Recap of Neighbor Joining

Here is the distance matrix, with marginal sums:

	Bovin	Pig	Rat	Mouse	Human	
Bovin	0	0.098	0.197	0.204	0.173	0.672
Pig	0.098	0	0.181	0.189	0.163	0.631
Rat	0.197	0.181	0	0.072	0.176	0.626
Mouse	0.204	0.189	0.072	0	0.204	0.669
Human	0.173	0.163	0.176	0.204	0	0.716
	0.672	0.631	0.626	0.669	0.716	

1. Compute $S_{i,j}$ for each pair i, j of species

$$S_{i,j} = (N-2)D_{i,j} - \sum D_{i,k} - \sum D_{j,k}$$

   The sums are taken over *all* entries in the row (or column).

2. Select the pair that gives the least quantity and join them in the tree, and replace them in the matrix with a new entry u

Recap of Neighbor Joining

3. Fill the matrix with new values $D_{k,u}$ where

$$D_{k,u} = \frac{1}{2}(D_{i,k} + D_{j,k} - D_{i,j})$$

and set the branch lengths

$$L_{i,u} = \frac{1}{2(N-2)}[(N-2)D_{i,j} + \sum_k D_{i,k} - \sum_k D_{j,k}]$$

$$L_{j,u} = \frac{1}{2(N-2)}[(N-2)D_{i,j} + \sum_k D_{j,k} - \sum_k D_{i,k}]$$

(The homework used *D* for these, which was confusing.)

4. Iterate until the tree is complete

Recap of Neighbor Joining

Here is our example worked out:

$D_{i,j}$	Bovin	Pig	Rat	Mouse	Human	
Bovin	0	0.098	0.197	0.204	0.173	0.672
Pig	0.098	0	0.181	0.189	0.163	0.631
Rat	0.197	0.181	0	0.072	0.176	0.626
Mouse	0.204	0.189	0.072	0	0.204	0.669
Human	0.173	0.163	0.176	0.204	0	0.716
	0.672	0.631	0.626	0.669	0.716	

Compute the matrix of $S_{i,j}$ values.

$S_{i,j}$	Bovin	Pig	Rat	Mouse	Human
Bovin		-1.009	-0.707	-0.729	-0.869
Pig			-0.714	-0.733	-0.858
Rat				-1.079	-0.814
Mouse					-0.773
Human					

Select the least value. This is between rat and mouse

Branch Lengths:

$L_{R,RM} = (1/6)(0.072*3 + 0.626 - 0.669) = 0.0288$

$L_{M,RM} = (1/6)(0.072*3 - 0.626 + 0.669) = 0.0432$

Recap of Neighbor Joining

Compute a new distance matrix:

$D_{i,j}$	Bovin	Pig	RM	Human	
Bovin	0	0.098	0.1645	0.173	0.4355
Pig	0.098	0	0.149	0.163	0.41
RM	0.1645	0.149	0	0.154	0.4675
Human	0.173	0.163	0.154	0	0.49
	0.4355	0.41	0.4675	0.49	

Compute the matrix of $S_{i,j}$ values.

$S_{i,j}$	Bovin	Pig	RM	Human
Bovin		-0.6495	-0.574	-0.5795
Pig			-0.5795	-0.574
RM				-0.6495
Human				

Select the least value. There is a tie (can you see why there has to be?). We'll join Bovin with Pig.

$L_{B,BP} = (1/4)(0.098*2 + 0.4355 - 0.41) = 0.0554$
$L_{P,BP} = (1/4)(0.098*2 - 0.4355 + 0.41) = 0.0426$

Recap of Neighbor Joining

Compute a new distance matrix:

$D_{i,j}$	BP	RM	Human	
BP	0	0.1082	0.119	0.2272
RM	0.1082	0	0.154	0.2622
Human	0.119	0.154	0	0.273
	0.2272	0.2622	0.273	

Compute the matrix of $S_{i,j}$ values.

$S_{i,j}$	BP	RM	Human
Bovin		-0.3812	-0.3812
RM			-0.3812
Human			

Select the least value. There is a tie (Because there is just one unrooted tree on 3 vertices). Let's join RM with Human, adding a new internal vertex A.

