2011 IEEE International Conference on Bioinformatics and Biomedicine

Inferring Transcript Phylogenies

Yann Christinat and Bernard M.E. Moret
Laboratory of Computational Biology and Bioinformatics
EPFL (Ecole Polytechnique Fédérale de Lausanne)
1015 Lausanne, Switzerland
Email: {yann.christinat,bernard.moret}@epfl.ch

***Abstract*—Alternative splicing, an unknown mechanism 20 years ago, is now recognized as a major mechanism for proteome and transcriptome diversity, particularly in mammals—some researchers conjecture that up to 90% of human genes are alternatively spliced. Despite much research on exon and intron evolution, little is known about the evolution of transcripts.**

In this paper, we present a model of transcript evolution and an associated algorithm to reconstruct transcript phylogenies. The evolution of the gene structure—exons and introns—is used as basis for the reconstruction of transcript phylogenies. We apply our model and reconstruction algorithm on two well-studied genes, MAG and PAX6, obtaining results consistent with current knowledge and thereby providing evidence that a phylogenetic analysis of transcripts is feasible and likely to be informative.

***Keywords*-alternative splicing; transcript; evolution;**

I. Introduction

Alternative splicing is a mechanism to produce different proteins from the same gene—the end of the paradigm "one gene, one protein." In many genomes, several, or even most, genes are split into pieces called exons, separated by regions called introns, and a splicing mechanism takes the transcribed string of exons and introns, removes the introns, and splices the exons to form a single continuous string that is then translated into a protein. In alternative splicing, a mechanism underestimated until 1990, the splicing produces variants in which some of the exons can be omitted (and occasionally even some of the introns retained), thereby causing different proteins to be produced. Alternative splicing exists to some degree in most eukaryotes, but is most frequent in the more complex lineages. Thus it is present, but limited, in plants and fungi, but quite common in vertebrates—some researchers have conjectured that up to 90% of the human genes are alternatively spliced [1]–[3]. Alternative splicing is now recognized as a major mechanism for proteome and transcriptome diversity [4], [5].

The implications of this shift in paradigm are significant. The basic model of transcriptome evolution—DNA modification at the gene level and alternative transcription start sites—is incomplete: any modification that affects the splicing mechanism has to be considered. However, while evolution of DNA at the gene level has been the subject of intense scrutiny for decades, very little is known regarding the changes in the splicing products of alternative splicing. Thus there is a need to define a model of evolution for transcripts, not at the nucleotide level, but at the splicing level—which exons (and introns) are included, which excluded?

A better understanding of the relationships among different transcripts would benefit annotation transfer. Different proteins from one gene may have different functions, may be localized to certain tissues, or may be present at different developmental stages. Knowledge of their evolution would help in assessing the function of their homologues. Transcript phylogenies would also contribute to next-gen sequencing methods, especially RNA-seq. For instance, the "DREAM6 Alternative Splicing Challenge" asks to reconstruct alternatively spliced mRNA transcripts from short mRNA-seq data without a reference genome, but using the transcriptomes of other organisms [6]. A transcript phylogeny would help in assessing the support value of a predicted transcript.

In this paper, we propose a model of transcript evolution and an associated algorithm to reconstruct transcript phylogenies.

II. Transcript Evolution

A. Background

Many studies have been published on the rate of exon insertion and deletion and on the statistics of different types of splicing, but few researchers so far have studied the evolution of transcripts [2]. Harr and Turner showed that most transcripts among *Mus* subspecies were novel [7]. Nurtdinov *et al.* compared the human and mouse transcriptomes and concluded that half of the genes give rise to species-specific isoforms and only three quarters of all isoforms are present in their orthologous genes [8]. Splicing is also affected by non-DNA events. Modification of the chromatin structure can yield changes in the expression of a given transcript and may even create a new transcript or silence an existing one [1]. Finally, a few groups studied the correlation between gene duplication and alternative splicing [9]–[11].

These studies indicate that alternative splicing is a fast-evolving mechanism and hint that most of the transcripts may be little more than evolutionary noise. These studies also indicate that groups of species share a significant number of transcripts, whose relationship can only be delineated with a more complete model.

B. A model of transcript evolution

In the description of alternative splicing, the simplest concepts are those of *constitutive* exons, which are part of

978-0-7695-4574-5/11 $26.00 © 2011 IEEE
DOI 10.1109/BIBM.2011.11

Table I
ALL EVOLUTIONARY CHANGES FOR TRANSCRIPTS FOR A GIVEN EXON AT THE GENE LEVEL. 1_C: CONSTITUTIVE EXON, 1_A: ALTERNATIVE EXON, 0: NO EXON

$\nearrow$	0
0	No transcript had this exon and none will have it.
1_A	Some transcripts had this exon and none will have it.
1_C	All transcripts had this exon and none will have it.

$\nearrow$	1_A
0	No transcript had this exon and some will have it.
1_A	Some transcripts had this exon and some will have it.
1_C	All transcripts had this exon and some will have it.

