SongSeq User Manual

SA+ Spreadsheets

The input to SongSeq is the Microsoft Excel spreadsheets generated by Sound Analysis Pro (SA+) software, which can be downloaded from (http://ofer.sci.ccny.cuny.edu/sound_analysis_pro).

It's very crucial for SongSeq NOT to alter the format of the spreadsheets that are generated by SA+. **Fig. 1** below shows a screenshot for a typical SA+ spreadsheet. Notice that the first column starts at "B" and the first row of data starts at row 5. This is how SA+ generates its data and it is very important to keep all the spreadsheets that are to be simulated by SongSeq in this same format.

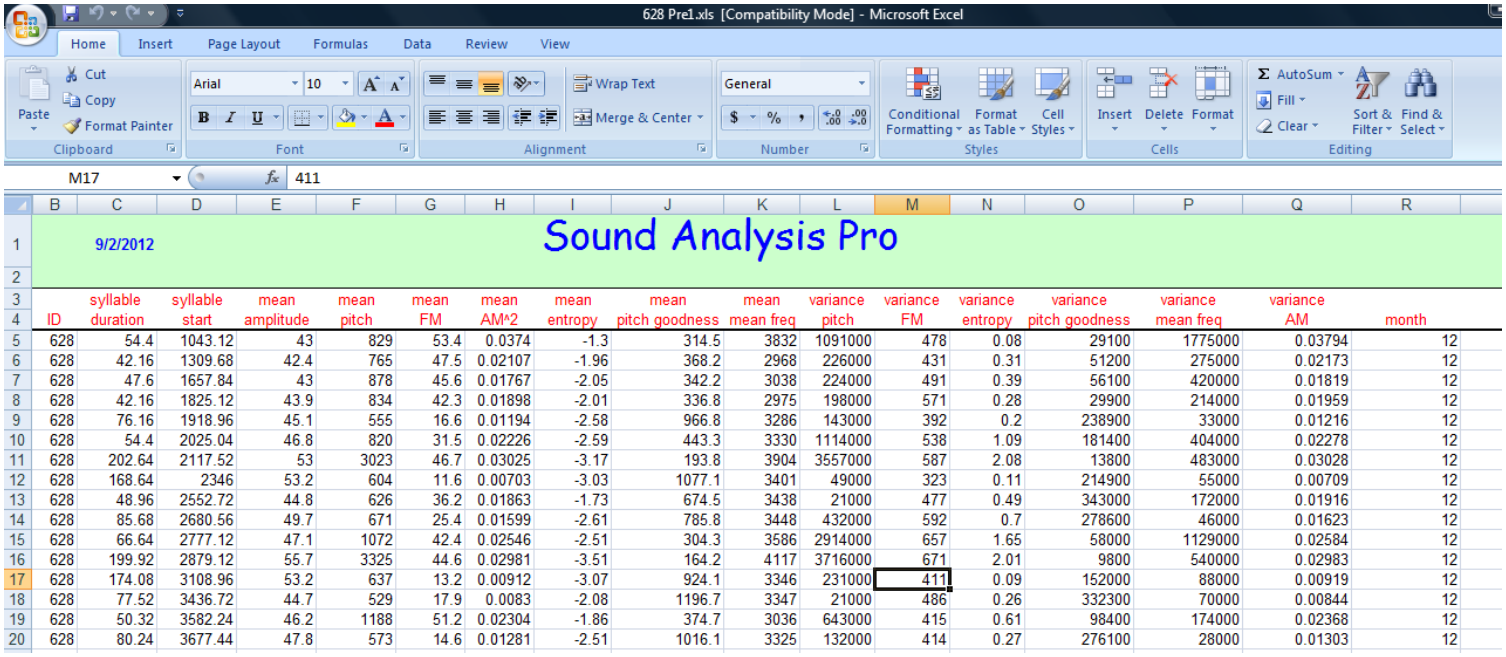

	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R
1		9/2/2012			Sound Analysis Pro												
2																	
3		syllable	syllable	mean	mean	mean	mean	mean	mean	mean	variance	variance	variance	variance	variance	variance	
4	ID	duration	start	amplitude	pitch	FM	AM^2	entropy	pitch goodness	mean freq	pitch	FM	entropy	pitch goodness	mean freq	AM	month
5	628	54.4	1043.12	43	829	53.4	0.0374	-1.3	314.5	3832	1091000	478	0.08	29100	1775000	0.03794	12
6	628	42.16	1309.68	42.4	765	47.5	0.02107	-1.96	368.2	2968	226000	431	0.31	51200	275000	0.02173	12
7	628	47.6	1657.84	43	878	45.6	0.01767	-2.05	342.2	3038	224000	491	0.39	56100	420000	0.01819	12
8	628	42.16	1825.12	43.9	834	42.3	0.01898	-2.01	336.8	2975	198000	571	0.28	29900	214000	0.01959	12
9	628	76.16	1918.96	45.1	555	16.6	0.01194	-2.58	966.8	3286	143000	392	0.2	238900	33000	0.01216	12
10	628	54.4	2025.04	46.8	820	31.5	0.02226	-2.59	443.3	3330	1114000	538	1.09	181400	404000	0.02278	12
11	628	202.64	2117.52	53	3023	46.7	0.03025	-3.17	193.8	3904	3557000	587	2.08	13800	483000	0.03028	12
12	628	168.64	2346	53.2	604	11.6	0.00703	-3.03	1077.1	3401	49000	323	0.11	214900	55000	0.00709	12
13	628	48.96	2552.72	44.8	626	36.2	0.01863	-1.73	674.5	3438	21000	477	0.49	343000	172000	0.01916	12
14	628	85.68	2680.56	49.7	671	25.4	0.01599	-2.61	785.8	3448	432000	592	0.7	278600	46000	0.01623	12
15	628	66.64	2777.12	47.1	1072	42.4	0.02546	-2.51	304.3	3586	2914000	657	1.65	58000	1129000	0.02584	12
16	628	199.92	2879.12	55.7	3325	44.6	0.02981	-3.51	164.2	4117	3716000	671	2.01	9800	540000	0.02983	12
17	628	174.08	3108.96	53.2	637	13.2	0.00912	-3.07	924.1	3346	231000	411	0.09	152000	88000	0.00919	12
18	628	77.52	3436.72	44.7	529	17.9	0.0083	-2.08	1196.7	3347	21000	486	0.26	332300	70000	0.00844	12
19	628	50.32	3582.24	46.2	1188	51.2	0.02304	-1.86	374.7	3036	643000	415	0.61	98400	174000	0.02368	12
20	628	80.24	3677.44	47.8	573	14.6	0.01281	-2.51	1016.1	3325	132000	414	0.27	276100	28000	0.01303	12

Fig. 1

It's preferred to have the Excel spreadsheets saved in ".xls" format rather than ".xlsx". This could be done in MS. Excel by saving the ".xlsx" spreadsheet as "Excel 97-2003 Workbook".

Make Template and Sequence

a- Template Selection

The sequencing process starts by choosing the "Make Template and Sequence" menu item (**Fig. 2**).

Fig. 2

The next step is designating a single spreadsheet as the template. This spreadsheet is a typical SA+ spreadsheet containing the acoustic variables. This is done by browsing for the spreadsheet on your computer. After that, select two acoustic features from the template as shown in **Fig. 3**. The acoustic features drop down menu lists all the features that SA+ generates.

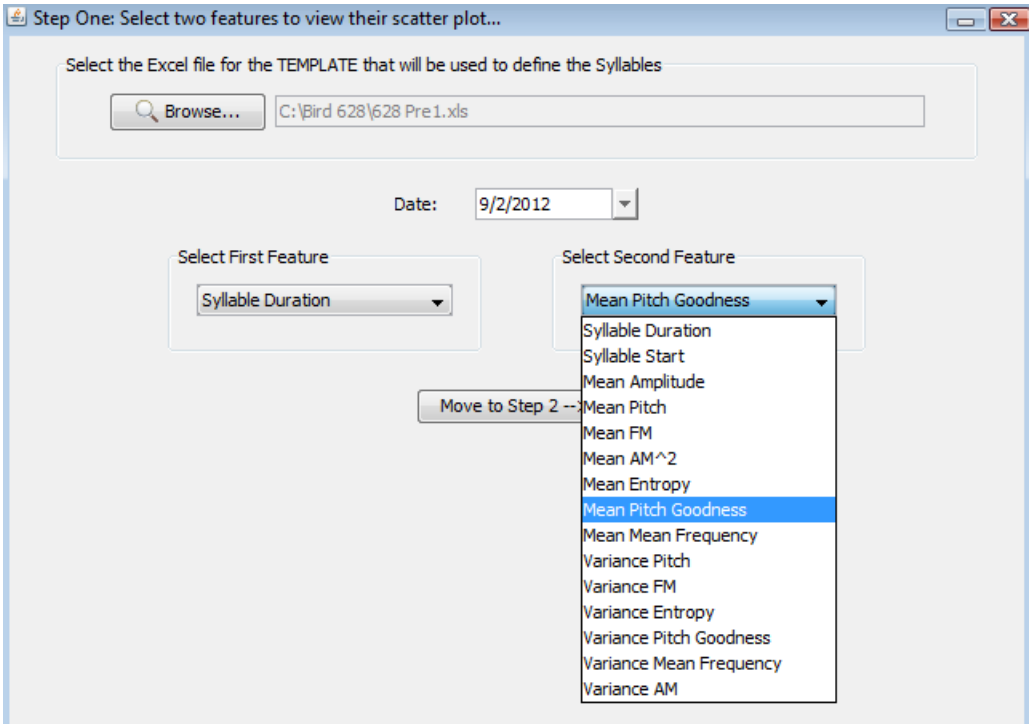

Fig. 3

b- Syllables Identification: Using One Pair of Features

Once the template and the two features are selected, you can move to the next step by clicking on "Move to Step 2 →" button (**Fig. 3**). This will pull up a window showing the 2D scatter plot of the template with the two features data points extracted from the spreadsheet as illustrated below (**Fig. 4**). **Fig. 4** also shows two tabbed panes on the right side that will be discussed below.

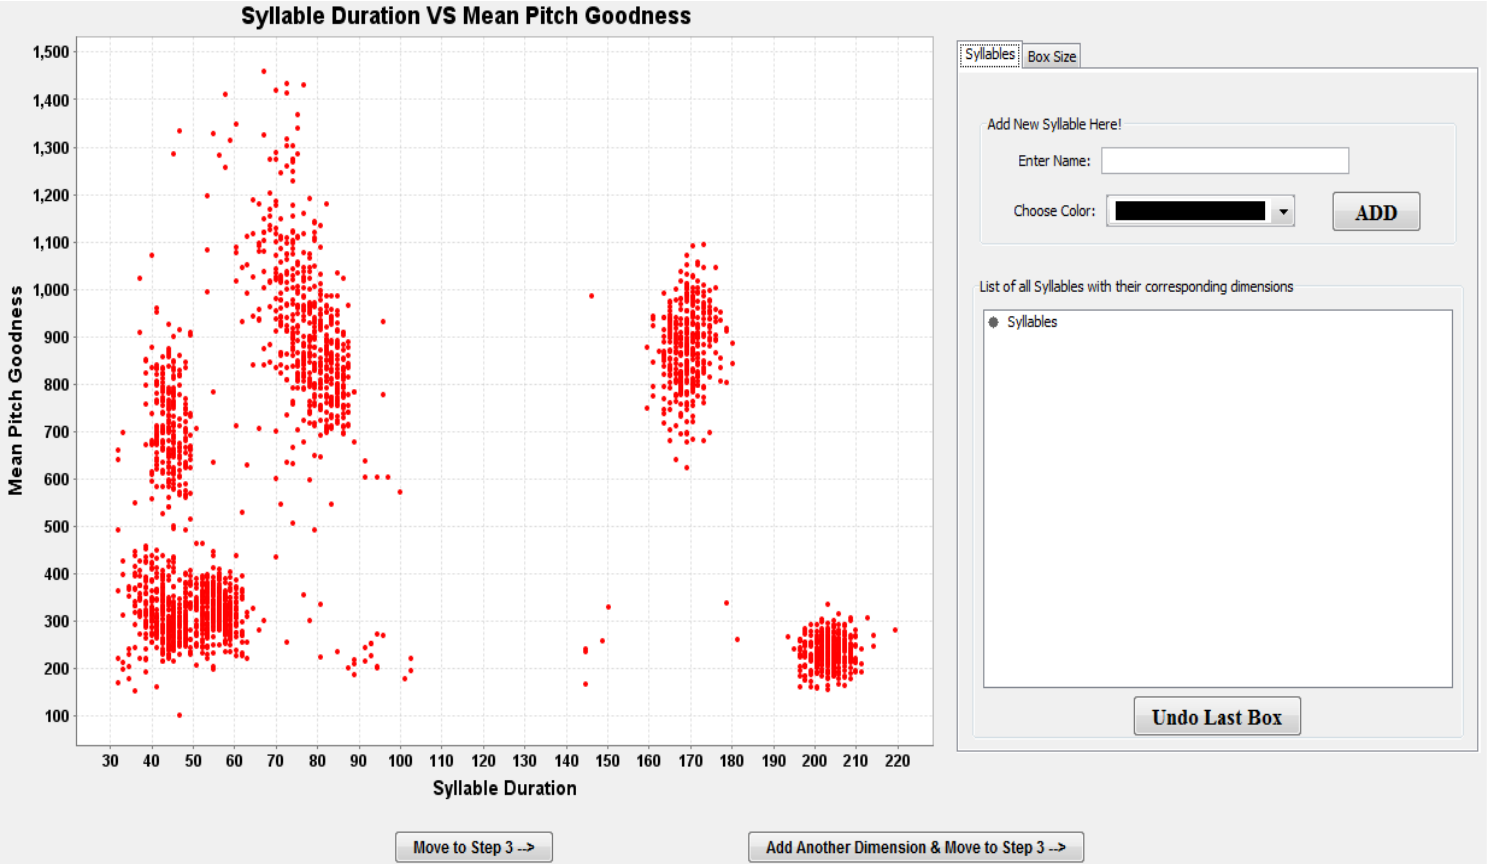

Fig. 4

The next step is using the graphical tools to associate the syllable clusters in **Fig. 4** with syllable labels. To add the first syllable (say "A"), enter its name in the text box on the right hand side, choose a color for it from the drop down menu of colors, and then click "ADD". The syllable will be added as a syllable node to the tree on the lower right corner of **Fig. 4**. SongSeq prevents adding the same syllable name twice and it prevents choosing the same color for more than one syllable.