$L_{A,RM} = (1/2)(0.154 + 0.2622 - 0.273) = 0.0716$

$L_{A,Human} = (1/2)(0.154 - 0.2622 + 0.273) = 0.0824$

and

$L_{BP,A} = (1/2)(0.1082 + 0.119 - 0.154) = 0.0366$

The Resulting Tree

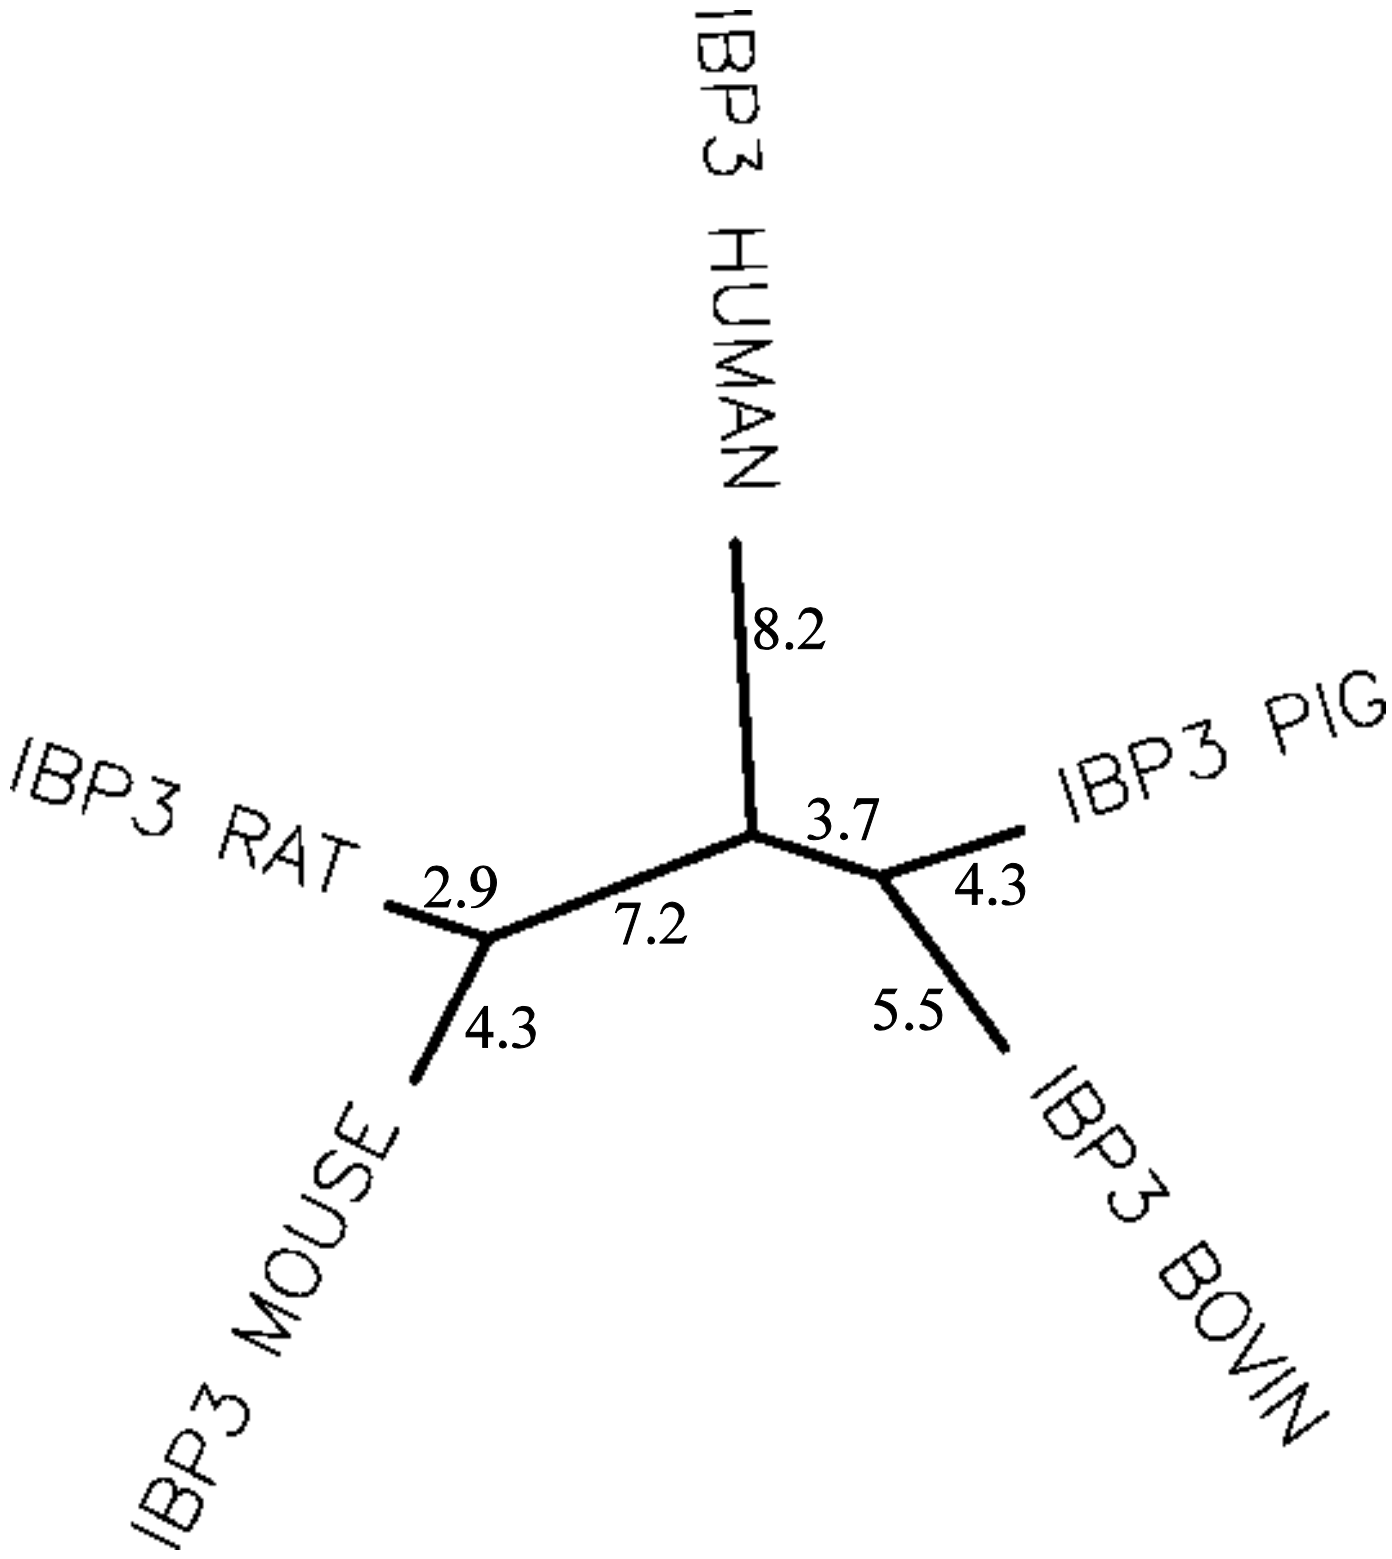

String Searching

Suppose we are given a sequence of nucleotides, such as ACGCGCAGGCA, and we wish to find all occurrences of that string in the human genome, about 3,000,000,000 bases long.

How long would that take?

- Without pre-processing, it would take at least 3,000,000,000 steps, because we need at least to look at every entry in the human genome.

- With pre-processing, perhaps we can set up some sort of indexing system to more rapidly search for these strings.

Indeed we can! The resulting search, after pre-processing, would take just 11 steps.

Pre-Processing???

The idea is that a popular string such as the human genome will be searched many, many times. Some pre-processing up front may be worth it, as every query will enjoy improved performance.

In fact, given:

- Search string S, of size |S|, and
- Query strings Q_1, Q_2, ..., of sizes $|Q_i|$,

We can create a data structure in time proportional to |S| so that each subsequent query takes time $|Q_i|$.

This data structure is called a *suffix tree*.

Sales Pitch

A search of MathSciNet for articles in mathematical journals revealed 9 matches since 2000, plus these in a biology database.

1: Lu B, Chen T.
A suffix tree approach to the interpretation of tandem mass spectra: applications to peptides of non-specific digestion and post-translational modifications.
Bioinformatics. 2003 Oct;19 Suppl 2:II113-II121.

2: Marsan L, Sagot MF.
Algorithms for extracting structured motifs using a suffix tree with an application to promoter and regulatory site consensus identification.
J Comput Biol. 2000;7(3-4):345-62.

3: Bejerano G, Yona G.
Variations on probabilistic suffix trees: statistical modeling and prediction of protein families.
Bioinformatics. 2001 Jan;17(1):23-43.

4: Dorohonceanu B, Nevill-Manning CG.
Accelerating protein classification using suffix trees.
Proc Int Conf Intell Syst Mol Biol. 2000;8:128-33.

