Ref. [40] compared to annotated protein-coding genes. As illustrated here, ncRNA candidates are distributed widely in the genome. For example, 16 860 of the RNAz predictions lie more than 10 kb from the nearest annotated protein-coding gene. It is plausible that the fraction of ncRNAs in this category is underestimated, since alignments tend to be less reliable in these regions, and the align-first strategy used by this screen depends on them. Of the 15 380 predictions that lie within annotated protein-coding genes, the majority (11 205) lie in introns flanked by coding exons, but a significant fraction lie in 5' and 3' UTRs, or their introns. (Some candidates are counted in multiple categories because of ambiguous gene annotation. Coding exons were not screened.) **(b)** Detailed view of a genomic region on human chromosome 7 identified in the ENCODE screen [47]. It contains overlapping predictions from RNAz (red) and EvoFold (green), as well as an unusually high AlifoldZ Z-score (-9.5). In addition, the region corresponds to an island of high conservation. The RACEfrags track verifies that an ncRNA element from this locus is expressed in testis. **(c)** The predicted structure of a portion of the element is shown. Adapted with permission from Ref. [40].

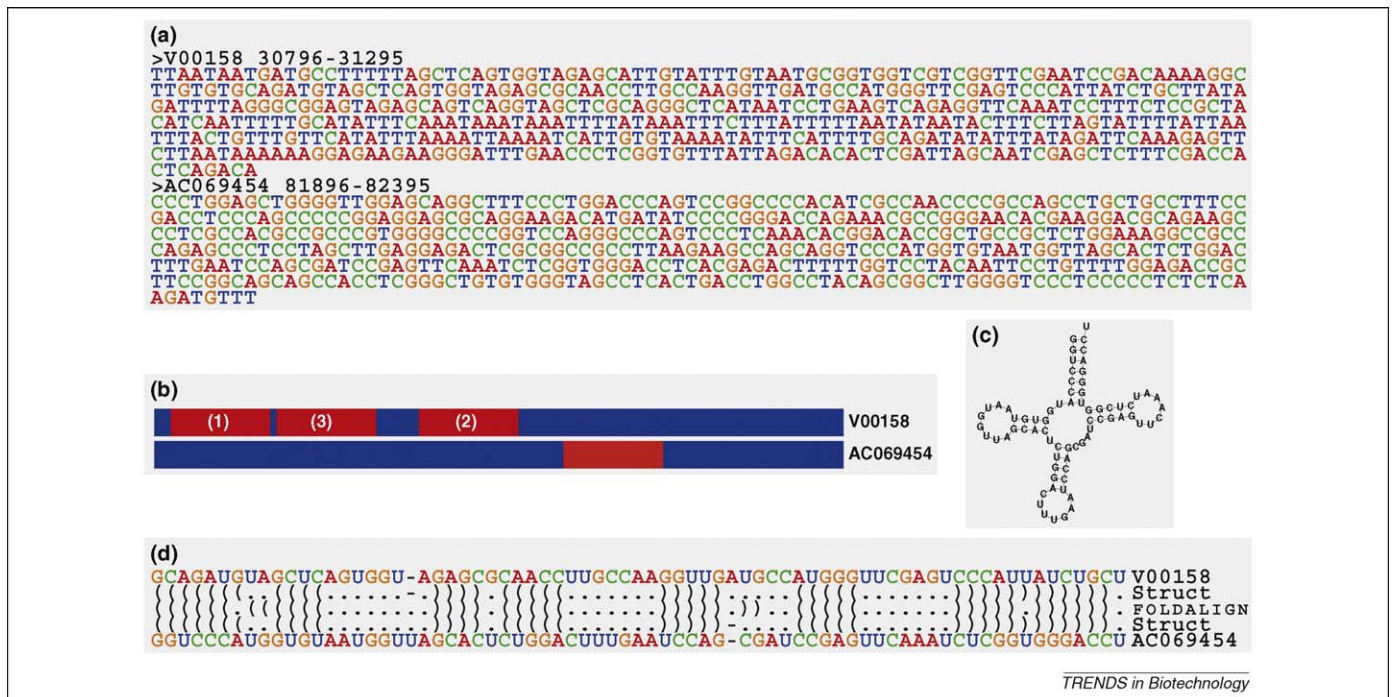
during cortical development and accelerated evolution along the human lineage. RNAz and EvoFold have since been applied in several other screens of diverse organisms. For example, an RNAz screen on porcine expressed sequence tag (EST) sequences has allowed for direct comparison of expression data in 92 cDNA libraries from 35 tissues and different developmental stages [46]. The screen has found that brain and developmental tissues host the highest relative numbers of predicted ncRNAs structures (~2% of genes expressed at detectable levels), whereas ncRNA content for the other tissues amounts to about 1%. A summary of other screens can be found on the RNA Structure in Genomic Sequence supplementary web-page: <http://genome.ku.dk/resources/rsgs>.