After the syllable is added, paint using mouse clicks onto the 2D scatter plot over the area that represents the syllable. Each mouse click paints a box onto the scatter plot and adds a "child node" under the syllable node in the tree, and this child node lists the dimensions of the two features covered under the 2D painted box. You can control the size of the painted boxes under the "Box Size" tabbed pane. This is illustrated in **Fig. 5**. You can choose a variety of box sizes while painting the same syllable. Also, if you mistakenly painted a box over the wrong region, you can undo the last box painted by clicking the "Undo Last Box" button. Continuous clicking of this button keeps removing the last painted box for the last syllable undergoing painting.

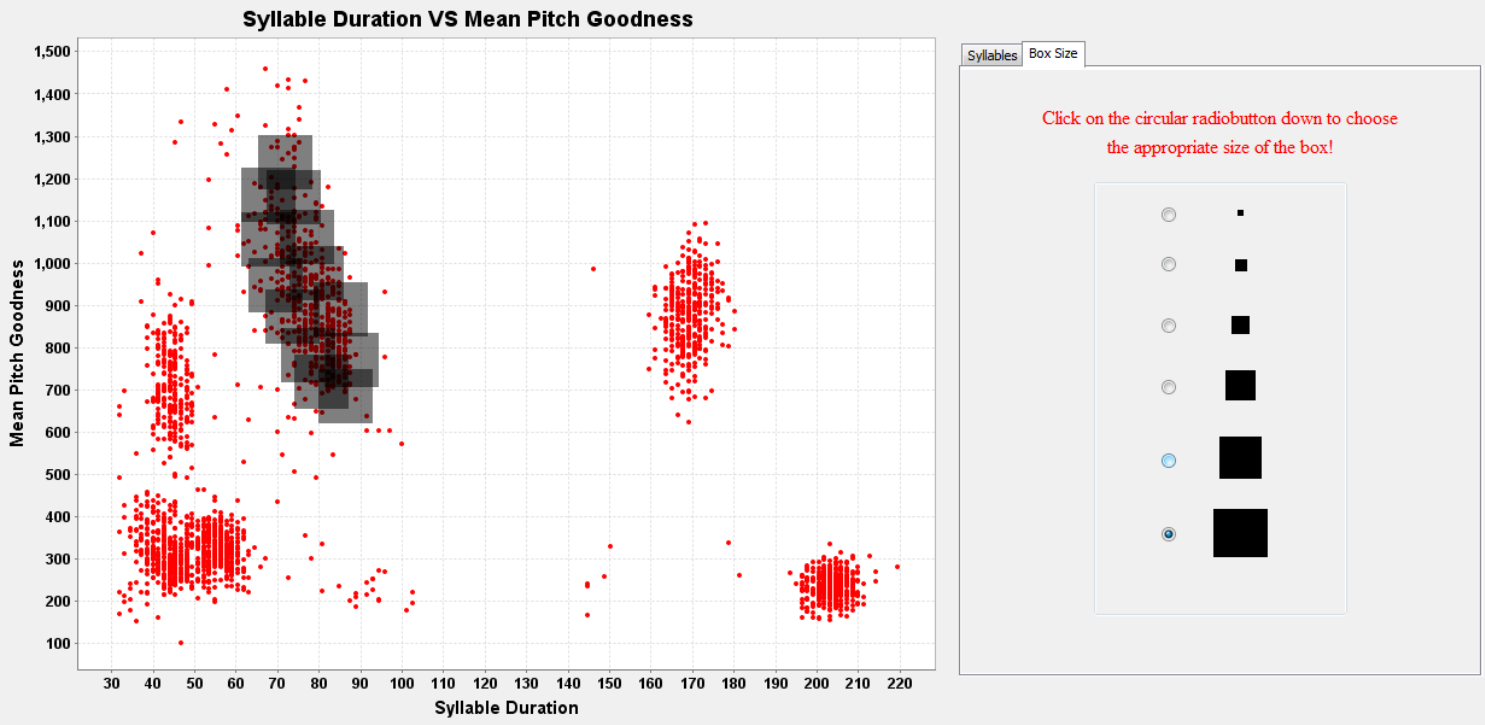

Fig. 5

After you define the dimensions of the first syllable, repeat the same for the remaining syllables. **Fig. 6** shows the 2D scatter plot painted for all syllables. Notice, the tree on the right side shows a list of all syllables, along with their associated colors and the number of boxes painted to define their boundaries. In order to see the dimensions of each of those boxes, press on the (+) sign near the syllable name on the tree.

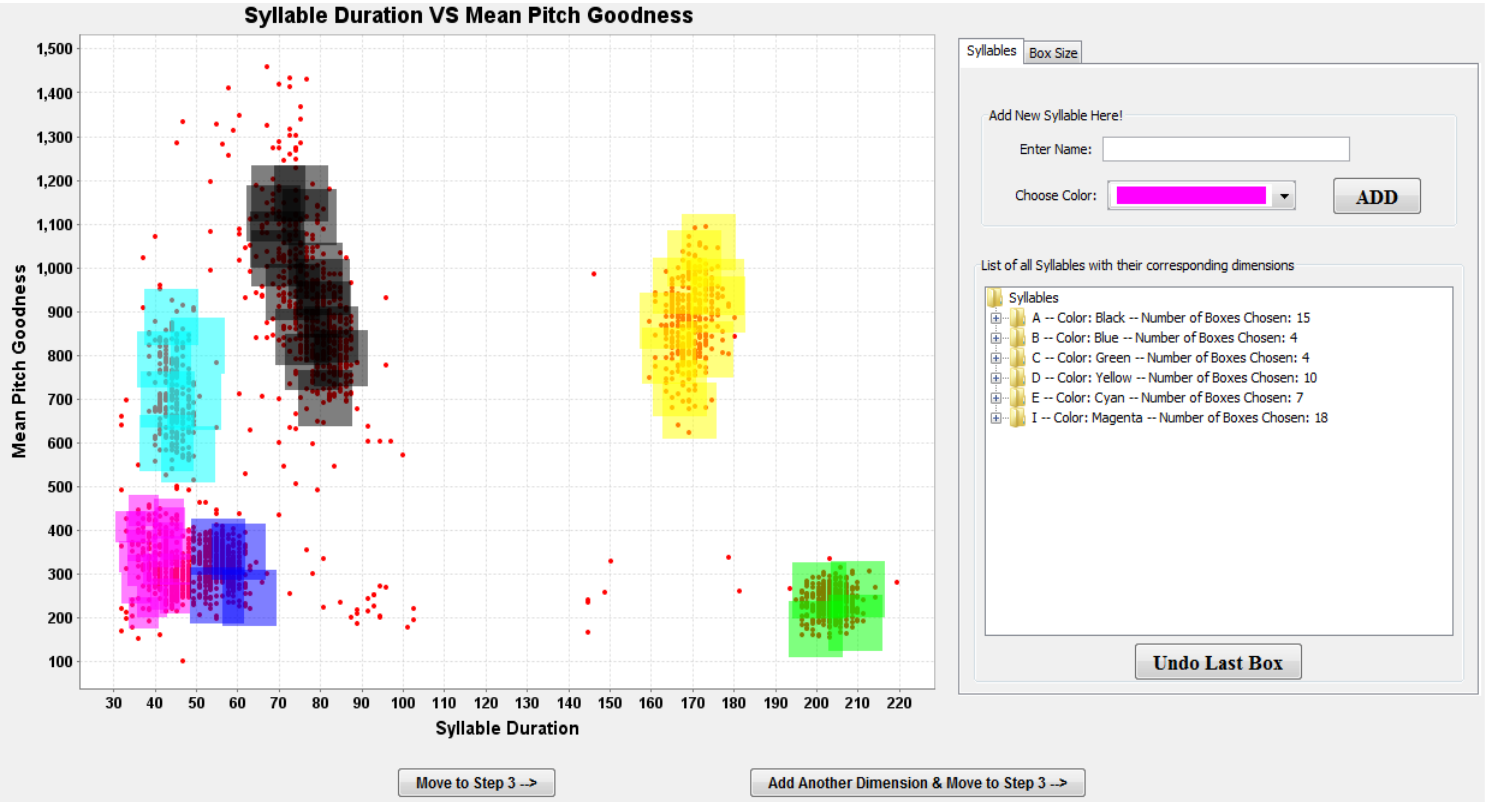

Fig. 6

The painted boxes are transparent. However, if you hover the mouse outside the scatter plot region and then hover it back inside the scatter plot region, the transparency of the boxes decreases. To return to the initial transparency level, minimize the SongSeq window, after that maximize it and then hover the mouse inside the scatter plot region.

If all syllables are easily identified with one pair of features, move next by clicking on "Move to Step 3 →". However, it may happen that some syllables are not easily identified with one pair of features (discussed later), and they can be better identified with a second pair of features. If this case happens, move next by clicking on "Add Another Dimension & Move to Step 3 →". Both cases will be illustrated below.

Click on the "Move to Step 3 →" button. This shows a panel similar to **Fig. 7** below.