A Suffix Tree

The following tree contains all suffixes of the string ATGCCAGCCACTTG

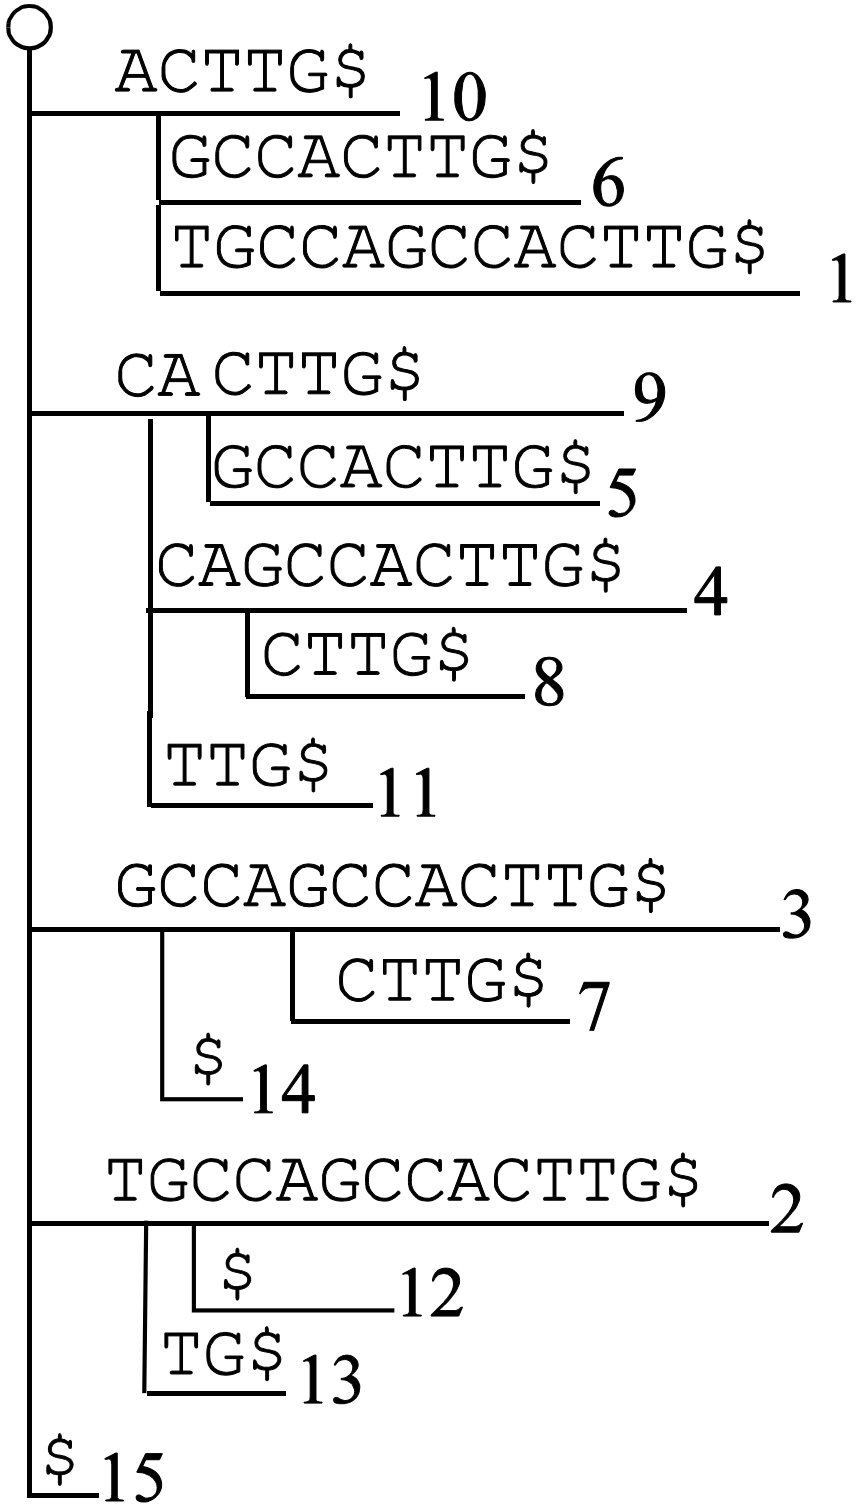

Some Uses for a Suffix Tree

1. Find all occurrences of the string “CCA”
   a. Start at the root
   b. Follow the string “CCA” as far as possible
      i. If we get stuck, the string does not occur
      ii. If we find the whole string, then the integers on all leaves below us tell where “CCA” may be found in the search string

2. Find all occurrences of “TGA”

3. How many times does “T” occur in the string?

4. Many more!

   a. See homework...

How does One Generate a Suffix Tree?

Naively:

1. Make a list of all suffixes in the search string, with "$" as an end marker

2. Write down the longest one (i.e., the whole string) on a single edge connected to the root

3. Go through the list from longest to shortest, adding the suffixes one at a time. As each string is added, walk down the tree as far as possible, searching for as long a prefix of that suffix as can be found, before adding a new branch to the tree.

E.g., if we add "ATGCTT" to our tree, then "AT":

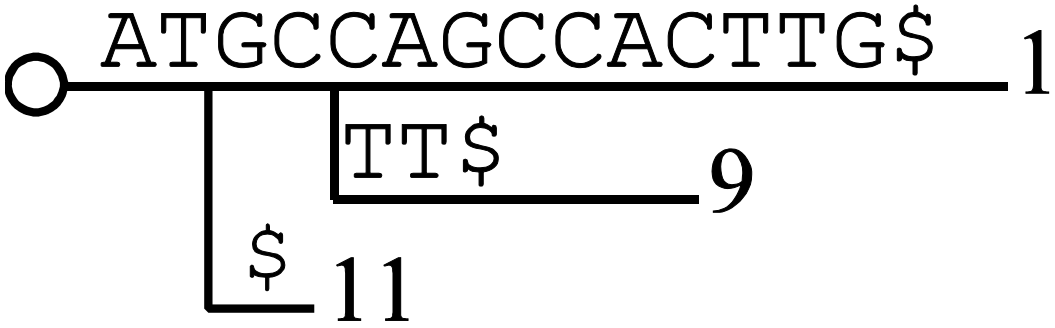

Suffix Trees are Fast

Our naive algorithm takes $O(n^2)$ steps in the worst case, which can be quite long if our search string is the human genome.

Martin Farach-Colton, right here at Rutgers / DIMACS, developed a linear-time algorithm for generating suffix trees.

Storage space for his tree is about 25 bytes per character in the search string.

WHAT!??

Edge label encoding is used...