It has proven difficult to estimate the reliability of these predictions. The standard method of comparing predictions to some gold standard is problematic, because the available true positives, that is, known ncRNAs, although generally well-recovered by these methods, are potentially not representative of yet undiscovered ncRNAs. Furthermore, there is no satisfactory, biologically justifiable set of true negatives, particularly in light of the widespread transcription observed in mammals. Instead, studies generally estimate their false discovery rate (FDR) by comparison with randomly shuffled sequences or alignments. Based on this approach, the human screens described above have been estimated to have FDRs of 15–50% for RNAz and up to 70% for EvoFold [47]. This emphasizes the need for continued development of ncRNA prediction methods, as well as for follow-up experiments.

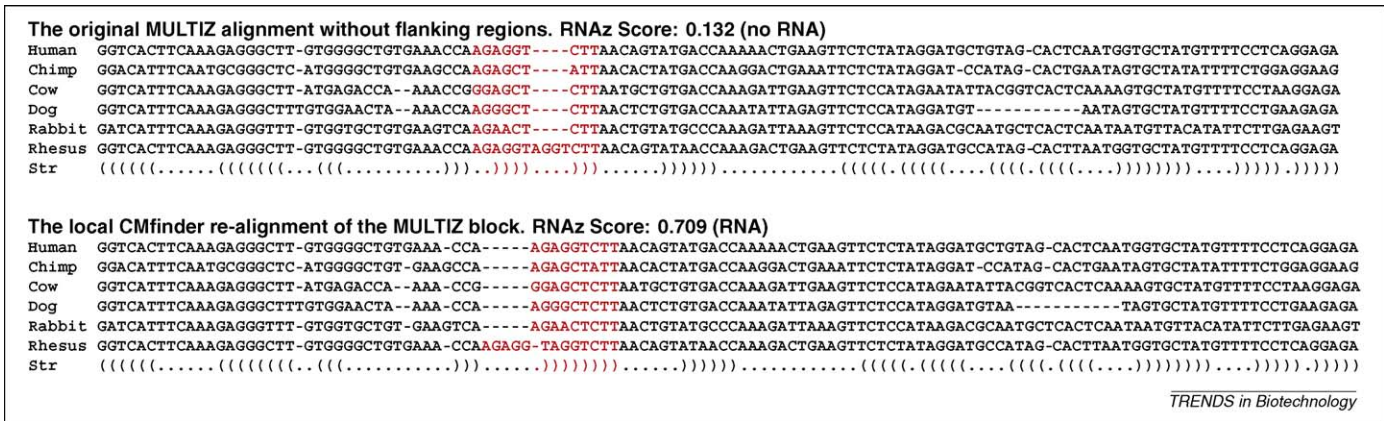
However, even if the true FDR in these studies were 90%, the remaining 10% still constitute thousands of novel ncRNAs in the human genome. One technical point that is worth noting is that the accuracy of the FDR estimation itself will depend on how well the randomized data reflect the complexities of real genomes. For example, the dinucleotide composition of an RNA sequence affects its stacking energy, thus randomization procedures that do not control for this yield biased FDR estimates [48]. Recent work has provided procedures that generate randomized multiple sequence alignments that approximately preserve dinucleotide-, gap-, and other statistics observed in the native alignments [48–50], thus providing a more appropriate null distribution for FDR estimation, but the FDR estimates reported in the ncRNA screens discussed here generally are not based on these improved methods.

### Genomic screening by pairwise local structural alignments

Although exploiting multiple alignments saves computational costs, there is the significant disadvantage that alignments might simply be incorrect, which will become increasingly likely as the sequence similarity drops. For example, the sensitivity of AlifoldZ drops sharply for pairwise sequence similarities below 60% [38]. This limitations can, at the expense of computational costs, be addressed within the above-mentioned joint approaches, for example, the Sankoff framework for simultaneously folding and aligning sequences [26]. Available methods using this strategy include FOLDALIGN [51,52], Dynalign [53,54],



**Figure 4.** The value of local, structural alignments. **(a)** Two genomic sequences of 500 nucleotides are shown. This segment of V00158, from *Eug. gracilis* chloroplast, contains three annotated tRNAs; AC069454, from human chromosome 17, contains one. The sequences are relatively dissimilar, as is apparent from the nucleotide color-coding; for example, the former has a G + C content of 35%, whereas the human sequence is 63% G + C. When comparing the sequences via standard pairwise BLAST (using default parameters), the result is “No significant similarity found.” **(b)** Comparison of the sequences using FOLDALIGN [51,52], a local, structural alignment tool, results in three pairwise structural alignments of the red region in humans to the three indicated *Eug. gracilis* regions, each corresponding to an annotated tRNA. The numbers in parentheses give the FOLDALIGN-score ranking. The sequence identities of the three pairwise alignments are: (1) 45%, (2) 48% and (3) 40%, respectively. **(c)** Secondary structure, as annotated in the tRNA database [81], of the human tRNA. **(d)** The structural alignment of region (3) with the human tRNA. The “Struct” lines give the structures of these two tRNAs, again from Ref. [81], where matching pairs of nested parentheses indicate canonical base pairs, dots indicate unpaired nucleotides, and dashes represent gaps. The middle line shows the consensus RNA structure predicted by FOLDALIGN from this pair of regions, which is in excellent agreement with Ref. [81].