$\nearrow$	1_C
0	No transcript had this exon and all will have it.
1_A	Some transcripts had this exon and all will have it.
1_C	All transcripts had this exon and all will have it.

every transcript, and of *cassette* exons, which may or may not be present in any given transcript. In general, any exon that is not constitutive is called *alternative*. There exists other types of splicing mechanisms, of which alternative 3'- or 5'-sites and intron retention are the most frequently cited [1], [3], [5], [12], [13]. Note that the definition of a constitutive exon requires that all transcripts for a given gene be known. If, however, alternative splicing is closer to a biased random sampling from the space of all possible isoforms (so that every possible splice form is produced at some or other time), then there may be no such thing as a constitutive exon. As the debate on this issue continues and as our aim is to provide a model against which to test various hypotheses regarding transcript evolution, we develop a model in which we consider the existence of constitutive exons as a given.

We thus model a transcript as a subset of the gene exons. We model alternative 3'- or 5'-sites as constitutive exons with two or more internal states—each state encoding for one particular configuration. Finally, we assimilate intron retention to cassette exons. We model transcript evolution as a two-level process. The gene structure, viewed in terms of its collection of exons and introns, constitutes one level, while the collection of transcripts obtained from that structure constitutes the other level. Modification of the gene structure affects the transcriptome, but modification of the transcriptome does not affect the gene structure. Peng and Li [14] showed that the status of an exon, constitutive or alternative, is conserved through tandem exon duplication, a finding that hints at a model of evolution where the status of an exon is encoded at the gene level. Consequently we have three possible exon states in our model: absent, constitutive, or alternative. We assume that all transitions—birth, death, and mutation between constitutive and alternative—are equally likely.

In addition to the model of exon evolution at the gene level, transcripts can gain or lose exons. Table I sums up the possible evolutionary changes at the transcript level, given the evolution of a particular gene exon.

Finally we assume that a transcript can undergo a lethal mutation or be subject to regulation and disappear at any time. By extension, new transcripts may also be created at any time during evolution.

This model yields a forest of *transcript trees*, which represents the evolution from ancestral transcripts to observed transcripts. Each transcript tree is a subtree of the gene tree, since all transcripts arise from that gene family and, if they evolve, must evolve on the same tree. If a new transcript arises from an existing one, the new transcript will be considered as the root of a new transcript tree.

Our basic model uses a fixed cost for the creation of new transcripts. Of course, the basic model does not assume that transcripts are created *ab initio*; rather, it postulates a hidden relationship with an unknown ancestor. New transcripts arise from existing ones and thus are the result of evolutionary changes that may legitimately correspond to different costs. We use a fixed cost for simplicity and also because it leads to a very efficient pruning of the search space. We have designed and implemented an extended model in which the creation of a new transcript is dynamically assigned a cost that corresponds to its evolution from its closest ancestor (a maximum parsimony approach). However, the dynamic cost computation prevents good pruning and makes the problem intractable for even medium-sized instances, as discussed in Section IV-D.

For the reconstruction algorithm, we used a maximum parsimony approach, using Dollo's parsimony—that is, an exon cannot be created twice [15], [16].

III. RESULTS

The algorithm was tested on two well-studied genes to assess the correctness of the model on biological data. Further testing was done on simulated data to test the algorithm itself.

A. Results on the MAG gene

The Myelin-Associated Glycoprotein (MAG) is a neuronal transmembrane glycoprotein which acts both as a ligand for an axonal receptor and as a receptor for an axonal signal [17]. It has been extensively studied and due to its short length and limited alternative splicing, it makes a perfect candidate for testing our algorithm.

Two main isoforms are known in mammals: L-MAG and S-MAG. The S-MAG is created by the inclusion of the penultimate exon that creates an early stop codon. In rodents the second exon is also alternatively spliced and occurs in both the S- and L- forms, yielding four transcripts in total [18]. Two major MAG isoforms have been observed in both zebrafish and fugu: L-tail (exon 9, from the left, is skipped) and XL-tail. The retention of the eighth intron in the fugu fish yields a third form (S-tail), which is not observed in the zebrafish [19].

1) Data: Transcripts and exons for the five species were compiled from the literature [17]–[21]. The sequences and the gene tree, as shown in Figure 2-A, were obtained from the Ensembl database [22].

We concatenated the gene's exons and aligned the resulting sequences using Mauve [23]. Every exon either

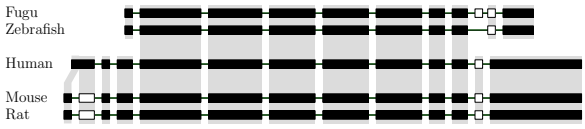

Figure 1. Orthologous exons of the MAG gene for the five species. A gray background indicates orthologous exons. Constitutive exons are shown in black.