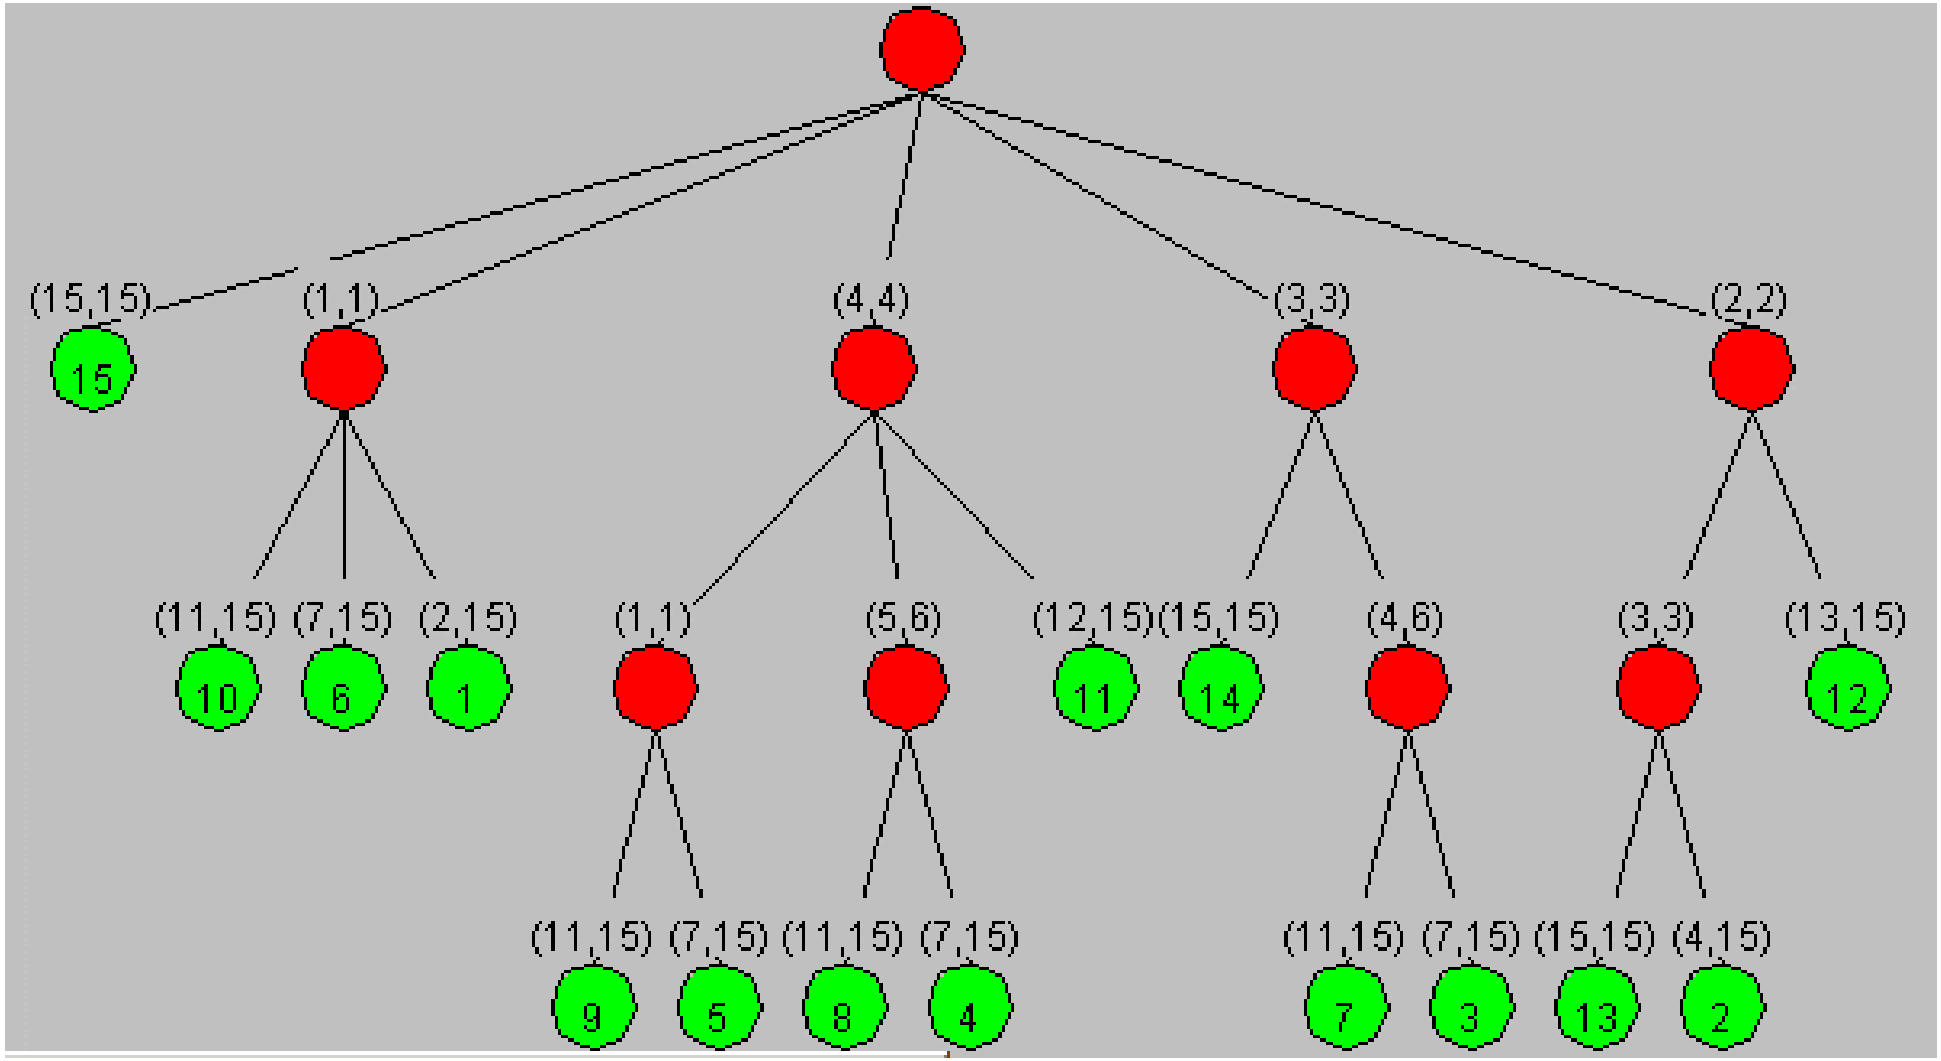

Your Turn

Let's generate a suffix tree for "CCGCATGCAAT"

Suffix Trees on Pairs of Strings

A small modification allows us to encode *two* search strings on a single suffix tree.

The trick is to use separate string termination symbols, such as “$” and “#.”

Here’s the tree for “ACCT”($) and “CCAT”(#)

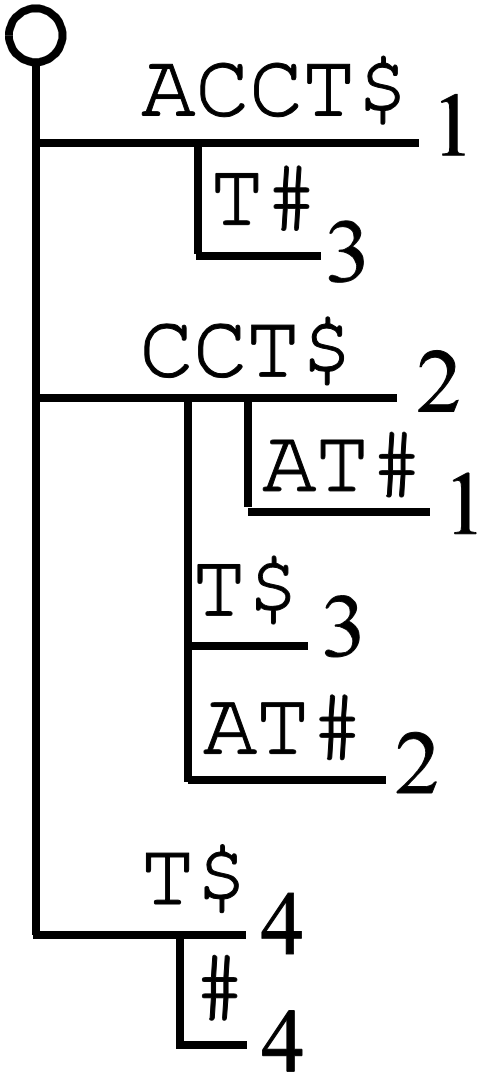

Handout #1 — Neighbor Joining Example

Here is our example worked out:

$D_{i,j}$	Bovin	Pig	Rat	Mouse	Human	
Bovin	0	0.098	0.197	0.204	0.173	0.672
Pig	0.098	0	0.181	0.189	0.163	0.631
Rat	0.197	0.181	0	0.072	0.176	0.626
Mouse	0.204	0.189	0.072	0	0.204	0.669
Human	0.173	0.163	0.176	0.204	0	0.716
	0.672	0.631	0.626	0.669	0.716	

Compute the matrix of $S_{i,j}$ values.

$S_{i,j}$	Bovin	Pig	Rat	Mouse	Human
Bovin		-1.009	-0.707	-0.729	-0.869
Pig			-0.714	-0.733	-0.858
Rat				-1.079	-0.814
Mouse					-0.773
Human					