**Figure 5.** Sequence-based alignments can obscure ncRNA structure. Even small changes in a sequence-based multiple alignment can significantly affect the recognition of ncRNAs. The upper panel shows a six-species MULTIZ alignment of a region that contains a box H/ACA snoRNA annotated in Rfam [82] (accession RF00402) that has been downloaded from the UCSC genome browser<sup>1</sup>. The use of RNAz to search for ncRNAs in this alignment resulted in a low score of 0.132, which suggests “no RNA”. In the lower panel, the same sequences have been aligned by CMfinder. Although only a few positions are realigned (key changes indicated in red), RNAz gives this structurally revised alignment a relatively high score of 0.709, correctly suggesting the presence of a conserved RNA structure. Matching parentheses in the structure lines (Str; bottom line in each group) indicate base pairs in the consensus secondary structure predicted by RNAz.

and LocARNA [32], which are all energy-based, as well as Consan [55] and StemLoc [56], which are SCFG-based. Furthermore, alternatives to the Sankoff framework have also been developed, such as SCARNA [57], which compares stem fragments to construct pairwise structural alignments.

The latest version of Dynalign [54] predicts ncRNAs using its own Sankoff-based pairwise structural alignments [53] and an SVM classifier trained on those alignments. In a pairwise analysis of *E. coli* and *Salmonella typhi*, its additional alignment flexibility appears to have resulted in improved sensitivity compared to RNAz, at least for regions that exhibit low sequence similarity (<50% identity). Using a sliding window approach based on genome-wide BLAST and MUMmer [58] alignments, approximately 1000 novel candidates have been reported in these two genomes [54].

To detect remotely homologous ncRNAs, it is desirable to extend such searches to directly perform local structural alignments, rather than relying on purely sequence-based tools such as BLAST. To illustrate this, we have examined a 500-nucleotide region from the human genome that contains one known tRNA and a 500-nucleotide region from the distant *Euglena gracilis* chloroplast genome that contains three tRNAs. BLAST reports no significant match between them. FOLDALIGN, however, detects one human sub-region that structurally aligns to three different locations within the chloroplast sequence, exactly corresponding to the tRNAs therein, despite the low (40–50%) sequence identities among these matching regions (Figure 4). Note the importance, and difficulty, of local alignment here, which allows a structure common to a small portion of each input sequence to be detected amidst extraneous flanking sequences; the algorithm must thus identify element boundaries as well as structure. The

FOLDALIGN screen takes about 2 min, however, compared to ~0.03 sec for BLAST, which further motivates the desire to accelerate these methods.

FOLDALIGN has also been used to screen regions of low sequence similarity between humans and mice [59]. These have been selected by choosing corresponding, but unaligned regions as identified via the UCSC genome browser MULTIZ alignments [60]. A total of 37 000 such regions have been screened, each in forward and reverse directions. The screen resulted in ~1300 candidates, with an FDR of ~50% as estimated from dinucleotide-controlled shuffled data. This particular screen took 5 months of computer time using 70 CPUs (Pentium IV, 2.4 GHz) with an earlier version of FOLDALIGN [61]. The current version of FOLDALIGN is estimated to be 20 times faster (i.e. 1 week on similar hardware), largely because of its ability to discard poor intermediate alignments, but even after this improvement in speed, an all-to-all comparison of two mammalian genomes remains highly impractical. Other applications of FOLDALIGN for genomic sequence analyses can be found at: <http://genome.ku.dk/resources/rsgs>.

#### Employing local multiple structural alignments for screening

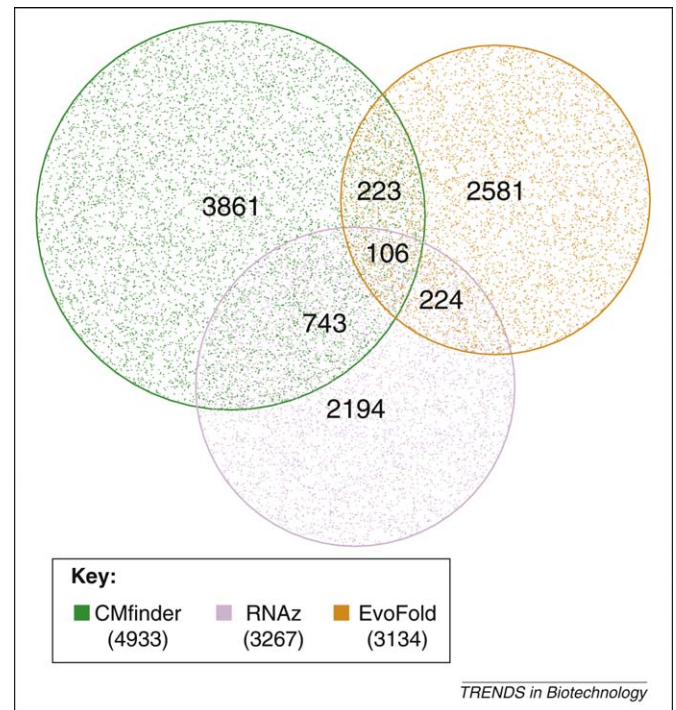
Using more than two organisms will potentially yield more accurate predictions. Extending the FOLDALIGN pairwise screen, CMfinder [62] was applied to the 44 ENCODE regions [3], which cover 1% of the human genome. CMfinder discovers structured RNA motifs in a set of unaligned sequences. It builds local alignments, that is, flanking regions that do not appear to contain RNA structures are ignored. Furthermore, entire sequences that do not appear to contain RNA structures are also ignored, which is valuable because the phylogenetic range of a given ncRNA is typically not known in advance. CMfinder builds an initial heuristic local alignment using sequence- and energy-based seed structures from which it derives a covariance model [20,22,63] (a specialized SCFG). The alignment and model are then refined jointly, using an approach akin to that described in Ref. [22], but extended to local alignments.