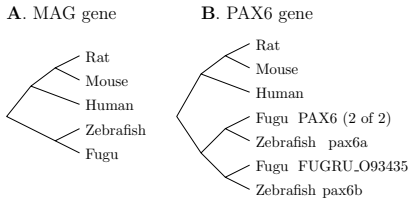

Figure 2. Gene trees for the MAG and the PAX6 gene.

was not aligned to any other exon or provided close to one hundred percent coverage of its ortholog. The only exception was the first human exon, which corresponds to the first two exons of the rodents. Such a situation might have posed a problem had the human exon been alternative, but fortunately it is a constitutive exon. The eighth intron of the fugu fish, which triggers an early stop codon, could not be aligned to any exon in any other species. Orthologous exons were then inferred from this alignment and are shown in Figure 1.

2) Results: We tested two setups. In the first experiment, we used a cost of infinity for exon gain/loss, $c_E = \infty$, whereas in the second we used a unit cost, $c_E = 1$. In both setups, the cost of transcript birth, c_B, and death, c_D, was set to a single parameter and varied. As shown in Figure 3, each experiment yielded solutions consistent with our biological knowledge of the isoforms. The S- and L-forms in mammals and the L- and XL-tail in fishes are clustered on their respective trees. However, the relationship between the fish and mammal isoforms is unclear. If the cost of exon gain is infinity, then the only relation between fishes and mammals is a link from the L-tail to the alternative L-form in rodents—but our model requires such a link, since it requires that all genes be connected. The same reasoning applies for $c_B = c_D \leq 2$ and $c_E = 1$. The cost of connecting a mammal transcript to a fish transcript is always greater than the cost of adding a new tree. The S-tail in the fugu fish is isolated and shows no relation to the S-form of mammals. Its distance to the mammal S-form is too great to allow the two to be clustered. There is no evidence that those two transcripts are biologically related, nor do their sequences align well—their only common feature is that both induce an early stop codon.

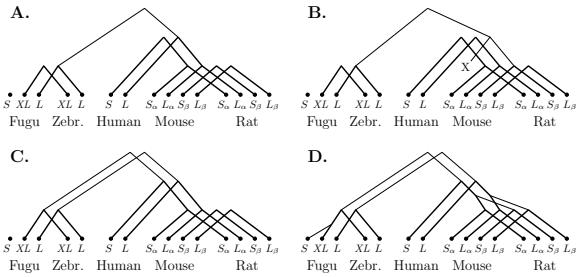

A. $c_B = c_D \leq 2$ and $c_E = 1$
B. $c_B = c_D > 0$ and $c_E = \infty$
C. $c_B = c_D > 2$ and $c_E = 1$
D. Extended model $c_D = 1$ and $c_E = 1$

Figure 3. Inferred transcript phylogenies for the MAG gene under different costs. Transcripts S_β and L_β in rodents represent the S- and L-forms in which exon 2 is skipped. Thicker trees contain similar transcripts. Only solutions with a minimal number of events are displayed.

3) Results with the extended model: Since the search space for the MAG instance is small, we were able to run the extended model on it. As seen in Figure 3-D, the result is the same as for the basic model using $c_B = c_D > 2$, except that newly created transcripts are linked to their closest ancestor.

B. Results on the PAX6 gene

The PAX6 gene is part of the well-studied paired box gene family (PAX), which encodes transcription factors for many developmental processes and is subject to heavy alternative splicing [24]–[26]—41 transcripts were found in a gene in the pigeon retina [27]. The canonical isoform is characterized by an N-terminal paired domain followed by a linker, a paired-type homeodomain, and a (P/S/T)-rich C-terminal domain, yielding a 422-amino-acid protein (437 in zebrafish). The gene undergoes alternative splicing and the best-known alternative isoform, +5a, differs from the canonical isoform by the inclusion of exon 5a. This 14-amino-acid insertion in the paired domain disrupts the DNA-binding ability of the N-terminal domain and enhances the binding of the C-terminal domain, thus creating a new set of interactions [28]. As can be seen in Figure 2-B, gene duplication occurred in the fish species leading to two PAX6 genes in the zebrafish and the fugu fish.

1) Transcripts and orthologous exons: Mammal transcripts were obtained from the Human-transcriptome Database for Alternative Splicing (H-DBAS) [29] and fish transcripts from the Ensembl database [22]. In the H-DBAS database, we considered only transcripts that were present in both the cDNA and mRNA databases, except in the case of *R. norvegicus*, where only the mRNA database was available. Similarly, with the Ensembl database, we used only transcripts that had CDS or UTR support. The gene tree was obtained from the Ensembl database and genes are thus named accordingly. The canonical and the +5a transcripts were identified through their protein product and the literature.

The literature on the PAX6 gene in the fugu fish is very sparse—we could find only one article, by Miles *et al.* [30], but that article does not corroborate the information in the Ensembl database. Thus we used the Ensembl data, as it is more recent, but we have no ground truth regarding the canonical or alternative isoforms.