Select the least value. This is between rat and mouse

Branch Lengths:

$L_{R,RM} = (1/6)(0.072*3 + 0.626 - 0.669) = 0.0288$

$L_{M,RM} = (1/6)(0.072*3 - 0.626 + 0.669) = 0.0432$

Compute a new distance matrix:

$D_{i,j}$	Bovin	Pig	RM	Human	
Bovin	0	0.098	0.1645	0.173	0.4355
Pig	0.098	0	0.149	0.163	0.41
RM	0.1645	0.149	0	0.154	0.4675
Human	0.173	0.163	0.154	0	0.49
	0.4355	0.41	0.4675	0.49	

Compute the matrix of $S_{i,j}$ values.

$S_{i,j}$	Bovin	Pig	RM	Human
Bovin		-0.6495	-0.574	-0.5795
Pig			-0.5795	-0.574
RM				-0.6495
Human				

Select the least value. There is a tie (can you see why there has to be?). We'll join Bovin with Pig.

$L_{B,BP} = (1/4)(0.098*2 + 0.4355 - 0.41) = 0.0554$

$L_{P,BP} = (1/4)(0.098*2 - 0.4355 + 0.41) = 0.0426$

Compute a new distance matrix:

$D_{i,j}$	BP	RM	Human	
BP	0	0.1082	0.119	0.2272
RM	0.1082	0	0.154	0.2622
Human	0.119	0.154	0	0.273
	0.2272	0.2622	0.273	

Compute the matrix of $S_{i,j}$ values.