<sup>1</sup> <http://genome.ucsc.edu/cgi-bin/hgTracks?org=human&db=hg18&position=chr3:180369542-180369721&rnaGene=pack&ensGene=pack>. This human sequence is annotated as a snoRNA pseudogene in the UCSC browser, although both it and its cow ortholog are high-scoring members of the manually annotated Rfam “seed” alignment for this family, and the other four sequences are included in the automatically generated Rfam “full” alignment.

Although the approach is free of any constraints imposed by fixed window sizes and pre-computed alignments, to control run time, there are some *a priori*, user-specified restrictions with regard to the complexities of motifs (e.g. number of hairpins) that can be discovered. In the CMfinder ENCODE screen [64], pre-computed MULTIZ sequence alignments for the ENCODE regions were used, but, in contrast to the align-first methods discussed earlier, they have been used only to indicate orthology, and not alignment at the nucleotide level. Repetitive sequences have been included, but exons and PhastCons [65] elements were omitted. The average sequence identity of the 56 017 input alignments was 50%, with an average sequence length of 155 nucleotides. The screen resulted in 6587 ncRNA candidates, with an estimated FDR of 50% [64]. About 60% of the candidates were located in non-coding parts of protein-coding genes. Of these, 83% were located in introns, 14% in 5' untranslated regions (UTRs) and the remainder in 3' UTRs, an interesting bias perhaps caused by the poor annotation of UTRs. A small number of candidates have been tested experimentally. Most of these show tissue-specific expression in humans. For one 67-nucleotide candidate in a 4-kb intron of the neuron-specific synapsin 3 gene, northern blot analysis has confirmed expression of a 2.8-kb ncRNA in the brain [64].

One key observation from this study was that the lower the identity of the input sequences, the more they were realigned, that is, the structure-driven alignments produced by CMfinder differed more from the MULTIZ sequence-based alignments as sequence identity declined. Although this trend is not surprising, its extent suggests that standard alignments are misleading unless well-conserved in primary sequence. The issue is significant because 25% of candidates had more than 50% of their positions realigned when compared to the original MULTIZ alignments. Furthermore, many regions with identities above 70% were realigned, and even small realignments can have an impact, as illustrated in Figure 5.

The RNAZ, EvoFold and CMfinder screens processed slightly different subsets of the ENCODE regions, but their predictions on the common subset can be compared directly [47]. For example, 330 genomic loci were predicted by RNAZ and EvoFold. This degree of overlap, as well as the other overlaps quantified in Figure 6, is highly unlikely to have occurred purely by chance, given the lengths of the predictions in comparison to the total size of the ENCODE regions. Nevertheless, this numerically modest level of concordance among methods appears somewhat disappointing. However, further investigation has shown that this was mainly to the result of different sensitivities in different regimes of sequence composition and sequence similarity [47]. Although EvoFold candidates primarily lie in highly conserved, AU-rich regions, RNAZ (and CMfinder) candidates tend to be slightly GC-rich and exhibit more sequence variation. We speculate that RNAZ and EvoFold are missing candidates because of their reliance on sequence-based alignments, particularly in diverged regions, whereas CMfinder favors such regions, but ignores evolutionary relationships exploited by EvoFold, and energetically stable candidates that lack strong covariance



**Figure 6.** Comparison of ENCODE scans. The Venn diagram compares RNA elements in the ENCODE regions predicted by the three screening methods discussed here, CMfinder, RNAz, and EvoFold. Only predictions from input data that are common to all three studies are tallied, that is, repeat, exon and PhastCons regions are excluded [64]. As illustrated here, of the 4933 candidates reported by CMfinder, 3861 were reported by neither of the other methods, while 106 candidates were reported by all three. Adapted with permission from [Ref. 64].

patterns favored by RNAz. Thus, overall, these methods can be considered as complementary, with each of them able to predict a fraction of the ~20 000 unique candidate loci found in total in the ENCODE regions.

Applications of CMfinder to prokaryotes (emphasizing riboswitch discovery) are also listed at <http://genome.ku.dk/resources/rsgs>.

We eagerly await future development of methods for local structural RNA alignments. For example, an extension of SCARNA [57], called SCARNA\_LM, has been published recently [66]. It creates local multiple alignments using the same principles of comparing fixed-size stem fragments that are separated into their 5' and 3' parts, respectively. We hope that other methods will follow as well.