The H-DBAS database conveniently indicates orthologous exons for the mouse, rat, and human. We ran an all-

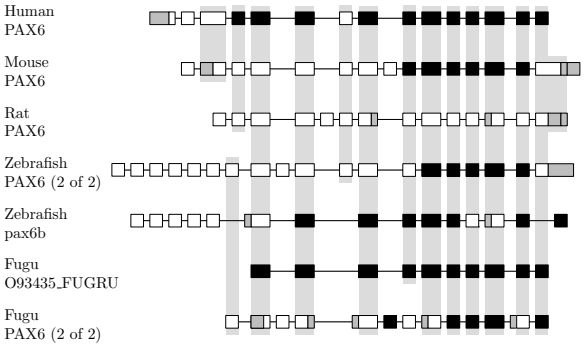

Figure 4. Orthologous exons of the PAX6 gene for the 7 genes of the 5 species. A gray background indicates orthologous exons. Constitutive exons are shown in black. Alternative 3'- and 5'-end are shown in gray. Note that only exons belonging to a transcript are shown.

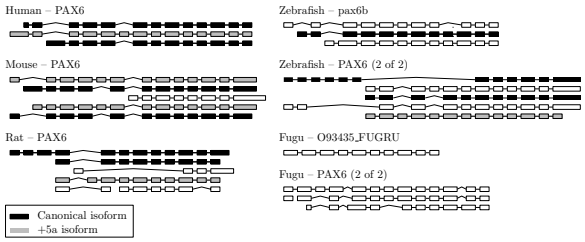

Figure 5. Transcripts for the PAX6 genes. The canonical isoforms are shown in black and the +5a isoforms in gray. Note that several transcripts are translated into the same isoform.

against-all semi-global alignment of all exons to confirm the mammalian orthologs and to obtain the orthologs for the two fishes. Orthologous exons are shown in Figure 4 and transcripts in Figure 5. In all species, we observe that several transcripts can produce the same isoform.

2) Results: As with the MAG gene, we tested two setups, with unit and infinite costs for exon gain or loss. However no solution could be found with $c_E = \infty$. Figure 6 reveals a correlation between c_B, c_D and the number of ancestral transcripts. A higher c_B affects the total number of trees. Any hypothesis should thus be tested under different parameters before drawing any conclusion. The best result uses $c_B = c_D = 5$, showing well-clustered isoforms within mammals. Note that the model imposes a link between all genes, so that the relevance of a single connection between fishes and mammals at low values of c_B is uncertain.

The number of solutions with minimum cost also increases along with c_B and c_D. The algorithm returns 36 solutions of equal cost for $c_B = c_D = 5$. However, if a random gene tree is used, this number gets about 10 times larger for the same parameters—a change that gives us confidence that phylogenetic information is indeed contained in the transcripts. Note that Figure 6 shows only solutions that have a minimum number of evolutionary events among solutions of minimum cost.

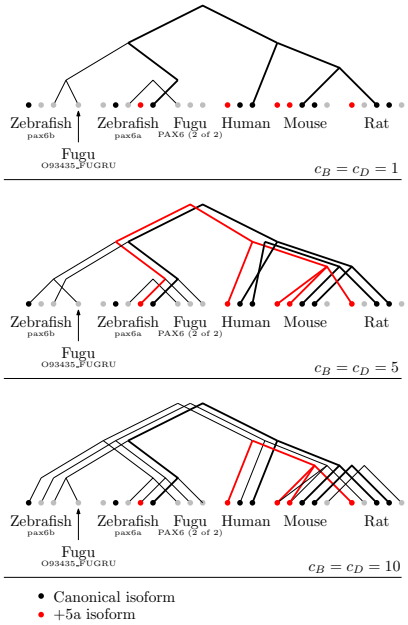

Figure 6. Transcript phylogenies for the PAX6 gene for different values of transcript birth, c_B, and death, c_D. Multiple solutions are superimposed. Thicker lines connect similar isoforms.

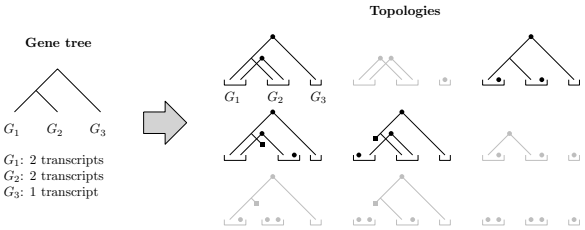

Figure 7. The 9 topologies corresponding to a trivial 3-gene example. In a topology, a dot represents a transcript birth and a square a transcript death. Grayed out topologies are not valid as some genes are unconnected.

C. Results on simulated data

In order to test the performance of the algorithm, we designed a simple scheme to generate transcripts with a given tree structure. Starting with n_T transcripts at the root and n_E random exons, each gene exon can either be born or die independently along the tree. The same evolutionary process applies to transcripts with exon gain or loss depending on the current set of gene exons.