$S_{i,j}$	BP	RM	Human
Bovin		-0.3812	-0.3812
RM			-0.3812
Human			

Select the least value. There is a tie (Because there is just one unrooted tree on 3 vertices). Let's join RM with Human, adding a new internal vertex A.

$L_{A,RM} = (1/2)(0.154 + 0.2622 - 0.273) = 0.0716$

$L_{A,Human} = (1/2)(0.154 - 0.2622 + 0.273) = 0.0824$

and

$L_{BP,A} = (1/2)(0.1082 + 0.119 - 0.154) = 0.0366$

And we are done.

The Resulting Tree

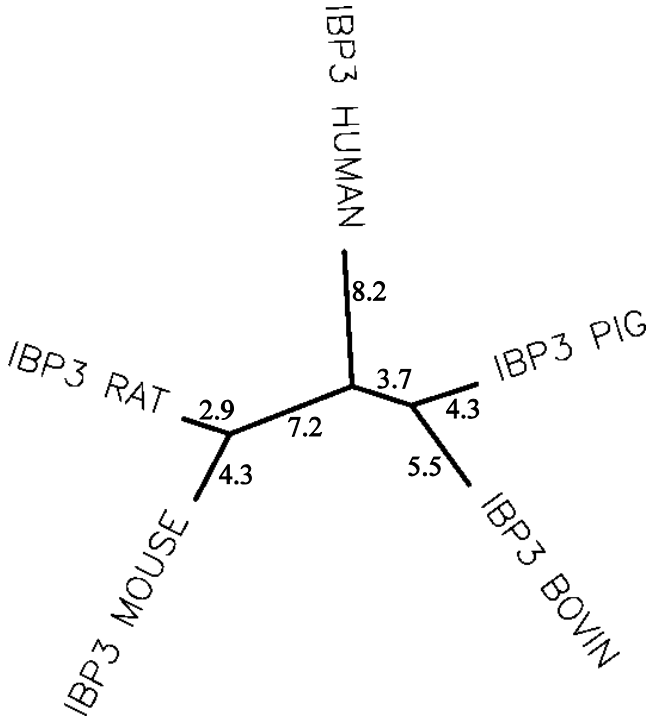

Handout #2 — Generating a Suffix Tree

Steps to Generate a Suffix Tree:

1. Make a list of all suffixes in the search string, with "$" as an end marker
2. Write down the longest one (i.e., the whole string) on a single edge connected to the root

3. Go through the list from longest to shortest, adding the suffixes one at a time. As each string is added, walk down the tree as far as possible, searching for as long a prefix of that suffix as can be found, before adding a new branch to the tree. Mark each new leaf with the index of where the prefix started.

E.g., if we add "ATGCTT" to our tree, then "AT":

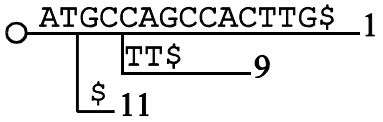

Why not try building a suffix tree for the string "CCGCATGCAAT"

Exercises — Suffix Trees

Presentation Problems:

1. Prove that every unrooted binary tree on at least 2 vertices has a pair of leaves adjacent to a common internal node. We called this a "cherry" in class.

2. Build the suffix tree for the search string "CACCTACCTCACC."

3. Show how to read, from the suffix tree in problem 2, where all occurrences of the string "ACC" can be found.

4. Here is a suffix tree for the string for one of the strings "CAAATGTGCATGC," "CCCGTCGTACCGT" or "TTCACAAACTGGT," represented with edge label compression. Which string?