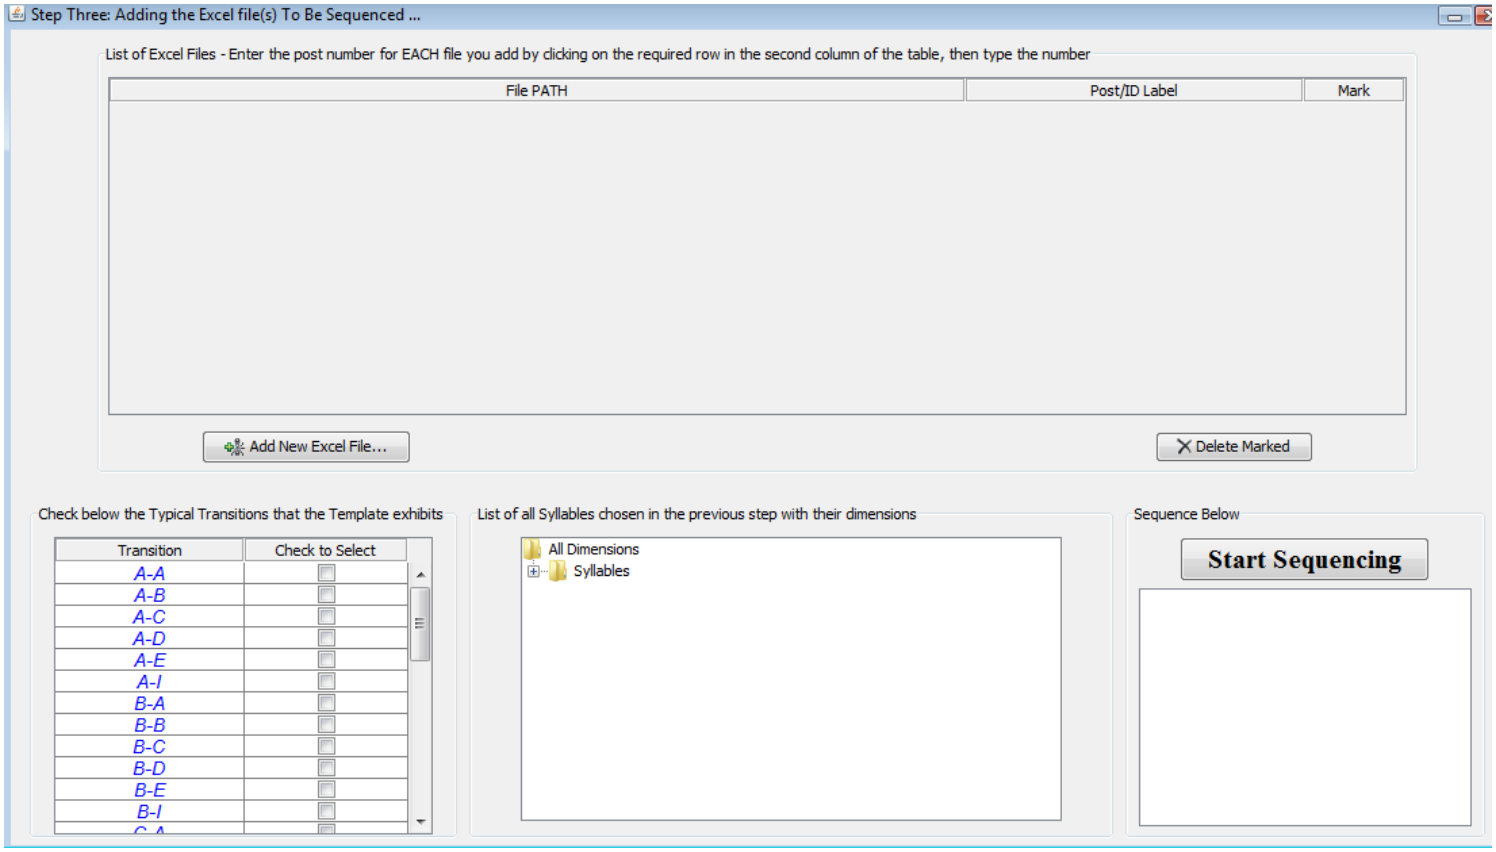

Fig. 7

c- Loading Target Files

The next step is choosing the target files (typical SA+ spreadsheets). This is done by clicking on the "Add New Excel File ..." button. The file chooser box that pops up allows you to select multiple excel files at once as illustrated in **Fig. 8**. We prefer to add the excel files in the same order as the chronological order of events that they represent. For instance, if file "Post1.xls" represents the first day after surgery and file "Post2.xls" represents the second day after surgery, it is better to upload "Post1.xls" before "Post2.xls". This is better and more convenient for reporting analysis at the end.

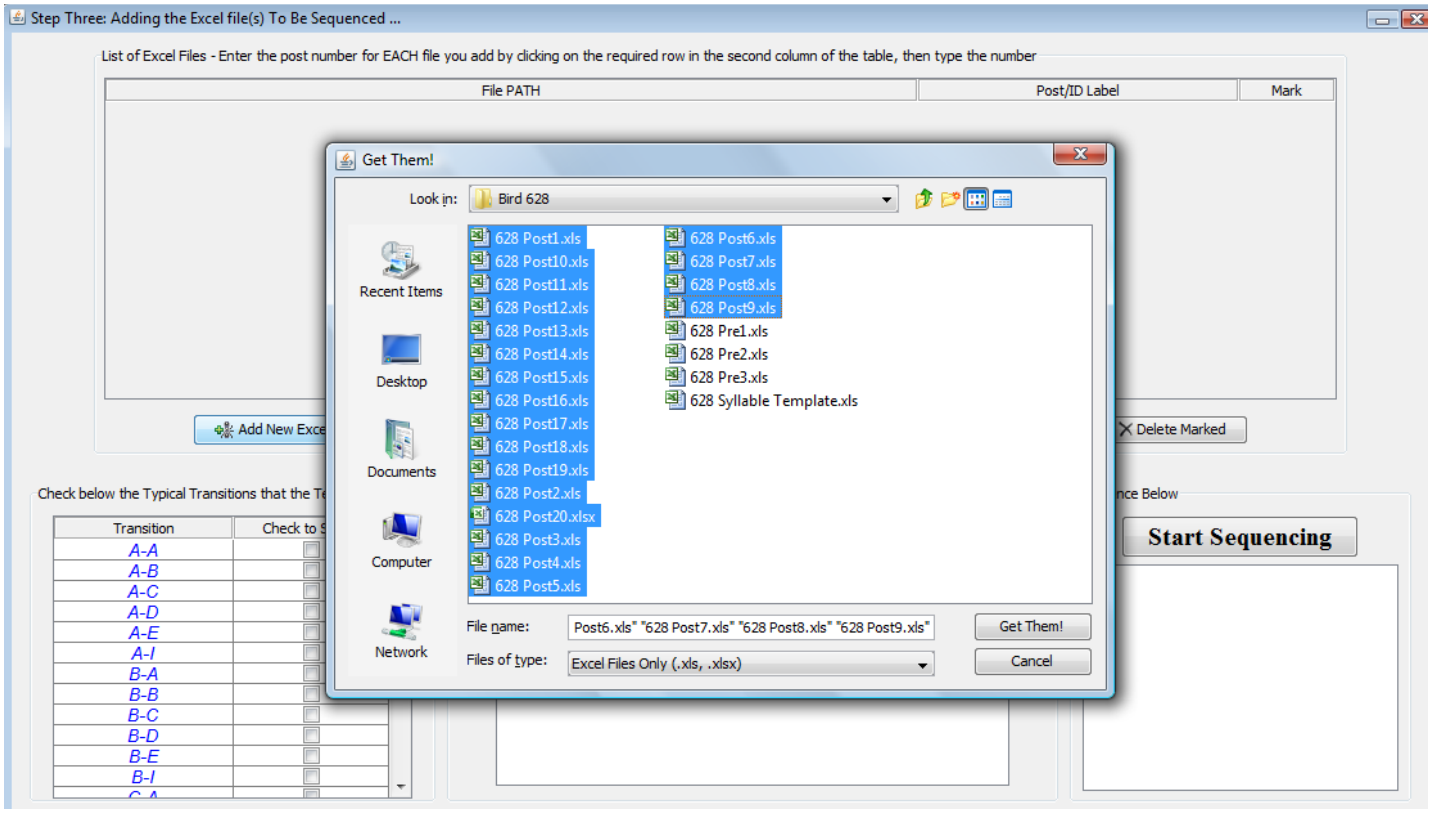

Fig. 8

Fig. 9 shows the upload of one pre-surgery and 13 post-surgery spreadsheets. If one or more target files are uploaded mistakenly, they can be removed by marking their corresponding rows on the third column of the table, then clicking "Delete Marked".

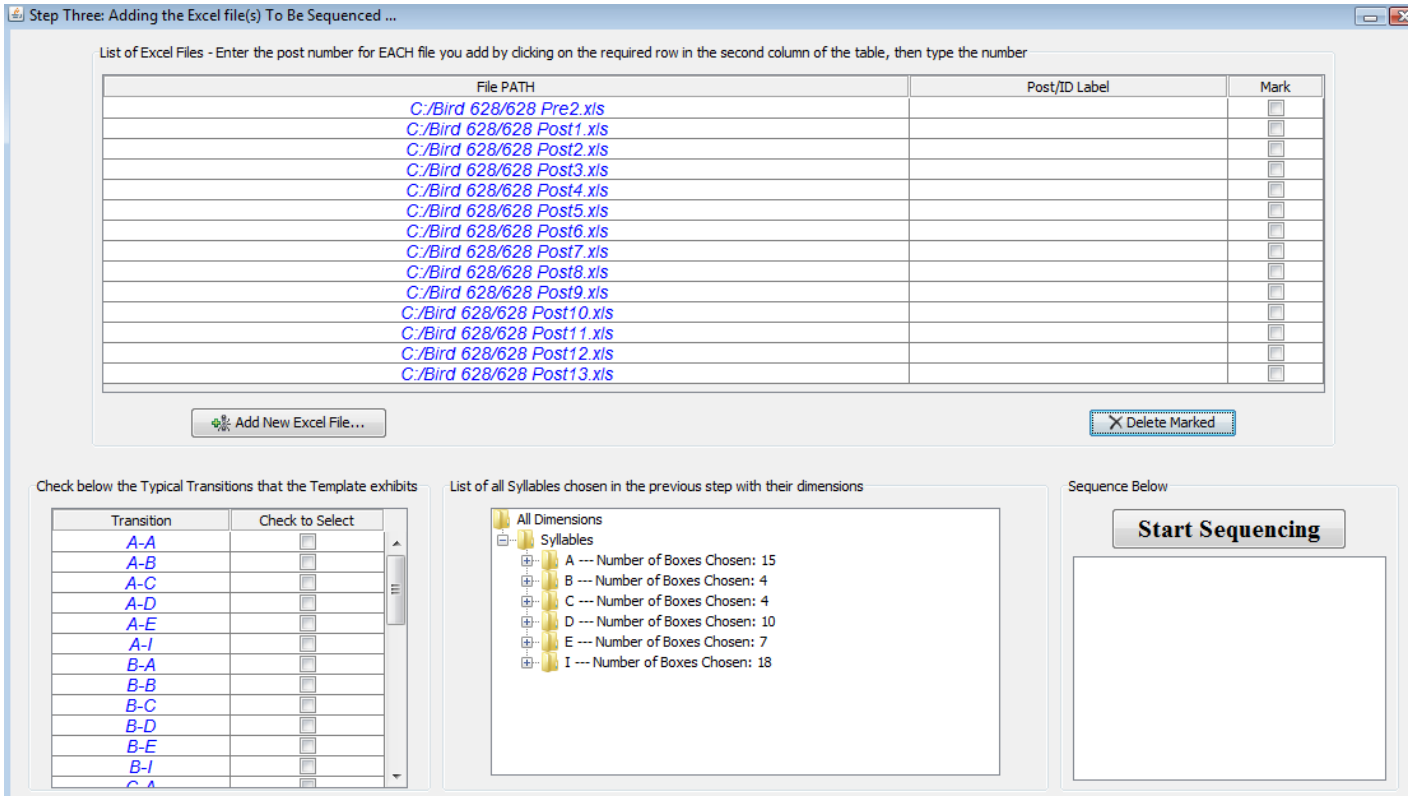

Fig. 9

The next step is uniquely identifying each target file uploaded by a unique name or ID. This is done by clicking on the second column of the table and typing in the corresponding row its unique name or ID (**Fig. 10**). The name or ID is totally up to the user, but no two target files can share the same name, and each target file must be designated by a name or ID. To fasten the labeling process, pressing the "Enter" on the keyboard moves the cursor from one row to another (this is instead of clicking the mouse on each corresponding row).