The reconstruction algorithm works by splitting the search space into topologies—a topology being a forest of transcript trees whose leaves are not assigned. Figure 7 illustrates the topology space on a simple 3-gene example. The algorithm first explores the topology space rapidly then finds the best leaf assignment on good candidates. (More details are given in Section IV.) The search on the topology space is optimal, but under unfavorable circumstances may explore the entire search space. For a given number of genes, we tested the algorithm on *caterpillar* trees (trees where one of the two children is always a leaf) with a random number of ancestral transcripts. Caterpillar

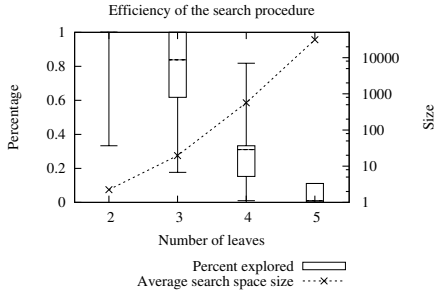

Figure 8. Percentage of topologies that get passed to the leaf assignment procedure. The average number of topologies per number of leaves is plotted on the second Y axis. Box plots and average sizes were computed on 1000 runs.

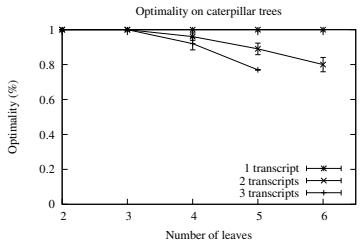

Figure 9. Optimality of the leaf assignment algorithm for caterpillar trees. Error bars show the standard deviation.

trees represent difficult instances since the depth of the tree is maximized for the number of leaves. Figure 8 shows that the percentage of topologies that get passed on to the leaf assignment procedure decreases quickly as the number of leaves increases. The size of the search space grows faster than exponential, but the search procedure reduces the growth rate of the number of refined topologies.

As the leaf assignment algorithm is not optimal, we tested how often it yielded the best solution. Given a gene tree and its associated transcripts, for every topology, every possible transcript assignment is generated. The best score is retained and tested against the solution proposed by the leaf assignment algorithm. We define optimality as the percentage of occurrences where the leaf assignment algorithm yielded the same score as the optimal solution during a single run. We performed one hundred runs, each with randomly generated hundred-exon genes, on difficult trees (caterpillar trees) and tested the optimality for different numbers of genes and of transcripts per gene. (The sizes of the instances are necessarily limited by the exhaustive search.) Figure 9 shows that, as expected, the optimality decreases as the number of leaves or transcripts increases.

We also tested the algorithm for scalability, since the running time grows faster than exponential in the worst case. Running the algorithm on the MAG gene takes a few seconds while running it on the PAX6 gene takes a few days. However our focus in this paper is to validate the concept of transcript phylogenies and show that reconstruction, within some limits, is possible. Past this step, we are confident that a heuristic can be designed to handle large problems with reasonable accuracy.

IV. Methods

The input of the algorithm is a gene tree with a set of leaf transcripts and orthologous exons. (Paralogous exons are considered as unrelated.)

The algorithm begins by reconstructing the state of the ancestral genes' exons—absent, alternative, or constitutive—using Sankoff's algorithm for the small parsimony problem [31], since exons can have more than 2 states. Without any further knowledge on exon evolution, we assumed that every transition has equal cost. A constraint is added to the algorithm to ensure that the result is consistent with Dollo parsimony—that an exon cannot be created more than once.

Transcript phylogenies are then reconstructed using a two-step algorithm. For each topology, a lower bound is computed. If the lower bound is higher than the minimum encountered so far, then the topology is discarded, since there could not exist a solution with this topology with a lower score than the current optimum. Otherwise the best solution for this topology is computed. We call this last step the *leaf assignment* step; it is the only part of the algorithm that makes use of the previously reconstructed evolution of the gene's exons.

Now, in order to prune the search space efficiently, we need to establish quickly a rather good solution. Since the algorithm explores the topology space in a deterministic, breadth-first search manner, it could, in the worst case, move from worst to best topology, improving the score at each step, and thus unable throughout the procedure to prune any part of the solution space. To make such a behavior extremely unlikely on any data, we establish initial solutions by randomly sampling the search space for each number of trees before the exploration of the search space starts and retaining the configuration with the lowest score as an upper bound.

When all topologies have been tested or scored, the algorithm returns all solutions of minimum cost. This process is described more precisely in Algorithm 1.

Our model makes no distinction between an event of zero cost and no event at all. Yet we would like to see only solutions that have the lowest number of events, so our algorithm uses (a version of) the number of events as a secondary criterion to rank the optimal solutions. For each tree in a given solution, we sum over all leaves the exons that are present in at least two leaves, then divide this value by the number of exons that are present in at least one leaf and by the number of leaves. The result is an index between 0 (all exons are unique to their leaves) and 1 (all leaves have the same exons).

A. Generating topologies

Topologies are generated with increasing numbers of trees. For each topology with t trees (t-topology), any edge removal yields a new topology with $t + 1$ trees. However, that process alone does not suffice to generate

Algorithm 1 Inference of transcript trees.