### Conclusions and perspectives

To date, annotation in the genome databases relates almost exclusively to protein coding genes. By contrast, the computational screens for ncRNAs described here suggest that a large number of functional ncRNAs remain to be found. This is consistent with the observation that most of the non-coding mammalian genome is transcribed. Computational *de novo* discovery of ncRNAs within genomic sequences is still in its infancy. Nevertheless, these screens provide a valuable starting point for subsequent studies of specific genomic regions and their functional characterization.

*In silico* screens are useful as the information generated can be correlated with other studies, or to obtain a more general picture of the putative RNA structure content. Currently, screens essentially only address RNA secondary



structure, with the numerous limitations discussed above, which leaves much room for improvement. Another major issue is their high computational cost (i.e. their time and memory requirements). Their high false discovery rates also make them inappropriate for systematic annotation and necessitate experimental follow-up. There is clearly a demand for faster, better computational approaches. Nonetheless, present methods have demonstrated their value and have opened an entirely new branch within comparative genomics. For instance, simple extrapolation from the ENCODE studies suggests prediction of ~2 million RNA structures within the human genome, albeit potentially with many structures per transcript, high FDR and other caveats discussed above. Computational methods should, however, be developed beyond searching for local RNA structures, in part, because some ncRNAs might have no structure (e.g. piRNAs [67]), or well-defined structures on only parts of their sequence (e.g. Evf-1 [68,69]; see also Ref.[70]). Additionally, ncRNAs are known to vary considerably in size, ranging from ~20 to ~100 000 nucleotides, which contributes significantly to the challenge at hand. Fusion transcripts also pose a serious computational challenge that has not yet been addressed. An example of such a transcript (the cDNA KIAA0510) that involves segments from different chromosomes was reported a few years ago [71], a finding echoed by the ENCODE project [3]. Experimental approaches continue to lead to novel discoveries. The lincRNAs [72] are one recent example. They were identified by a combination of CHIP-Seq and custom tiling array technologies, and include 1600 large, multiexon RNAs that are expressed in multiple mouse cell types and show evidence of purifying selection. Information obtained from such studies can in turn be incorporated into computational methods, which then might lead to further discoveries of novel ncRNAs. Methods to predict mRNA-like ncRNAs promise to contribute to this [73]. Another challenge is to unravel the significance of conserved structures in UTRs of protein-coding genes, which constitute a sizeable fraction of predictions. For example, CMfinder candidate r-6354 ([64], see [http://genome.ku.dk/resources/cm\\_f\\_encode/pages/candidate.php?id=r-6354](http://genome.ku.dk/resources/cm_f_encode/pages/candidate.php?id=r-6354)) which is found in a 5' UTR in several mammals, might have a role in regulating its own mRNA, but is also a predicted microRNA precursor. In this situation, comparison with known elements in UTR databases and/or attempting *de novo* motif finding in UTRs are plausible approaches [74,75].

Putative ncRNAs are available from various sources, including the UCSC browser [60] and (with some degree of overlap) RNAdb [76]. In our online supplement at <http://genome.ku.dk/resources/rsgs>, we provide links to files that define the genomic locations of many of the ncRNAs candidates described in this paper. These files may be passed directly to the UCSC genome browser, so that the candidates may be viewed in context with other genomic features, such as gene predictions and sequence-based multispecies conservation information.

Novel strategies are also emerging for post-clustering of predicted RNA structures. In particular, LocARNA [32], a Sankoff-based approach for structural alignment, has grouped structurally related ncRNAs found by RNAz in a screen of *Ciona intestinalis* and *Ciona savignyi*. The cluster-

ing has identified known RNA families, such as tRNAs, microRNAs, and spliceosomal RNAs, and has suggested additional groups that consist of novel, so-far unannotated RNAs. Moreover, clustering based on FoldalignM [77] has detected novel relationships among microRNA family members [78]. Another clustering approach that uses RNAforester has suggested some of the structural characteristics that distinguish microRNA precursors from the large number of other transcribed RNA stem-loops in mammalian genomes [79].

The increased availability of genome sequences has exposed sufficient nucleotide variation to allow the detection of patterns of RNA structure in sequence-based alignments, but has also revealed that sequence-based alignments alone might not be sufficient for detection of ncRNAs. The computational screens performed so far have pushed the limits of comparative genomics and pointed (perhaps unsurprisingly) towards structural alignments as an essential co-requisite for detection of ncRNAs. As we have seen, this is not only the case for those orthologous genomic regions whose sequences are poorly conserved across related organisms, but might also be true for small portions of otherwise well-conserved regions. This in turn poses new challenges in addressing comparative genomics from an RNA perspective. Furthermore, despite remarkable progress in computer hardware and algorithms, even faster tools are highly desirable. With this in mind, tools for integrating new genomic sequences with existing *de novo* ncRNA predictions would be valuable, both for annotating the new genome and for strengthening or dismissing previously marginal predictions. These predictions only need to be carried out in limited regions, therefore, they should be fast.

Even though further advances in methods for predicting RNA structures within genomic sequences are expected, the field is also developing tools to predict RNA interactions, with many utilizing the same principles as for RNA folding. The ability to predict RNA interactions will be an essential component in assigning a putative functional context for predicted RNA structures in the future.