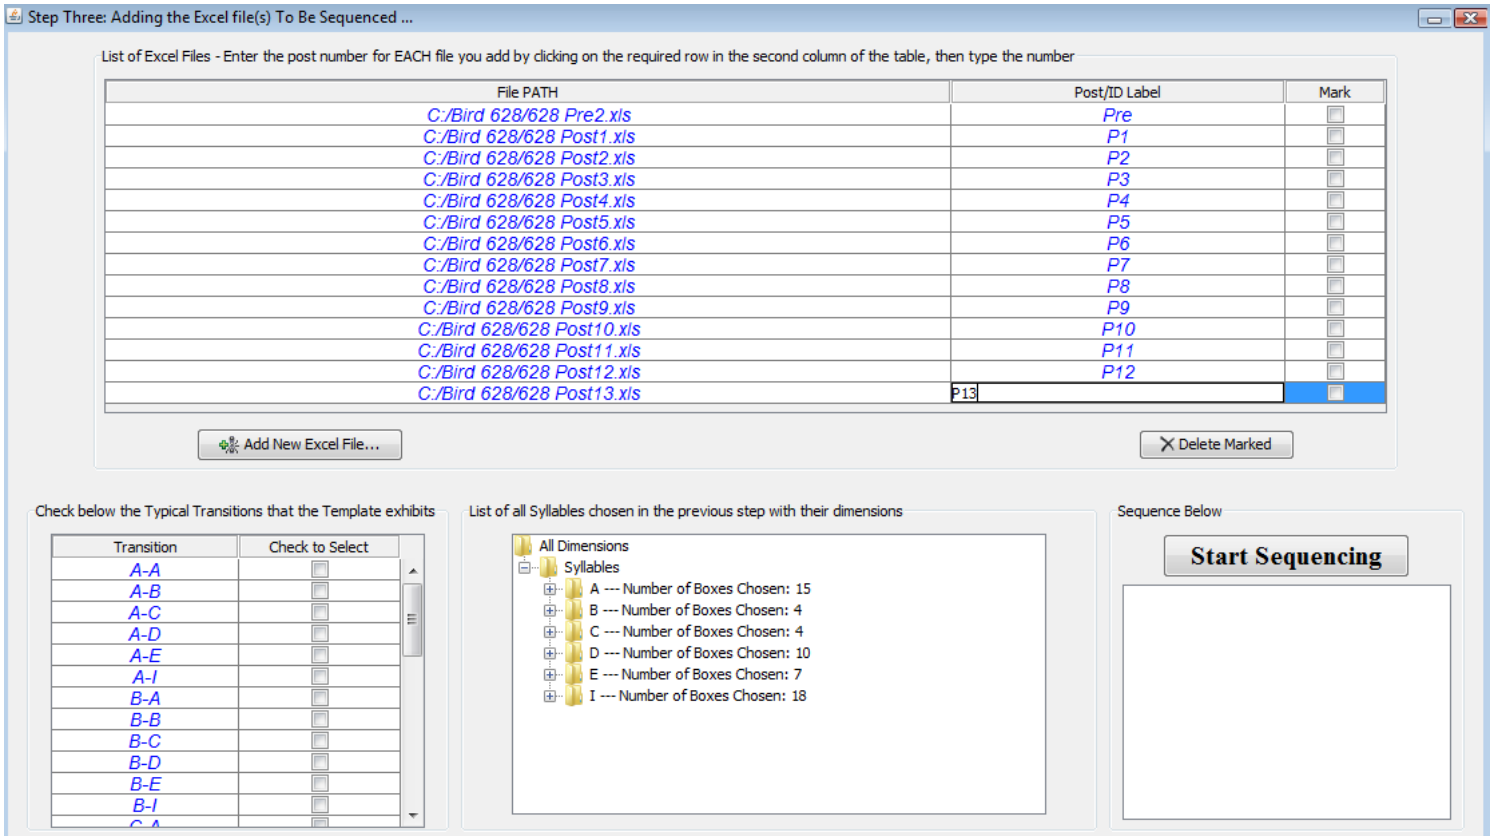

Fig. 10

Next, the typical transitions are chosen by marking them from the checkboxes in the lower left panel of **Fig. 10**. SongSeq lists all possible combinations of syllable transitions based on the syllable names entered in the first step. For this particular example, the typical transitions chosen are "I-I", "I-A", "A-I", "A-B", "B-C", "C-D", "D-E" and "E-A" (**Fig. 11**). When the typical transitions are chosen, sequencing the target files can start by clicking "Start Sequencing" button. The text area under the "Start Sequencing" button shows the progress of the sequencing process by writing the name of the target file currently being simulated. This is illustrated in **Fig. 11** (the screenshot was taken after SongSeq finished running files Pre, P1, P2 and P3, while P4 was being simulated).

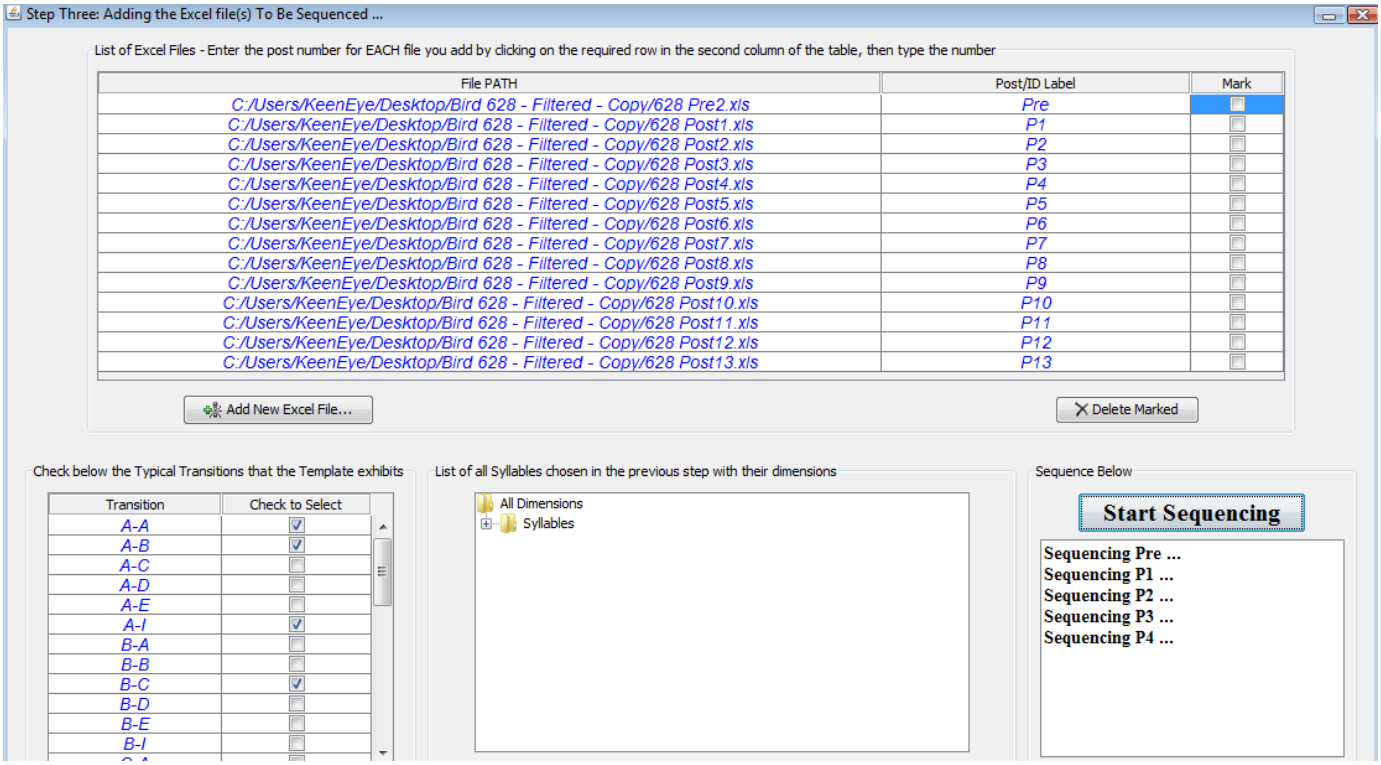

Fig. 11

d- Viewing Results

Once SongSeq finishes sequencing all the target files, a window that aims at summarizing the results appears (**Fig. 12**). Initially, the window is empty showing no results. However, a drop down menu at the top shows the names for all the target files sequenced. The results for each of those target files can be accessed by choosing the target file name from the drop down menu.

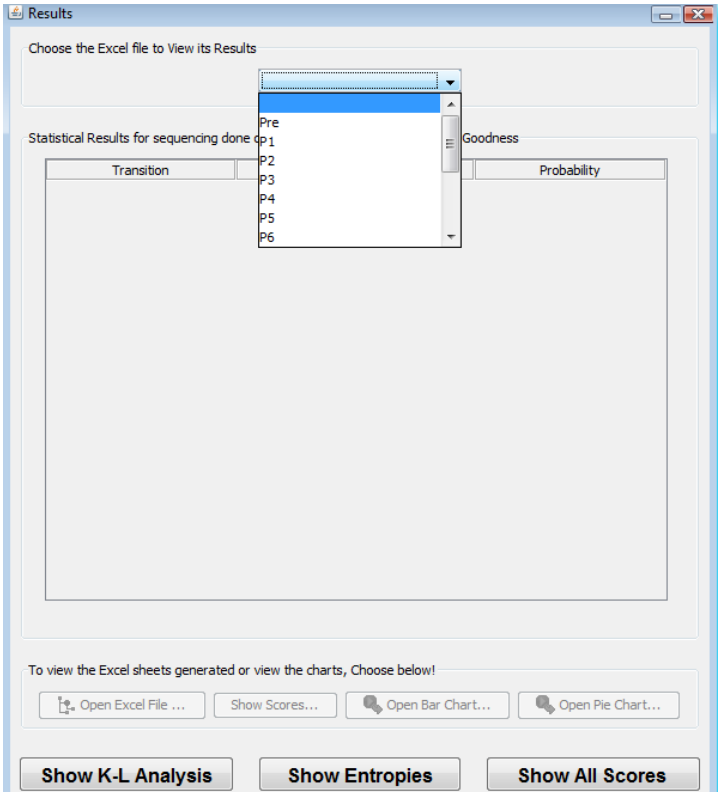

Fig. 12

Fig. 13 shows the results for the selected Pre day. The table lists all the encountered transitions, the number of occurrences of each transition and the corresponding probability of each transition.

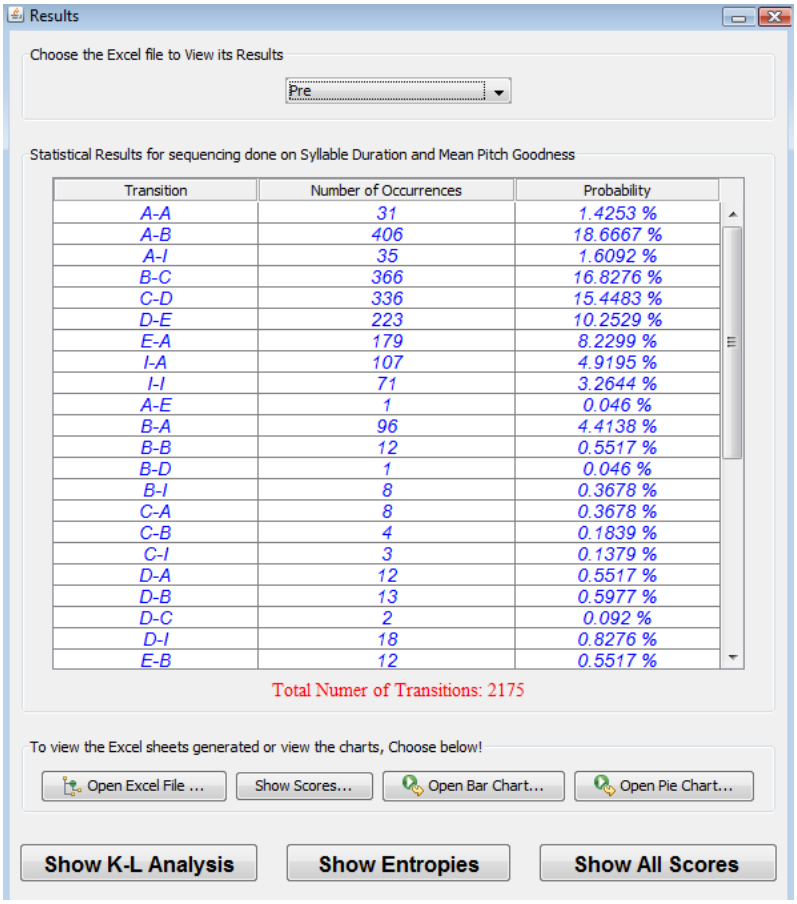

Fig. 13

A bar chart of all the transitions can be viewed by clicking "Open Bar Chart" (**Fig. 14**). The blue bars represent the typical transitions, that is, the transitions chosen by the user in the step of **Fig. 11**. The yellow bars represent "atypical transitions", that is, transitions involving syllables the bird sung but not specified as typical by the user. These transitions have the form "X-X" where X can be any syllable name, but where "X-X" is not a user-specified typical transition. Finally, the red bars represent "NMS" transitions. These transitions involve the data points that are not painted by the user (the red dots on the scatter plot). NMS transitions have the forms "X-N", "N-X" or "N-N" where X can be any syllable name and N represents NEW or NMS syllable.

Notice that the transitions listed in the table of **Fig. 13** occur in the same order as the transitions in the bar chart of **Fig. 14**. Results can also be shown in a pie chart by clicking "Open Pie Chart" (**Fig. 15**).