$S_{MIN} \leftarrow$ minimum score of the random walk
$i \leftarrow$ minimum number of trees
while $c_B \cdot i \leq S_{MIN}$ and $i \leq$ max number of trees **do**
  $TOPS \leftarrow$ List of all topologies with i trees
  $t \leftarrow$ first item in $TOPS$
  while t is not null **do**
    if LOWERBOUND(t) $\leq S_{MIN}$ **then**
      $S \rightarrow$ Score of the best leaf assignment
      if $S < S_{MIN}$ **then**
        $S_{MIN} \leftarrow S$
      end if
    end if
    $t \leftarrow$ next item in $TOPS$
  end while
  $i \leftarrow i + 1$
end while
return all solutions that have a score equal to S_{MIN}

all $(t + 1)$-topologies. Therefore, once that procedure has been applied to all t-topologies (and duplicates have been removed), we use branch-swapping to generate the remaining $(t + 1)$-topologies. A branch swap disconnects edge (n_1, p_1) from tree t_1 and edge (n_2, p_2) from tree t_2 and creates two new edges: (n_1, p_2) and (n_2, p_1). Here n_1 and n_2, and also p_1 and p_2, represent transcripts from the same ancestral gene. The algorithm again checks for duplicates, as it searches for all branch swaps on the set of $(t+1)$-topologies until no new topology can be generated.

B. Scoring solutions and topologies

The score of a particular solution is composed of two parts, the first reflecting the structure of the trees and the second, S_F, providing the parsimony score of the trees. The cost of creating or losing a transcript is a constant and thus we have

$$S = (c_B \cdot N_{tree} + c_D \cdot N_{death}) + S_F \qquad (1)$$

where N_{tree} is the number of trees, N_{death} the total number of transcript losses, and S_F the sum of the maximum parsimony scores of each tree. S_F is the only quantity that depends on the evolution of the gene's exons. c_B and c_D are parameters that control the cost of transcript birth and death.

A lower bound for topologies can be computed by considering the first part of the scoring function, $c_B \cdot N_{tree} + c_D \cdot N_{death}$. This value does not depend on the transcripts, but only on the topology. However, a better lower bound can be computed by adding a lower bound on the S_F score. For each tree, we compute the best leaf assignment as if all transcripts were available, corresponding to a topology with a single nontrivial tree. (In the real leaf assignment procedure, of course, trees compete for transcripts.) The sum of these values is a true lower bound.

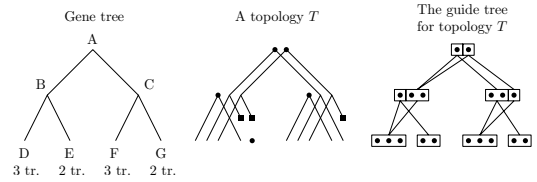

Figure 10. The guide tree for a given topology. **Topology**: There are two ancestral transcripts at A. Another transcript is created at B, two are lost between B and E, and a new transcript is created at E. A similar reasoning applies to C, F, and G. **Guide tree**: Boxes in the guide tree indicate a specific subtree. For instance, at B, we have only two kinds of subtrees: $B(D, E)$ and B$(D, -)$, thus we have two boxes. Within a box, each dot represents a transcript. At B we have two transcripts with subtree $B(D, -)$, thus there are two dots in the second box.

C. Leaf assignment procedure

Given a topology, leaf assignment remains challenging: given N genes and k transcripts per gene, a topology can lead to $k!^{N-1}$ possible leaf assignments. To tackle this problem, we combine a bottom-up dynamic programming algorithm with Sankoff's algorithm for the small parsimony problem.

We define a *state* as an ordered list of transcripts for a given gene. Each transcript t in a state has pointers to its left and right children, $l(t)$ and $r(t)$—if any. The ordering of the transcripts is the same in two states of the same gene, but the pointers change to reflect phylogenetic relationships. The number of transcripts of ancestral genes (inner nodes) is defined by the topology.

For each ancestral gene, every possible state is generated. If a gene has k_{LR} transcripts that have two children, k_L transcripts with a single left child, k_R transcripts with a single right child, and both children of the gene have n transcripts, then we have up to

$$\binom{n}{k_{LR}} \cdot \frac{n!}{(n-k_{LR})!} \cdot \binom{n-k_{LR}}{k_L} \cdot \binom{n-k_{LR}}{k_R} \qquad (2)$$

possible states. The product of the first two terms of Equation 2 is the number of possible assignments for the k_{LR} transcripts that have two children, while the last two terms compute the number of assignments for the transcripts that have only child. Since the first part selected k_{LR} elements, there remains only $n - k_{LR}$ elements to choose from.

However, this number is constrained by the topology: a transcript in some state cannot be connected to any transcript in its child's state—the subtrees have to match. We represent these constraints through a guide tree, which indicates the possible interactions for each transcript. Figure 10 illustrates the guide tree for an example topology. There could be up to 18 states at node A without the topology constraints, but these reduce the number down to just 4.