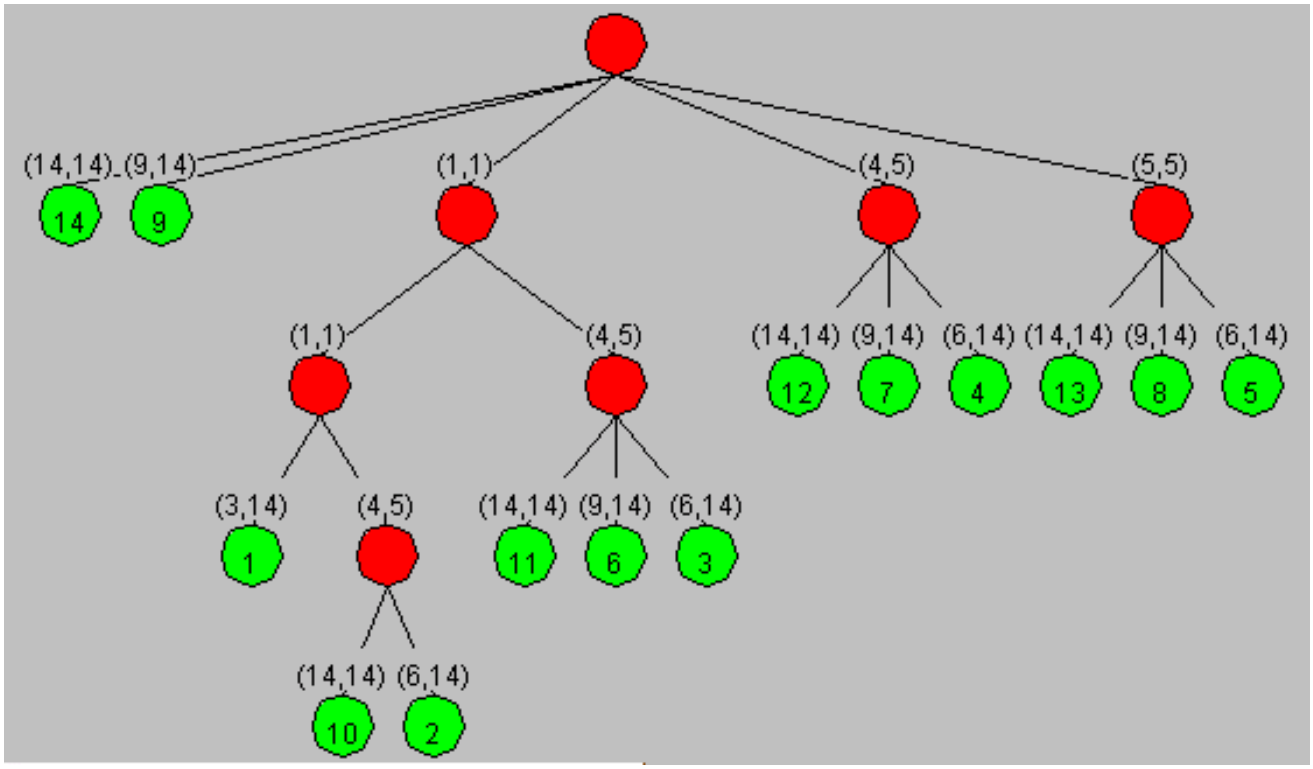

5. How would you use the edge-label compressed suffix tree given above to find all occurrences of "CGT" in the search string?

6. Here is a suffix tree for some search string with details graciously omitted. Suppose that each edge has a string of length exactly 2, except for the edges at the leaves, which are of varying lengths. What is the length of the longest string that occurs at least 8 times in the search string? How many such strings are there? Do you see a neat dynamic programming way to do this?

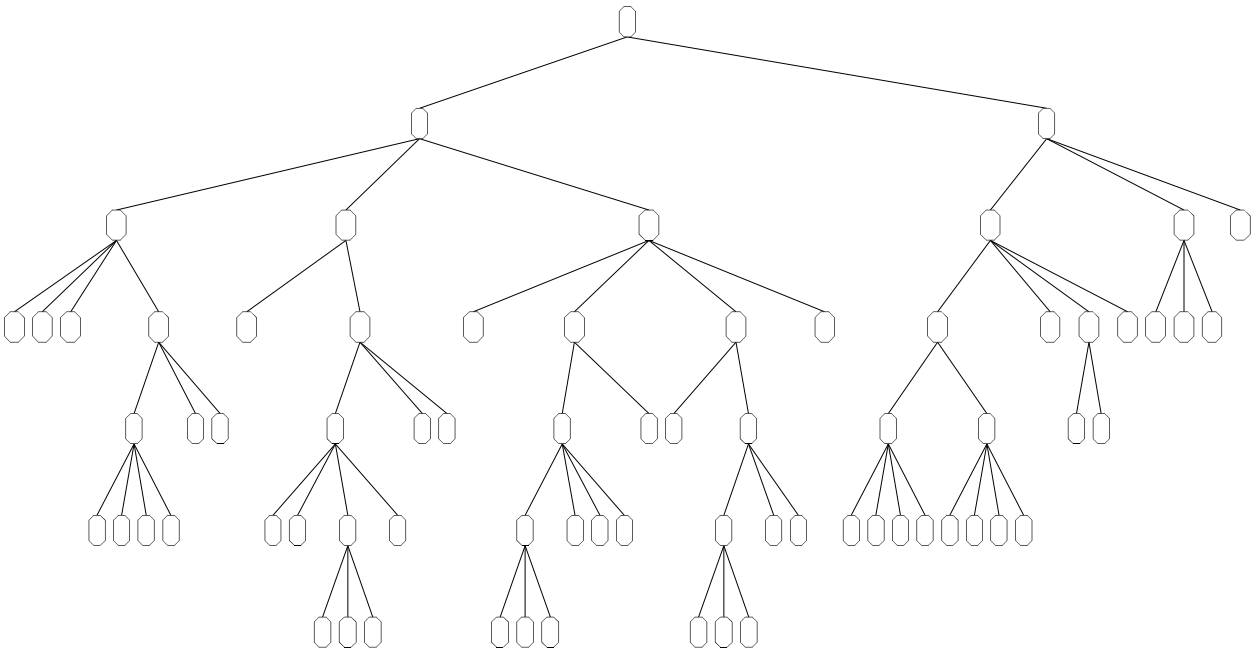

7. Build the suffix tree that encodes the *two* strings "CAGAGT" and "GAGGA." Show how to use your string to find the longest common substring to those two strings.