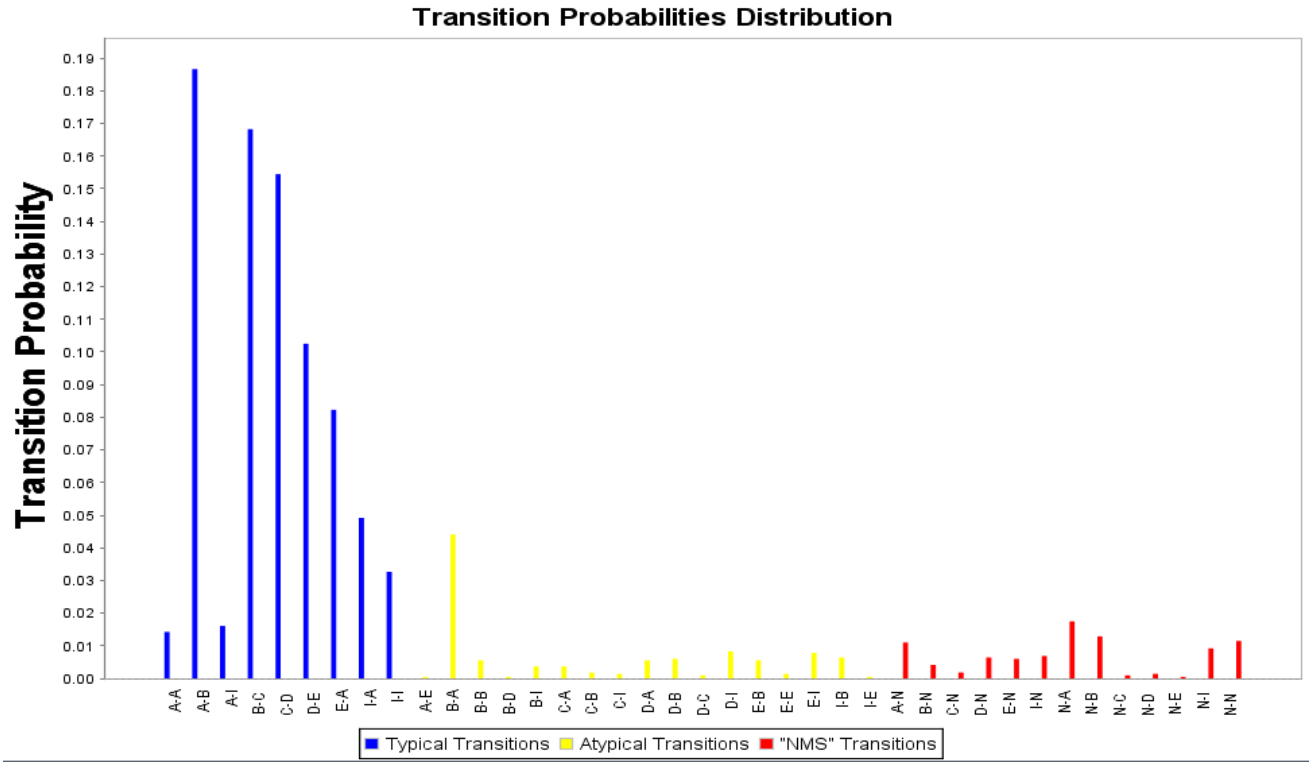
Transition Probabilities Distribution
Transition Probability
0.19
0.18
0.17
0.16
0.15
0.14
0.13
0.12
0.11
0.10
0.09
0.08
0.07
0.06
0.05
0.04
0.03
0.02
0.01
0.00
A-A
A-B
A-I
B-C
C-D
D-E
E-A
I-A
I-I
A-E
B-A
B-B
B-D
B-I
C-A
C-B
C-I
D-A
D-B
D-C
D-I
E-B
E-E
E-I
I-B
I-E
A-N
B-N
C-N
D-N
E-N
I-N
N-A
N-B
N-C
N-D
N-E
N-I
N-N
Typical Transitions
Atypical Transitions
"NMS" Transitions

Fig. 14

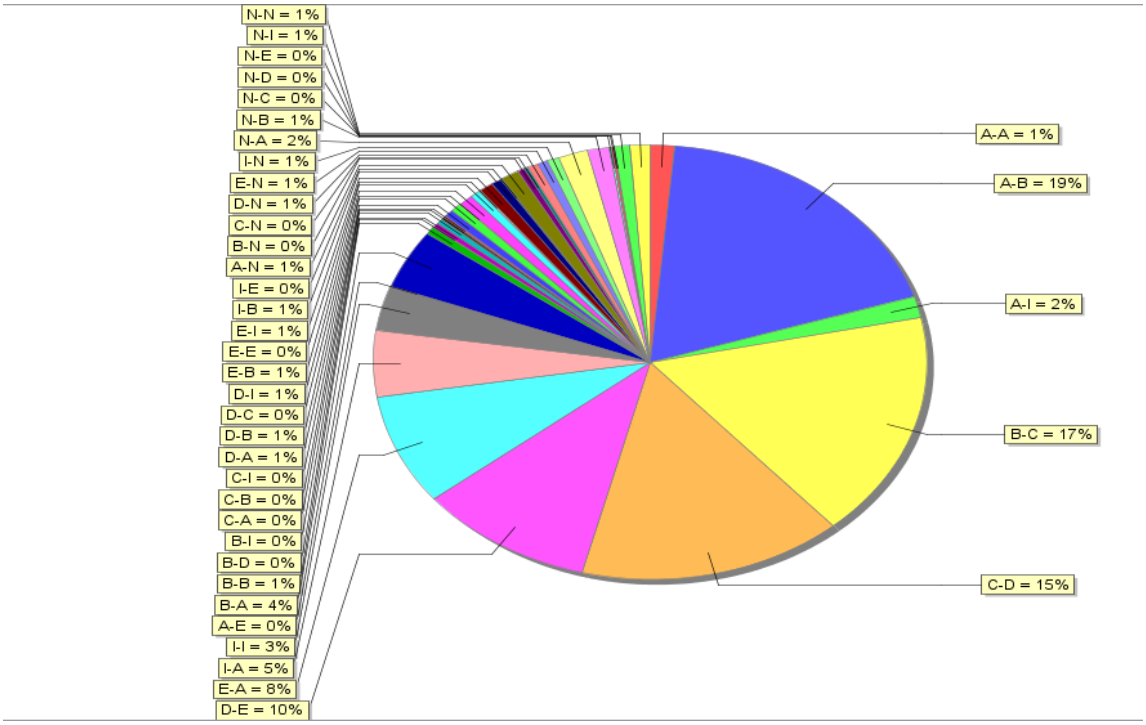
N-N = 1%
N-I = 1%
N-E = 0%
N-D = 0%
N-C = 0%
N-B = 1%
N-A = 2%
I-N = 1%
E-N = 1%
D-N = 1%
C-N = 0%
B-N = 0%
A-N = 1%
I-E = 0%
I-B = 1%
E-I = 1%
E-E = 0%
E-B = 1%
D-I = 1%
D-C = 0%
D-B = 1%
D-A = 1%
C-I = 0%
C-B = 0%
C-A = 0%
B-I = 0%
B-D = 0%
B-B = 1%
B-A = 4%
A-E = 0%
I-I = 3%
I-A = 5%
E-A = 8%
D-E = 10%
A-A = 1%
A-B = 19%
A-I = 2%
B-C = 17%
C-D = 15%

Fig. 15

Linearity, consistency, and stereotypy scores for the target file selected can be visualized under the "Show Scores" button (**Fig. 16**). For every .wav file in the spreadsheet of day 12, the number of transitions is calculated as well as the number of transition types (the number of different transitions). Corresponding linearity, consistency and stereotypy scores are then generated based on these values. For any score that is larger than 0.85, the entry in the table is colored red, indicating that the song for this .wav file is highly stereotyped. Average linearity, consistency and stereotypy scores are then calculated and shown at the top.

Average Linearity Score = 0.6497

Average Consistency Score = 0.805

Average Stereotypy Score = 0.7274

.Wav File	Number of Transitions	Number of Transition Types	Linearity Score	Consistency Score	Stereotypy Score
00001.wav	3	3	0.6667	1.0	0.8333
00002.wav	23	10	0.7	0.8696	0.7848
00003.wav	18	11	0.6364	0.7778	0.7071
00004.wav	21	10	0.7	0.7143	0.7071
00005.wav	28	10	0.6	0.8929	0.7464
00006.wav	48	11	0.5455	0.875	0.7102
00007.wav	16	10	0.7	0.6875	0.6938
00008.wav	26	11	0.5455	0.8077	0.6766
00009.wav	18	10	0.7	0.6667	0.6833
00010.wav	17	8	0.75	0.7647	0.7574
00011.wav	22	9	0.6667	0.8636	0.7652
00012.wav	28	10	0.7	0.8571	0.7786
00013.wav	25	10	0.6	0.88	0.74
00014.wav	23	10	0.7	0.7826	0.7413
00015.wav	31	13	0.5385	0.7742	0.6563
00016.wav	18	9	0.7778	0.8889	0.8333
00017.wav	15	11	0.6364	0.8	0.7182
00018.wav	17	8	0.875	0.7647	0.8199
00019.wav	17	10	0.7	0.8235	0.7618
00020.wav	34	14	0.5	0.7941	0.6471
00021.wav	23	9	0.6667	0.8696	0.7681
00022.wav	17	10	0.7	0.7059	0.7029
00023.wav	27	13	0.4615	0.6296	0.5456
00024.wav	22	12	0.5833	0.6364	0.6098
00025.wav	17	9	0.6667	0.7647	0.7157
00026.wav	39	18	0.3889	0.6667	0.5278
00027.wav	26	11	0.6364	0.8077	0.722
00028.wav	21	9	0.7778	0.9048	0.8413
00029.wav	19	12	0.5833	0.6842	0.6338

Fig. 16

All of the above results are written onto the target file's spreadsheet. This can be seen by clicking "Open Excel File" button. **Fig. 17** shows a screen shot of the Pre day target excel spreadsheet after the simulation. The screenshot captures only its internal spreadsheets names. SongSeq creates multiple spreadsheets within the target spreadsheet. A spreadsheet is created for each syllable entered by the user, along with a spreadsheet for the "NEW" or "NMS" syllables. Next these syllable spreadsheets are merged into a final "Sequenced" spreadsheet that is created to contain all the notes in the order they were sung. This allows for the analysis of syllable transition probabilities.

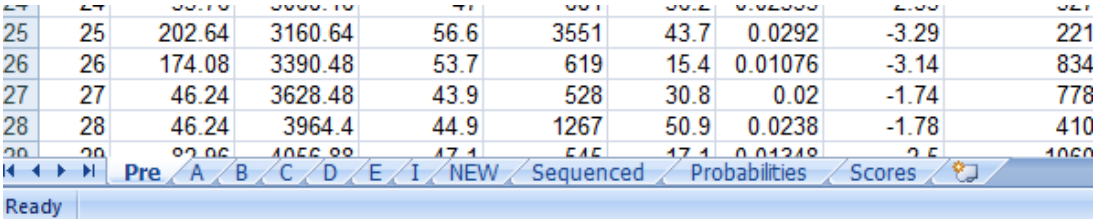

Fig. 17

The transition probabilities shown in **Fig 13**, **14** and **15** are written under the "Probabilities" spreadsheet (**Fig. 18**). In the same order as in the table of **Fig. 13**, the data are written onto their corresponding rows and columns in the spreadsheet. This enables user manipulation of the data on Excel apart from SongSeq. The typical transitions are written first, then atypical transitions followed by NMS transitions.

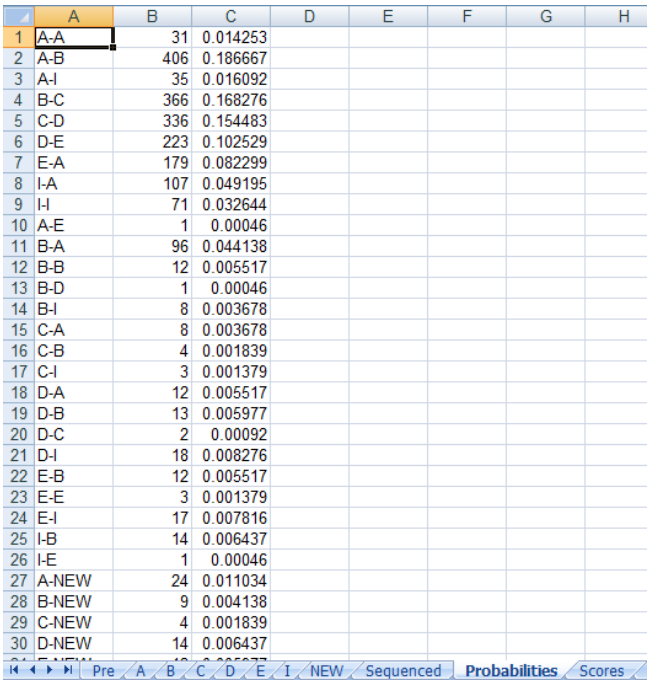

	A	B	C
1	A-A	31	0.014253
2	A-B	406	0.186667
3	A-I	35	0.016092
4	B-C	366	0.168276
5	C-D	336	0.154483
6	D-E	223	0.102529
7	E-A	179	0.082299
8	I-A	107	0.049195
9	I-I	71	0.032644
10	A-E	1	0.00046
11	B-A	96	0.044138
12	B-B	12	0.005517
13	B-D	1	0.00046
14	B-I	8	0.003678
15	C-A	8	0.003678
16	C-B	4	0.001839
17	C-I	3	0.001379
18	D-A	12	0.005517
19	D-B	13	0.005977
20	D-C	2	0.00092
21	D-I	18	0.008276
22	E-B	12	0.005517
23	E-E	3	0.001379
24	E-I	17	0.007816
25	I-B	14	0.006437
26	I-E	1	0.00046
27	A-NEW	24	0.011034
28	B-NEW	9	0.004138
29	C-NEW	4	0.001839
30	D-NEW	14	0.006437

Pre | A | B | C | D | E | I | NEW | Sequenced | Probabilities | Scores

Fig. 18

Similarly, the scores (**Fig. 16**) are written under the spreadsheet labeled "Scores" as shown in **Fig. 19** below.