Our algorithm traverses the gene tree in postorder; at each node N, it computes the parsimony score of each state. Given a gene N, its parent P, and its two children

L and R, the score for a given state s of N is given by

$$S(s_N, s_P) = \min_{s_R \in R, s_L \in L} \{S(s_R, s_N) + S(s_L, s_N) + \sum_{t \in T} \text{minMP}(t) \text{ s.t. } T = roots(s_N, s_P)\} \quad (3)$$

where $roots(s_A, s_{parent})$ contains all transcripts of s_A that have no ancestor in s_{parent}. (If s_{parent} is null then it is the set of transcripts in s_A.)

minMP(t) is the parsimony score of transcript t as inferred by Sankoff's algorithm. A profile is built for each exon and the score of exon i in state u is computed by

$$t_{iu} = \min_{x \in E}\{c(u, x) + r_{ix}\} + \min_{y \in E}\{c(u, y) + l_{iy}\}$$

where l and r are the left and right children of transcript t, $c(a, b)$ is the cost of evolving from a to b and E is the set of all exon states. In our case we have $c(a, b) = 0$ for $a = b$ and $c(a, b) = 1$ otherwise. However the cost function must be slightly modified to account for exon evolution at the gene level. If a constitutive exon was gained or an exon was lost (at the gene level), then we set the cost of the change to zero. Additionally, if exon i is absent in the gene, then for all transcripts in the gene we have $t_{ix} = \infty, \forall x > 0$. Note that the left and right children of t depend on the choice of s_L and s_R. A similar equation can be derived in case of single-child transcripts. minMP(t) is then the sum over all exons of $\min_u t_{iu}$.

The quantities of 3 are built up during the postorder traversal; once the score of every state at the root of the gene tree has been computed, the minimum score is retained. Backtracking from the root to the leaves will then produce all optimal transcript phylogenies.

D. An extended model

The extended model sets a dynamic cost for transcript birth, but retains a constant cost for transcript death. Given a topology, the best leaf assignment is computed and backtracking is used to reconstruct the ancestral transcripts' sequences. Creation of new transcripts is assigned a cost that corresponds to its evolution from its closest ancestor. The birth cost is added only once the leaf assignment procedure has terminated and thus has no influence on the transcript assignment, except in case of multiple solutions, where only those that have a minimum birth cost will be selected.

Developing a good lower bound on the birth cost remains a challenge. This cost can vary between zero and the number of exons, so that simply using the lowest possible value would produce very loose bounds and thus be of no help in the search. (On simulated data and our two test genes, the search procedure using a zero value as a bound on the birth cost always had to look at every topology.)

V. Conclusion

In this study we addressed the lack of evolutionary model for alternative splicing by presenting a two-level model of transcript evolution and an algorithm to reconstruct transcript phylogenies. Our work opens the door for the study of transcript evolution, as it provides tools for testing evolutionary hypotheses.

We presented two models. The basic model uses a fixed cost for the creation of new transcripts—an unrealistic assumption, but one that greatly decreases the computational cost. The extended model assigns a cost dynamically by finding the closest (least cost) ancestor; however, the dynamic nature of the cost defeats our pruning strategy and the problem became intractable for medium-sized instances.

Results on a well-studied gene, MAG, showed that the extended model yielded results similar to those obtained with the basic model. Good clustering of known isoforms was achieved with the basic model for both gene families (MAG and PAX6) we studied.

Future work involves a faster version of the algorithm, and eventually approximation methods to enable us to use the extended model on large problems.

Acknowledgments

We would like to thank Prof. Marc Robinson-Réchavi and Dr. Eyal Privman for fruitful discussions that helped us to make our model more biologically plausible.

References

[1] H. Keren, G. Lev-Maor, and G. Ast, "Alternative splicing and evolution: diversification, exon definition and function." *Nature reviews. Genetics*, vol. 11, no. 5, pp. 345–55, May 2010.

[2] I. I. Artamonova and M. S. Gelfand, "Comparative genomics and evolution of alternative splicing: the pessimists' science." *Chemical reviews*, vol. 107, no. 8, pp. 3407–30, Aug. 2007.

[3] E. Kim, A. Magen, and G. Ast, "Different levels of alternative splicing among eukaryotes." *Nucleic acids research*, vol. 35, no. 1, pp. 125–31, Jan. 2007.

[4] P. M. Harrison, A. Kumar, N. Lang, M. Snyder, and M. Gerstein, "A question of size : the eukaryotic proteome and the problems in defining it." *Nucleic acids research*, vol. 30, no. 5, pp. 1083–1090, 2002.

[5] B. Modrek and C. Lee, "A genomic view of alternative splicing." *Nature genetics*, vol. 30, no. 1, pp. 13–9, Jan. 2002.

[6] (2011, Jun.) DREAM6 Alternative Splicing Challenge. [Online]. Available: http://www.the-dream-project.org/challanges/dream6-alternative-splicing-challenge

[7] B. Harr and L. M. Turner, "Genome-wide analysis of alternative splicing evolution among Mus subspecies." *Molecular ecology*, vol. 19 Suppl 1, pp. 228–39, Mar. 2010.

[8] R. N. Nurtdinov, "Low conservation of alternative splicing patterns in the human and mouse genomes." *Human Molecular Genetics*, vol. 12, no. 11, pp. 1313–1320, Jun. 2003.

[9] J. Roux and M. Robinson-Rechavi, "Age-dependent gain of alternative splice forms and biased duplication explain the relation between splicing and duplication." *Genome Research*, vol. 21, no. 3, p. 357, 2011.