#### Note added in proof

See Kavanaugh [84] for a recent success with the single-sequence folding energy approach, as well as some investigation of the subtleties inherent in sliding window approaches.

#### Acknowledgements

We thank the anonymous referees for numerous helpful suggestions. This work was supported in part by the Danish Research Council for Technology and Production, The Lundbeck Foundation and Danish Center for Scientific Computation, NIEHS Grant P30ES07033, and Austrian GEN-AU project "Regulatory ncRNAs."

#### References

- 1 Eddy, S.R. (2001) Non-coding RNA genes and the modern RNA world. *Nat. Rev. Genet.* 2, 919–929
- 2 International Human Genome Sequencing Consortium (2004) Finishing the euchromatic sequence of the human genome. *Nature* 431, 931–945
- 3 ENCODE Project Consortium (2004) The ENCODE (ENCyclopedia Of DNA Elements) Project. *Science*, 306, 636–640

- 4 Mattick, J.S. (2009) The genetic signatures of noncoding RNAs. *PLoS Genetics* 5, e1000459
- 5 Eddy, S.R. (2002) Computational genomics of noncoding RNA genes. *Cell* 109, 137–140
- 6 Rivas, E. and Eddy, S.R. (2000) Secondary structure alone is generally not statistically significant for the detection of noncoding RNAs. *Bioinformatics* 16, 583–605
- 7 Workman, C.T. and Krogh, A. (1999) No evidence that mRNAs have lower folding free energies than random sequences with the same dinucleotide distribution. *Nucleic Acids Res.* 27, 4816–4822
- 8 Bonnet, E. *et al.* (2004). Evidence that microRNA precursors, unlike other non-coding RNAs, have lower folding free energies than random sequences. *Bioinformatics* 20, 2911–2917
- 9 Lindow, M. and Gorodkin, J. (2007) Principles and limitations of computational microRNA gene and target finding. *DNA Cell Biol.* 26, 339–351
- 10 Pace, N.R. *et al.* (1999) Probing RNA structure, function, and history by comparative analysis, In *The RNA World*, (2nd edition), CSHL Press, pp. 113–141
- 11 Cloonan, N. and Grimmond, S. (2008) Transcriptome content and dynamics at single-nucleotide resolution. *Genome Biol.* 9, 234
- 12 Nussinov, R. and Jacobson, A.B. (1980) Fast algorithm for predicting the secondary structure of single-stranded RNA. *Proc. Natl. Acad. Sci. U. S. A.* 77, 6309–6313
- 13 Eddy, S.R. (2004) How do RNA folding algorithms work? *Nat. Biotechnol.* 22, 1457–1458
- 14 Klein, R.J. *et al.* (2002) Noncoding RNA genes identified in AT-rich hyperthermophiles. *Proc. Natl. Acad. Sci. U. S. A.* 99, 7542–7547
- 15 Schattner, P. (2002) Searching for RNA genes using base-composition statistics. *Nucleic Acids Res.* 30, 2076–2082
- 16 Larsson, P. *et al.* (2008) *De novo* search for non-coding RNA genes in the AT-rich genome of *Dictyostelium discoideum*: performance of Markov-dependent genome feature scoring. *Genome Res.* 18, 888–899
- 17 Menzel, P. *et al.* (2009) The tedious task of finding homologous non-coding RNA genes. *RNA* 15 (12), doi:10.1261/rna.1556009
- 18 Zuker, M. *et al.* (1999) Algorithms and thermodynamics for RNA secondary structure prediction: A practical guide. In *RNA Biochemistry and Biotechnology (NATO ASI Series, Vol. 70)* (Barciszewski, J. and Clark, B.F.C., eds), pp. 11–43, Kluwer Academic Publishers
- 19 Hofacker, I.L. *et al.* (1994) Fast folding and comparison of RNA secondary structures (The Vienna RNA Package). *Monatshfte für Chemie (Chemical Monthly)* 125, 167–188
- 20 Durbin, R. *et al.* (1998) *Biological Sequence Analysis: Probabilistic Models of Proteins and Nucleic Acids*, Cambridge University Press
- 21 Dowell, R.D. and Eddy, S.R. (2004) Evaluation of several lightweight stochastic context-free grammars for RNA secondary structure prediction. *BMC Bioinformatics* 5, 71
- 22 Eddy, S.R. and Durbin, R. (1994) RNA sequence analysis using covariance models. *Nucleic Acids Res.* 22, 2079–2088
- 23 Sakakibara, Y. *et al.* (1994) Stochastic context-free grammars for tRNA modeling. *Nucleic Acids Res.* 22, 5112–5120
- 24 Andronescu, M. *et al.* (2007) Efficient parameter estimation for RNA secondary structure prediction. *Bioinformatics* 23, i19–28
- 25 Gardner, P.P. and Giegerich, R. (2004) A comprehensive comparison of comparative RNA structure prediction approaches. *BMC Bioinformatics* 5, 140
- 26 Sankoff, D.D. (1985) Simultaneous solution of the RNA folding, alignment and protosequence problems. *SIAM J. Appl. Math.* 45, 810–825
- 27 Seemann, S. *et al.* (2008) Unifying evolutionary and thermodynamic information for RNA folding of multiple alignments. *Nucleic Acids Res.* 36, 6355–6362
- 28 Janssen, S. *et al.* (2008) Shape based indexing for faster search of RNA family databases. *BMC Bioinformatics* 9, 131