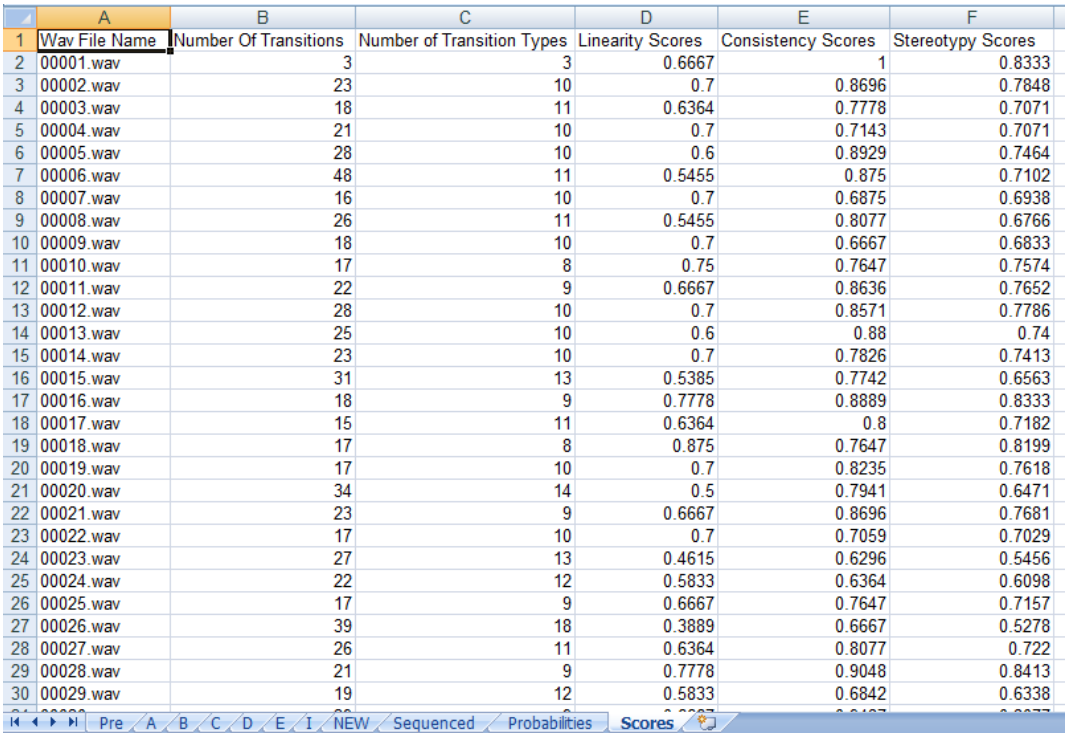

	A	B	C	D	E	F
1	Wav File Name	Number Of Transitions	Number of Transition Types	Linearity Scores	Consistency Scores	Stereotypy Scores
2	00001.wav	3	3	0.6667	1	0.8333
3	00002.wav	23	10	0.7	0.8696	0.7848
4	00003.wav	18	11	0.6364	0.7778	0.7071
5	00004.wav	21	10	0.7	0.7143	0.7071
6	00005.wav	28	10	0.6	0.8929	0.7464
7	00006.wav	48	11	0.5455	0.875	0.7102
8	00007.wav	16	10	0.7	0.6875	0.6938
9	00008.wav	26	11	0.5455	0.8077	0.6766
10	00009.wav	18	10	0.7	0.6667	0.6833
11	00010.wav	17	8	0.75	0.7647	0.7574
12	00011.wav	22	9	0.6667	0.8636	0.7652
13	00012.wav	28	10	0.7	0.8571	0.7786
14	00013.wav	25	10	0.6	0.88	0.74
15	00014.wav	23	10	0.7	0.7826	0.7413
16	00015.wav	31	13	0.5385	0.7742	0.6563
17	00016.wav	18	9	0.7778	0.8889	0.8333
18	00017.wav	15	11	0.6364	0.8	0.7182
19	00018.wav	17	8	0.875	0.7647	0.8199
20	00019.wav	17	10	0.7	0.8235	0.7618
21	00020.wav	34	14	0.5	0.7941	0.6471
22	00021.wav	23	9	0.6667	0.8696	0.7681
23	00022.wav	17	10	0.7	0.7059	0.7029
24	00023.wav	27	13	0.4615	0.6296	0.5456
25	00024.wav	22	12	0.5833	0.6364	0.6098
26	00025.wav	17	9	0.6667	0.7647	0.7157
27	00026.wav	39	18	0.3889	0.6667	0.5278
28	00027.wav	26	11	0.6364	0.8077	0.722
29	00028.wav	21	9	0.7778	0.9048	0.8413
30	00029.wav	19	12	0.5833	0.6842	0.6338

Pre | A | B | C | D | E | I | NEW | Sequenced | Probabilities | Scores

Fig. 19

The generated spreadsheets described above should be left unaltered by the user; this is required for other modules in SongSeq to operate successfully. One can create a copy of the spreadsheets in a separate file and manipulate the data externally.

The KL analysis, entropy analysis and the average scores for all the target files simulated can be seen by pressing on the buttons “Show K-L Analysis”, “Show Entropies” and “Show All Scores” buttons respectively (**Fig. 13**). The details of this analysis are explained in Daou et al (2012).

e- Identification using Two Pairs of Features

In **Fig. 6**, notice that syllables B (blue) and I (magenta) are very close to each other, and thereby not very efficiently discriminated. If such a situation occurs, there can be a better identification of the syllables with a second pair of features. In order to explore that, click on “Add Another Dimension and Move to Step 3” button. This brings a window similar to **Fig. 20**. One can identify the first set of syllables using one feature pair, then move to a second feature pair to improve the identification of the remaining syllables. In this process, and after all the syllables are identified with the first 2 acoustic features, two new acoustic features are chosen along with a subset of the previously added syllables. Only this chosen subset of syllables (which can include all syllables) can be painted again in the new 2D feature space. **Fig. 20** shows the selection of syllable duration versus mean entropy acoustic features. It also shows the selection of B and I since their clusters in **Fig. 6** are not well isolated, and as we will see shortly syllable duration versus mean entropy pulls B and I apart. Syllables A, C, D and E are not chosen here because they are discriminated nicely with the first two acoustic features.

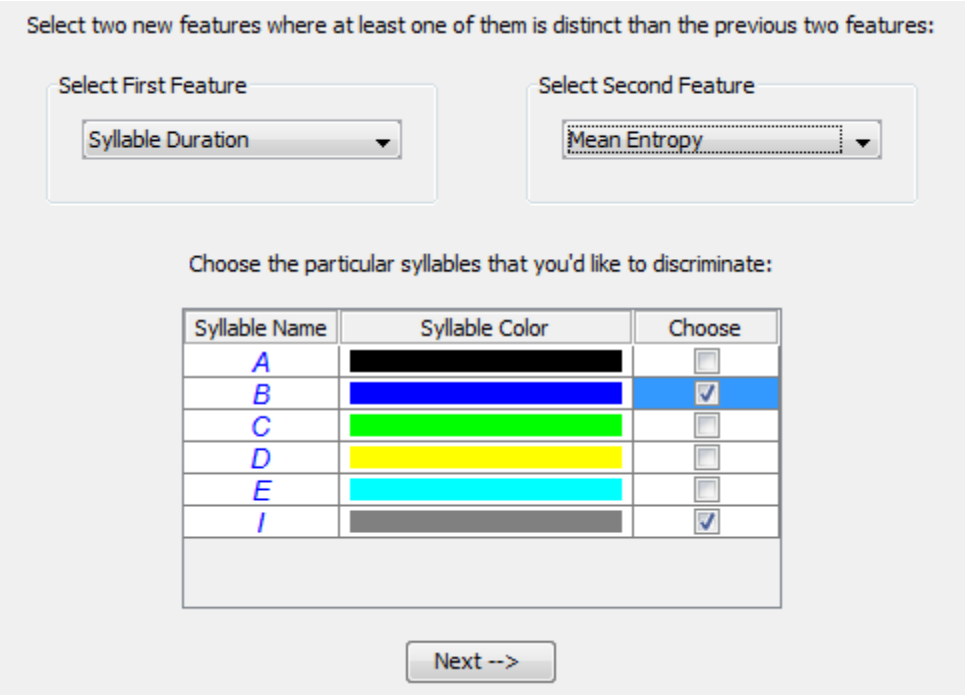

Fig. 20

Once the second pair of features is chosen, a new 2D scatter plot appears. The data points are now color coded according to the first step of identification (**Fig. 21**). For instance, any data points that were within the black boxes defining the borders of syllable A in the scatter plot of **Fig. 6**, will be painted in the same color (black) in the scatter plot of **Fig. 21**. In Fig. 21 we see that the data points for B (blue) and I (magenta) are pulled apart in the new 2D feature space, and it is clear that some data points previously misidentified as B actually cluster better with I (blue points in the bottom right of the magenta cluster). Syllables A (black), C (green) and D (yellow) remain well isolated, but syllable E (gray) has a very similar mean entropy as the introductory notes and therefore their corresponding clusters overlap.

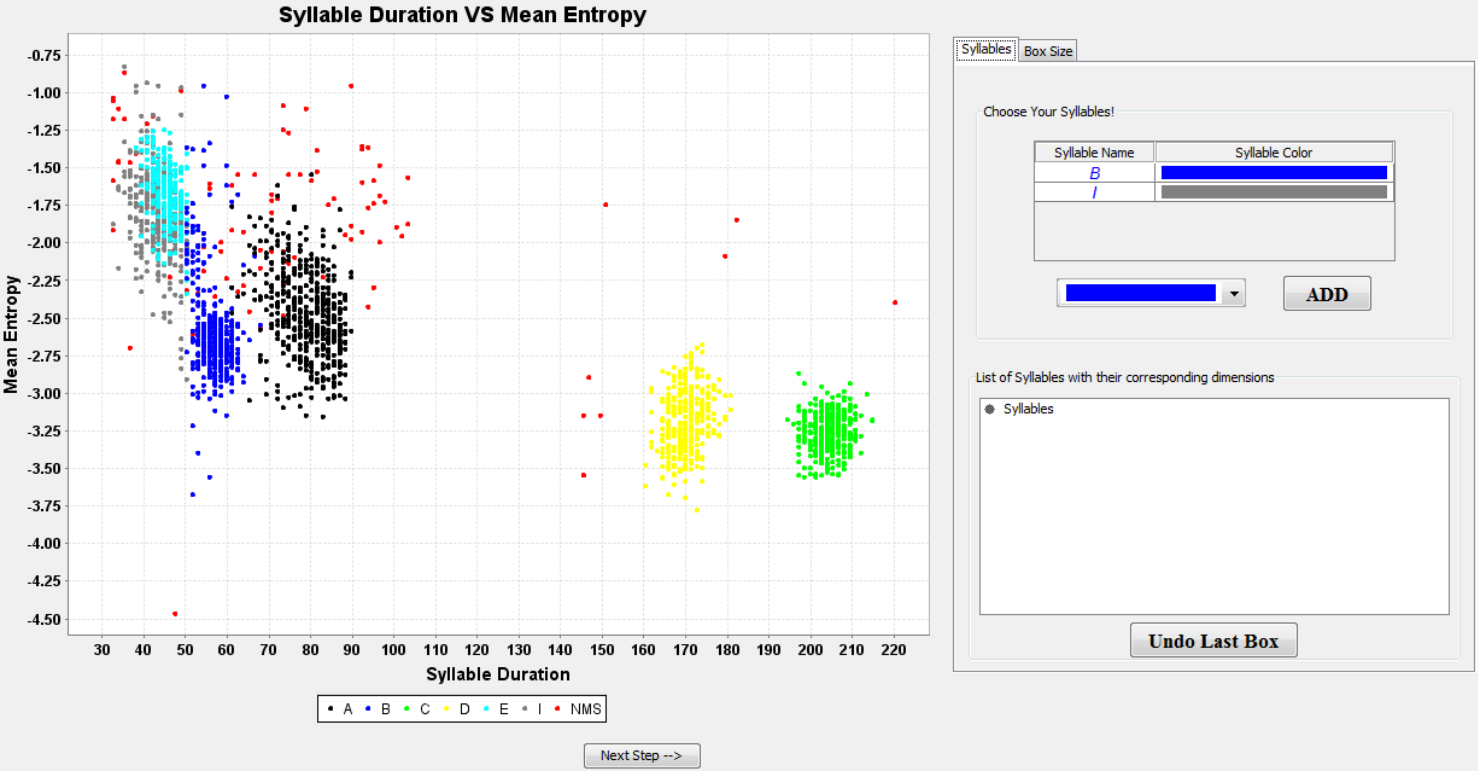

Fig. 21

Next, the user paints in the same way as before onto the new scatter plot, but now only for the syllables (B and I) that were specified for this second round of discrimination (**Fig. 22**). For these specified syllables, a data point is considered as syllable X if and only if either of the following two conditions hold: (1) it belongs to one of the boxes defining syllable X in the first 2D scatter plot AND it belongs to one of the boxes defining syllable X in the second 2D scatter plot, (2) it belongs to one of the boxes defining syllable X in the second 2D scatter plot BUT is not a data point of a non-specified syllable (that is, a data point that is colored with black (A), green (C), yellow (D) or gray (E)). Although the clusters for E and I overlap in the second scatter plot, the points originally labeled as E retain that identification since E was not selected for repainting.