[10] Z. Su, J. Wang, J. Yu, X. Huang, and X. Gu, "Evolution of alternative splicing after gene duplication." *Genome research*, vol. 16, no. 2, pp. 182–9, Feb. 2006.

[11] N. M. Kopelman, D. Lancet, and I. Yanai, "Alternative splicing and gene duplication are inversely correlated evolutionary mechanisms." *Nature genetics*, vol. 37, no. 6, pp. 588–9, Jun. 2005.

[12] J. Lu *et al.*, "Alternative splicing in teleost fish genomes: same-species and cross-species analysis and comparisons." *Molecular genetics and genomics : MGG*, pp. 531–539, Apr. 2010.

[13] A. J. Matlin, F. Clark, and C. W. J. Smith, "Understanding alternative splicing: towards a cellular code." *Nature reviews. Molecular cell biology*, vol. 6, no. 5, pp. 386–98, May 2005.

[14] T. Peng and Y. Li, "Tandem exon duplication tends to propagate rather than to create de novo alternative splicing." *Biochemical and biophysical research communications*, vol. 383, no. 2, pp. 163–6, May 2009.

[15] L. Dollo, "Les lois de l'évolution." *Bull Soc Belge Geol Pal Hydr*, vol. 7, pp. 164–166, 1893.

[16] A. V. Alekseyenko, C. J. Lee, and M. a. Suchard, "Wagner and Dollo: a stochastic duet by composing two parsimonious solos." *Systematic biology*, vol. 57, no. 5, pp. 772–84, Oct. 2008.

[17] R. H. Quarles, "Myelin-associated glycoprotein (MAG): past, present and beyond." *Journal of neurochemistry*, vol. 100, no. 6, pp. 1431–48, Mar. 2007.

[18] C. Lai *et al.*, "Two forms of 1B236/myelin-associated glycoprotein, a cell adhesion molecule for postnatal neural development, are produced by alternative splicing." *Proceedings of the National Academy of Sciences of the United States of America*, vol. 84, no. 12, pp. 4337–41, Jun. 1987.

[19] F. Lehmann, H. Gäthje, S. r. Kelm, and F. Dietz, "Evolution of sialic acid-binding proteins: molecular cloning and expression of fish siglec-4." *Glycobiology*, vol. 14, no. 11, pp. 959–68, 2004.

[20] N. Fujita *et al.*, "Developmentally regulated alternative splicing of brain myelin-associated glycoprotein mRNA is lacking in the quaking mouse." *FEBS letters*, vol. 232, no. 2, pp. 323–7, May 1988.

[21] J. Georgiou, M. B. Tropak, and J. C. Roder, "Myelin-Associated Glycoprotein Gene." *Myelin Biology and Disorders*, vol. 1, pp. 421–467, 2004.

[22] P. Flicek *et al.*, "Ensembl's 10th year." *Nucleic acids research*, vol. 38, no. Database issue, pp. D557–62, Jan. 2010.

[23] A. E. Darling, B. Mau, and N. T. Perna, "progressiveMauve: multiple genome alignment with gene gain, loss and rearrangement." *PLoS ONE*, vol. 5, no. 6, p. e11147, 06 2010.

[24] L. Z. Holland and S. Short, "Alternative splicing in development and function of chordate endocrine systems: a focus on pax genes." *Integrative and Comparative Biology*, vol. 50, no. 1, pp. 22–34, May 2010.

[25] S. Short and L. Z. Holland, "The evolution of alternative splicing in the Pax family: the view from the Basal chordate amphioxus." *Journal of molecular evolution*, vol. 66, no. 6, pp. 605–20, Jun. 2008.

[26] S. Nornes *et al.*, "Zebrafish contains two pax6 genes involved in eye development." *Mechanisms of development*, vol. 77, no. 2, pp. 185–96, Oct. 1998.

[27] D. Bandah *et al.*, "A complex expression pattern of Pax6 in the pigeon retina." *Investigative ophthalmology & visual science*, vol. 48, no. 6, pp. 2503–9, Jun. 2007.

[28] J. A. Epstein *et al.*, "Two independent and interactive DNA-binding subdomains of the Pax6 paired domain are regulated by alternative splicing." *Genes & Development*, vol. 8, no. 17, pp. 2022–2034, Sep. 1994.

[29] J. Takeda *et al.*, "H-DBAS: human-transcriptome database for alternative splicing: update 2010." *Nucleic acids research*, vol. 38, no. Database issue, pp. D86–90, Jan. 2010.

[30] C. Miles, "Complete sequencing of the Fugu WAGR region from WT1 to PAX6: Dramatic compaction and conservation of synteny with human chromosome 11p13." *Proceedings of the National Academy of Sciences*, vol. 95, no. 22, pp. 13 068–13 072, Oct. 1998.

[31] D. Sankoff, "Minimal Mutation Trees of Sequences." *SIAM Journal on Applied Mathematics*, vol. 28, no. 1, pp. 35–42, 1975.