- 29 Reeder, J. *et al.* (2007) Locomotif: from graphical motif description to RNA motif search. *Bioinformatics* 23, i392–400
- 30 Marquez, S.M. *et al.* (2005) Structural implications of novel diversity in eucaryal RNase P RNA. *RNA* 11, 739–751
- 31 Nawrocki, E.P. *et al.* (2009) Infernal 1.0: inference of RNA alignments. *Bioinformatics* 25, 1335–1337
- 32 Will, S. *et al.* (2007) Inferring noncoding RNA families and classes by means of genome-scale structure-based clustering. *PLoS Comput. Biol.* 3, e65
- 33 Rivas, E. and Eddy, S.R. (2001) Noncoding RNA gene detection using comparative sequence analysis. *BMC Bioinformatics* 2, 8
- 34 Altschul, S.F. *et al.* (1990) Basic local alignment search tool. *J. Mol. Biol.* 215, 403–410
- 35 Rivas, E. *et al.* (2001) Computational identification of noncoding RNAs in *E. coli* by comparative genomics. *Curr. Biol.* 11, 1369–1373
- 36 McCutcheon, J.P. and Eddy, S.R. (2003) Computational identification of non-coding RNAs in *Saccharomyces cerevisiae* by comparative genomics. *Nucleic Acids Res.* 31, 4119–4128
- 37 Clote, P. *et al.* (2005) Structural RNA has lower folding energy than random RNA of the same dinucleotide frequency. *RNA* 11, 578–591
- 38 Washietl, S. and Hofacker, I.L. (2004) Consensus folding of aligned sequences as a new measure for the detection of functional RNAs by comparative genomics. *J. Mol. Biol.* 342, 19–30
- 39 Washietl, S. *et al.* (2005) Fast and reliable prediction of noncoding RNAs. *Proc. Natl. Acad. Sci. U. S. A.* 102, 2454–2459
- 40 Washietl, S. *et al.* (2005) Mapping of conserved RNA secondary structures predicts thousands of functional noncoding RNAs in the human genome. *Nat. Biotechnol.* 23, 1383–1390
- 41 Knudsen, B. and Hein, J. (1999) RNA secondary structure prediction using stochastic context-free grammars and evolutionary history. *Bioinformatics* 15, 446–454
- 42 Knudsen, B. and Hein, J. (2003) Pfold: RNA secondary structure prediction using stochastic context-free grammars. *Nucleic Acids Res.* 31, 3423–3428
- 43 Pedersen, J.S. *et al.* (2006) Identification and classification of conserved RNA secondary structures in the human genome. *PLoS Comput. Biol.* 2, e33
- 44 Blanchette, M. *et al.* (2004) Aligning multiple genomic sequences with the threaded blockset aligner. *Genome Res.* 14, 708–715
- 45 Pollard, K.S. *et al.* (2006) An RNA gene expressed during cortical development evolved rapidly in humans. *Nature* 443, 167–172
- 46 Seemann, S.E. *et al.* (2007) Detection of RNA structures in porcine EST data and related mammals. *BMC Genomics* 8, 316
- 47 Washietl, S. *et al.* (2007) Structured RNAs in the ENCODE selected regions of the human genome. *Genome Res.* 17, 852–864
- 48 Babak, T. *et al.* (2007) Considerations in the identification of functional RNA structural elements in genomic alignments. *BMC Bioinformatics* 8, 33
- 49 Gesell, T. and Washietl, S. (2008) Dinucleotide controlled null models for comparative RNA gene prediction. *BMC Bioinformatics* 9, 248
- 50 Anandam, P. *et al.* (2009) Multiperm: shuffling multiple sequence alignments while approximately preserving dinucleotide frequencies. *Bioinformatics* 25, 668–669
- 51 Gorodkin, J. *et al.* (1997) Finding the most significant common sequence and structure motifs in a set of RNA sequences. *Nucleic Acids Res.* 25, 3724–3732
- 52 Havgaard, J.H. *et al.* (2007) Fast pairwise structural RNA alignments by pruning of the dynamical programming matrix. *PLoS Comput. Biol.* 3, 1896–1908
- 53 Mathews, D.H. and Turner, D.H. (2002) Dynalign: an algorithm for finding the secondary structure common to two RNA sequences. *J. Mol. Biol.* 317, 191–203
- 54 Uzilov, A.V. *et al.* (2006) Detection of non-coding RNAs on the basis of predicted secondary structure formation free energy change. *BMC Bioinformatics* 7, 173
- 55 Dowell, R.D. and Eddy, S.R. (2006) Efficient pairwise RNA structure prediction and alignment using sequence alignment constraints. *BMC Bioinformatics* 7, 400
- 56 Holmes, I. (2005) Accelerated probabilistic inference of RNA structure evolution. *BMC Bioinformatics* 6, 73
- 57 Tabei, Y. *et al.* (2006) SCARNA: fast and accurate structural alignment of RNA sequences by matching fixed-length stem fragments. *Bioinformatics* 22, 1723–1729
- 58 Kurtz, S. *et al.* (2004) Versatile and open software for comparing large genomes. *Genome Biol.* 5, R12
- 59 Torarinsson, E. *et al.* (2006) Thousands of corresponding human and mouse genomic regions unalignable in primary sequence contain common RNA structure. *Genome Res.* 16, 885–889. Erratum: *Genome Res.* 16, 1439
- 60 Kuhn, R.M. *et al.* (2009) The UCSC Genome Browser Database: update 2009. *Nucleic Acids Res.* 37 (Database issue), D755–761