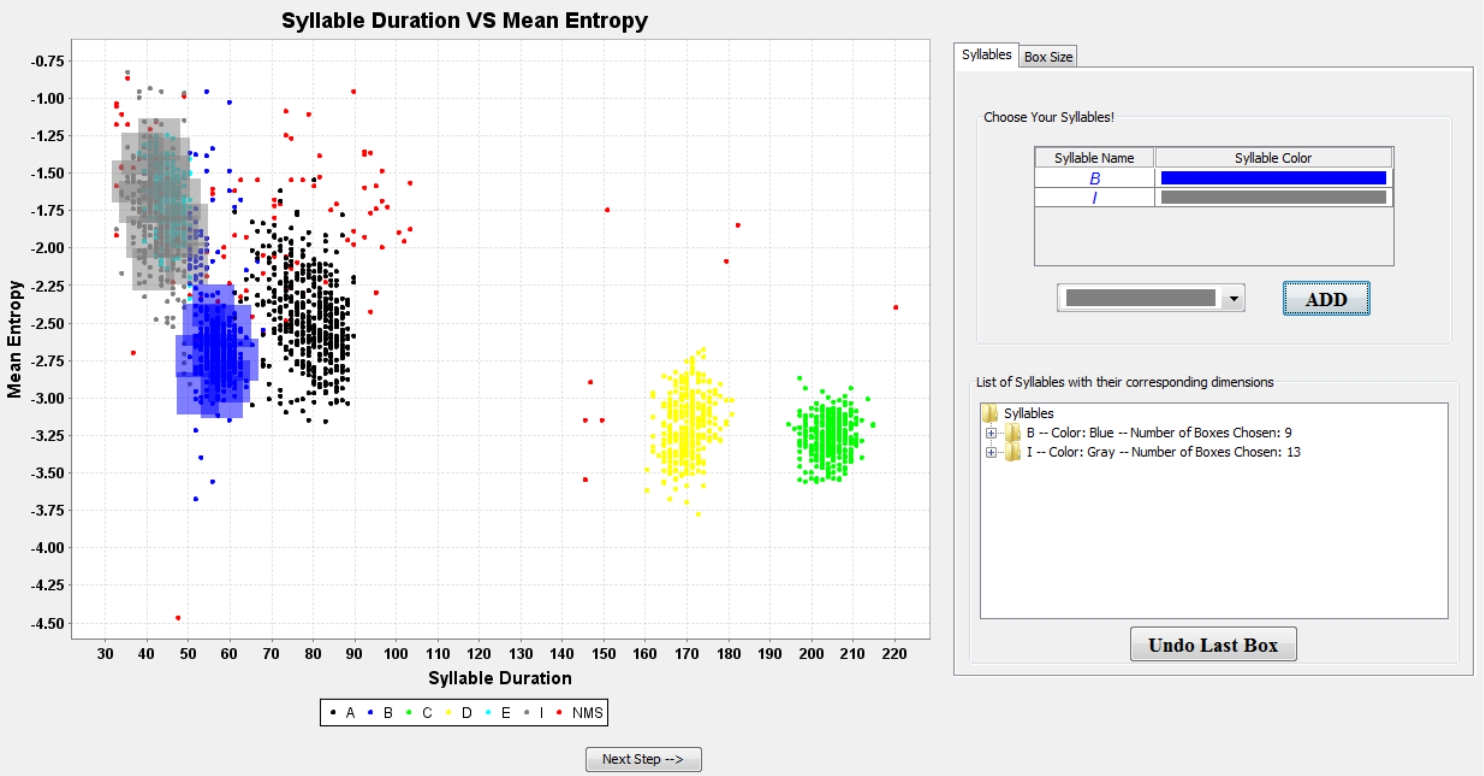

Fig. 22

The next step in choosing the target files and visualizing the results follows the same as previously discussed (**Figs 7** through **19**).

f- Modified Template Spreadsheet

Just as SongSeq modifies each of the target files spreadsheets as described above, the template spreadsheet is also modified, but in a different way. **Fig. 23** shows a screenshot of the template spreadsheet after simulation, zoomed at the spreadsheet's taskbar showing its internal generated spreadsheets. Three spreadsheets are generated by SongSeq, named "Syllable Dimensions", "Typical Transitions" and "Axis Ranges".

22	628	205.36	3874.64	56.9	3287	43.2	0.03228	-3.51
23	628	168.64	4108.56	52.8	583	14.6	0.0088	-3.19
24	628	32.64	847.28	41.4	940	48.6	0.02205	-2.13
25	628	38.08	1066.24	42.2	900	44	0.02481	-1.99
26	628	38.08	1266.16	43	982	40.5	0.02575	-2.06
27	628	68	1472.88	42.3	523	19.3	0.00996	-2.75
28	628	44.88	1610.24	43.1	1130	41.7	0.01942	-2.07

Sheet1 | Syllable Dimensions | Typical Transitions | Axis Ranges

Fig. 23

The "Syllable Dimensions" spreadsheet (**Fig. 24**) contains information of syllable names, the colors used to paint each syllable onto the scatter plot, the dimensions of the pair of features defined by each painted box, and other information required for JAVA processing in the "Load Template & Sequence" section described later, particularly the location on the screen (in pixels) of each mouse click used to paint a box and the width of the box painted.

	A	B	C	D	E	F	G	H	I	J
1	Syllable Name	Syllable Color	Syllable Duration FROM	Syllable Duration TO	Mean Pitch Goodness FROM	Mean Pitch Goodness TO	e.getX()	e.getY()	Window Width	
2	A	black	74.554	87.637	637.899	766.34	250	282	45	
3	A	black	69.031	82.113	777.757	906.199	231	233	45	
4	A	black	71.357	84.439	720.672	849.114	239	253	45	
5	A	black	78.334	91.416	729.235	857.676	263	250	45	
6	A	black	76.008	89.09	783.466	911.907	255	231	45	
7	A	black	74.264	87.346	840.551	968.992	249	211	45	
8	A	black	73.973	87.055	891.927	1020.369	248	193	45	
9	A	black	72.229	85.311	929.033	1057.474	242	180	45	
10	A	black	66.705	79.787	880.51	1008.952	223	197	45	
11	A	black	63.216	76.299	957.575	1086.016	211	170	45	
12	A	black	62.926	76.008	1000.389	1128.83	210	155	45	
13	A	black	62.054	75.136	1060.328	1188.769	207	134	45	
14	A	black	63.216	76.299	1105.996	1234.437	211	118	45	
15	A	black	68.74	81.822	1105.996	1234.437	230	118	45	
16	A	black	70.775	83.857	1051.765	1180.207	237	137	45	
17	B	blue	56.239	69.322	181.219	309.66	187	442	45	
18	B	blue	48.39	61.472	186.927	315.369	160	440	45	
19	B	blue	48.681	61.763	298.243	426.685	161	401	45	
20	B	blue	53.623	66.705	286.826	415.268	178	405	45	
21	C	green	193.167	206.249	109.863	238.304	658	467	45	
22	C	green	194.039	207.122	198.344	326.786	661	436	45	
23	C	green	203.342	216.425	201.199	329.64	693	435	45	
24	C	green	202.761	215.843	124.134	252.575	691	462	45	
25	D	yellow	162.642	175.724	609.357	737.798	553	292	45	
26	D	yellow	160.316	173.398	660.733	789.174	545	274	45	
27	D	yellow	157.7	170.782	734.944	863.385	536	248	45	
28	D	yellow	157.118	170.201	814.863	943.304	534	220	45	
29	D	yellow	160.316	173.398	894.782	1023.223	545	192	45	
30	D	yellow	163.805	176.887	957.575	1086.016	557	170	45	

Sheet1 | Syllable Dimensions | Typical Transitions | Axis Ranges

Fig. 24

The typical transitions are written onto the spreadsheet "Typical Transitions" (**Fig. 25**).

	A	B	C	D	E	F	G	H
1	A-A							
2	A-B							
3	A-I							
4	B-C							
5	C-D							
6	D-E							
7	E-A							
8	I-A							
9	I-I							

Sheet1 | Syllable Dimensions | Typical Transitions | Axis Ranges

Fig. 25

“Axis Ranges” spreadsheet (**Fig. 26**) contains information about the x (first feature) and y (second feature) axis of the template’s 2D scatter plot, particularly, the range of both axis. This information is required for the “Load Template & Sequence” module.

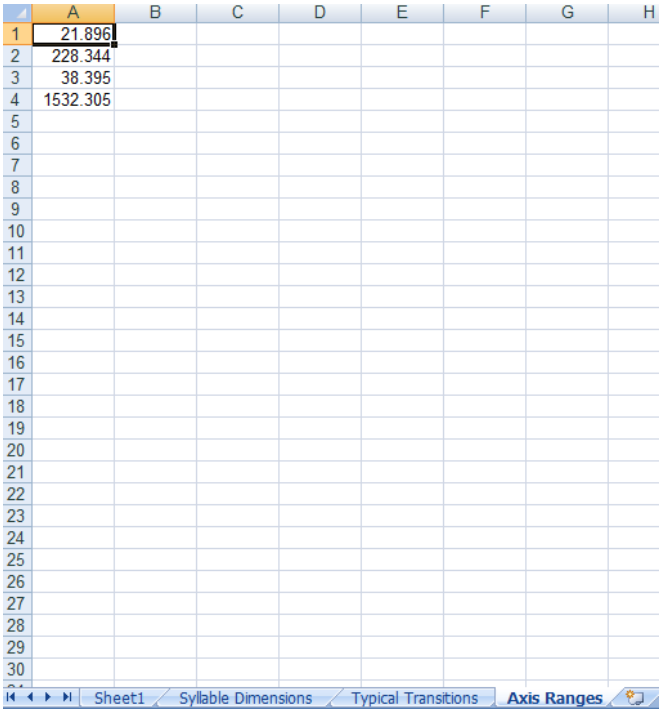

Fig. 26

It’s important that the information in the modified spreadsheets (target’s and template’s) to be left unaltered by the user. This information is required by other modules of SongSeq, and any user modification of the spreadsheets might result in unsuccessful execution of tasks by SongSeq. It’s recommended to create a copy of the spreadsheet that wishes user manipulation.

Load Template and Sequence

After the boundaries for all syllable clusters are painted and defined through the "Make Template and Sequence" module, and after SongSeq sequences few target files, it may happen that one would like to run the same simulation (with the same template and syllable boundaries) over some other target files at a later time (for example, after you get data from a post-surgery day). It would be inconvenient to paint the syllable boundaries again, because each time one paints it is most likely that the boxes will be painted in a different way. To overcome this problem, "Load Template and Sequence" module allows you to load the previously simulated template (that has undergone painting previously) thereby maintaining the same painted boxes and syllable boundaries. This of course requires that the new target files to be simulated are to be compared with the designated template spreadsheet.

Fig. 27 shows the first step in launching the "Load Template and Sequence" module.

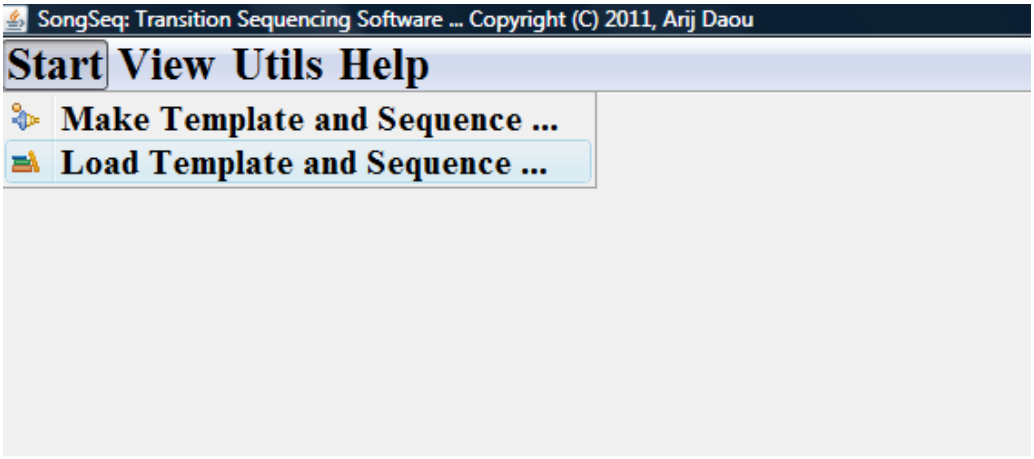

Fig. 27

The next step is browsing for the previously simulated template (MAIN) excel spreadsheet (**Fig. 28**).