- 61 Havgaard, J.H. *et al.* (2005) Pairwise local structural alignment of RNA sequences with sequence similarity less than 40%. *Bioinformatics* 21, 1815–1824
- 62 Yao, Z. *et al.* (2006) CMfinder—a covariance model based RNA motif finding algorithm. *Bioinformatics* 22, 445–452
- 63 Eddy, S.R. (2002) A memory-efficient dynamic programming algorithm for optimal alignment of a sequence to an RNA secondary structure. *BMC Bioinformatics* 3, 18
- 64 Torarinsson, E. *et al.* (2008) Comparative genomics beyond sequence-based alignments: RNA structures in the ENCODE regions. *Genome Res.* 18, 242–251
- 65 Siepel, A. *et al.* (2005) Evolutionarily conserved elements in vertebrate, insect, worm, and yeast genomes. *Genome Res.* 15, 1034–1050
- 66 Tabei, Y. and Asai, K. (2009) A local multiple alignment method for detection of non-coding RNA sequences. *Bioinformatics* 25, 1498–1505
- 67 Lau, N.C. *et al.* (2006) Characterization of the piRNA complex from rat testes. *Science* 313, 363–367
- 68 Faedo, A. *et al.* (2004) Identification and characterization of a novel transcript down-regulated in Dlx1/Dlx2 and up-regulated in Pax6 mutant telencephalon. *Dev. Dyn.* 231, 614–620
- 69 Kohtz, J.D. and Fishell, G. (2004) Developmental regulation of EVF-1, a novel non-coding RNA transcribed upstream of the mouse Dlx6 gene. *Gene Expr. Patterns* 4, 407–412
- 70 Athanasius F. Bompfünewerer Consortium *et al.* (2007) RNAs everywhere: genome-wide annotation of structured RNAs. *J. Exp. Zool. B Mol. Dev. Evol.* 308, 1–25
- 71 Claverie, J.M. (2005) Fewer genes, more noncoding RNA. *Science* 309, 1529–1530
- 72 Guttman, M. *et al.* (2009) Chromatin signature reveals over a thousand highly conserved large non-coding RNAs in mammals. *Nature* 458, 223–227
- 73 Hiller, M. *et al.* (2009) Conserved introns reveal novel transcripts in *Drosophila melanogaster*. *Genome Res.* 19, 1289–1300
- 74 Pavesi, G. *et al.* (2004) RNAProfile: an algorithm for finding conserved secondary structure motifs in unaligned RNA sequences. *Nucleic Acids Res.* 32, 3258–3269
- 75 Mignone, F. *et al.* (2005) UTRdb and UTRsite: a collection of sequences and regulatory motifs of the untranslated regions of eukaryotic mRNAs. *Nucleic Acids Res.* 33 (Database issue), D141–146
- 76 Pang, K.C. *et al.* (2007) RNAdb 2.0—an expanded database of mammalian non-coding RNAs. *Nucleic Acids Res.* 35 (Database issue), D178–182
- 77 Torarinsson, E. *et al.* (2007) Multiple structural alignment and clustering of RNA sequences. *Bioinformatics* 15, 926–932
- 78 Kaczkowski, B. *et al.* (2009) Structural profiles of human miRNA families from pairwise clustering. *Bioinformatics* 25, 291–294
- 79 Ritchie, W. *et al.* (2007) RNA stem-loops: to be or not to be cleaved by RNase III. *RNA* 13, 457–462
- 80 Gorodkin, J. and Knudsen, B. (2000) RNA Informatik. *Naturens Verden* 11–12, 2–9
- 81 Sprinzl, M. and Vassilenko, K.S. (2005) Compilation of tRNA sequences and sequences of tRNA genes. *Nucleic Acids Res.* 33 (Database issue), D139–140
- 82 Gardner, P.P. *et al.* (2009) Rfam: updates to the RNA families database. *Nucleic Acids Res.* 37 (Database issue), D136–140
- 83 Yao, Z. *et al.* (2007) A computational pipeline for high-throughput discovery of cis-regulatory noncoding RNA in prokaryotes. *PLoS Comput. Biol.* 3, e126
- 84 Kavanaugh, L.A. *et al.* (2009) Non-coding RNA prediction and verification in *Saccharomyces cerevisiae*. *PLoS Genet* 5, e1000321 DOI: 10.1371/journal.pgen.1000321