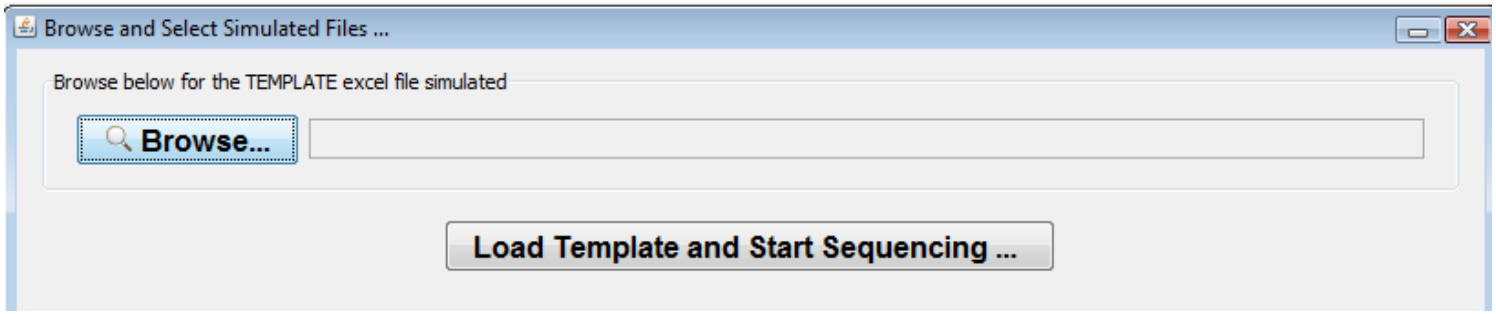

Fig. 28

Fig. 29 below shows the generated window after loading the same template used in the simulation above. The scatter plot with the chosen two features (**Fig. 3**) and the painted boxes (**Figs. 4**, **5**, and **6**) are shown. Also shown on the right side of the panel are the syllable names and their associated colors, along with the typical transitions entered by the user previously. This information is extracted from the spreadsheet of the template as described before (**Figs. 24** through **26**).

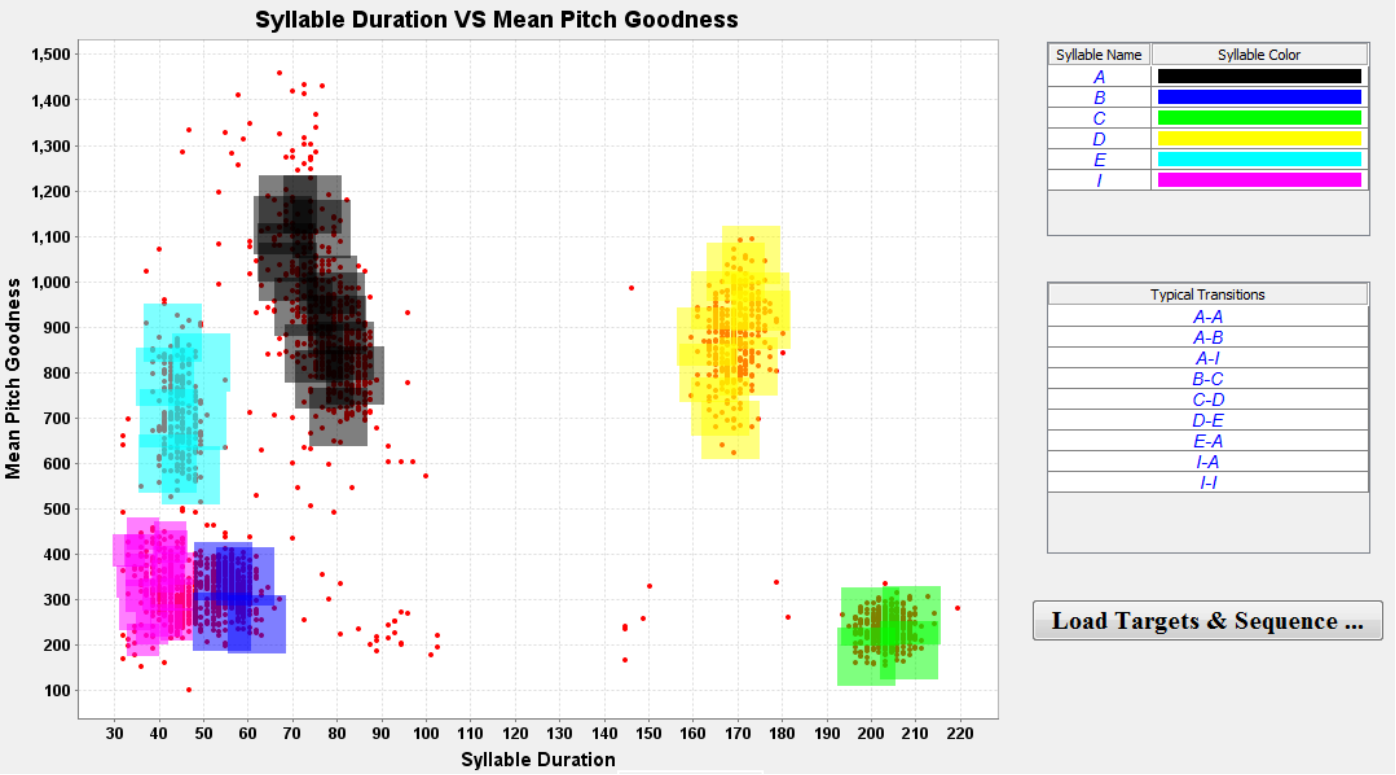

Fig. 29

The window above (**Fig. 29**) just gives the user an idea of what had been done before. No extra painting can be done here and no new syllables or typical transitions can be added. Next, the target files can be added through "Load Targets & Sequence". This will load **Fig. 30** below. In this step, the target files are uploaded and processed in the same way as discussed before (**Figs. 7** through **10**).

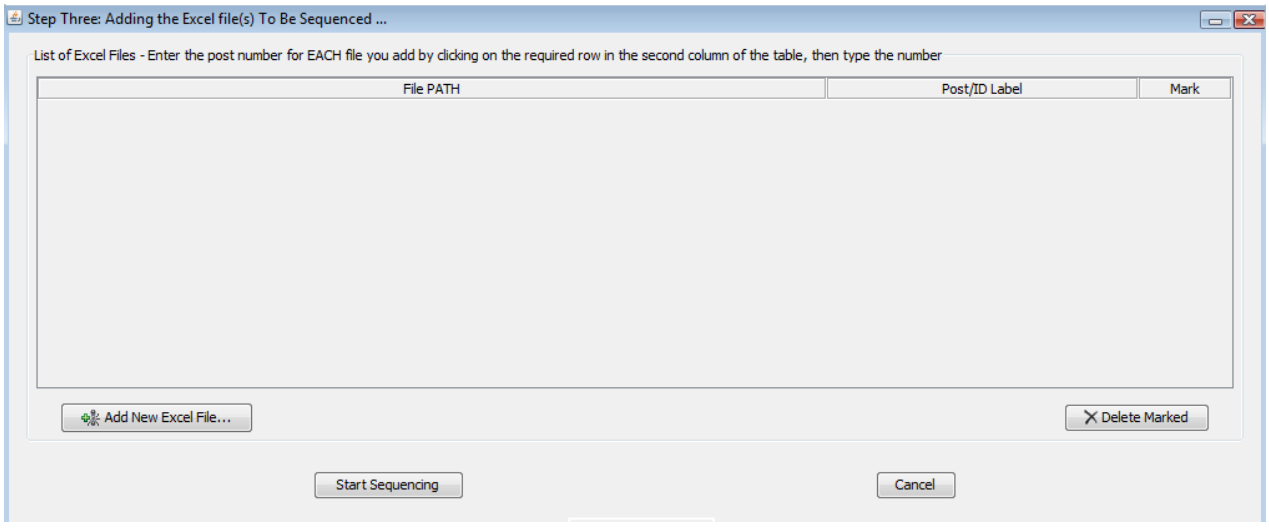

Fig. 30

View Scatter Plots

"Scatter Plots" module (**Fig. 31**) provides an efficient tool allowing the user to browse easily between the different pairs of features to find the best two features that discriminate the syllables efficiently.

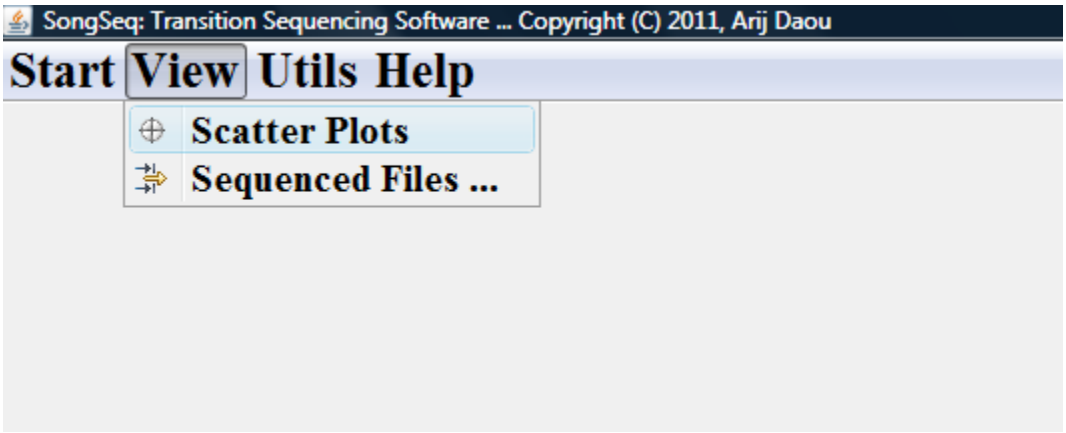

Fig. 31

After the spreadsheet is browsed for (**Fig. 32**), the user can choose any two acoustic features to view their scatter plot. The window in Fig. 32 remains allowing the user to keep choosing different combinations of features. It's very easy and convenient using this module to look at the different pairs of features.

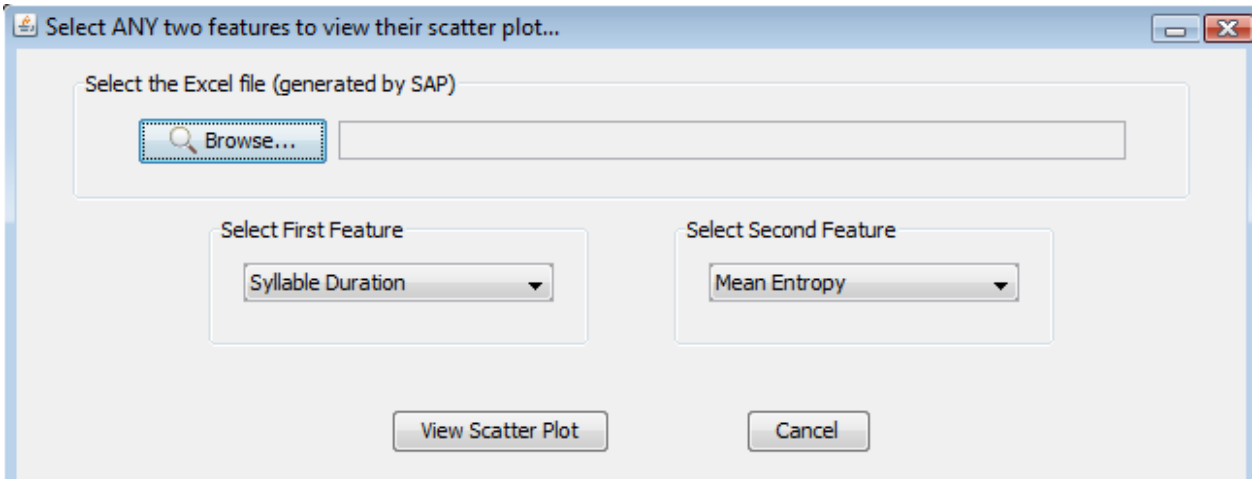

Fig. 32

Fig. 33 shows a sample scatter plot window. Also, the panel gives the option to save the scatter plot as an image using the button "Save Graph".

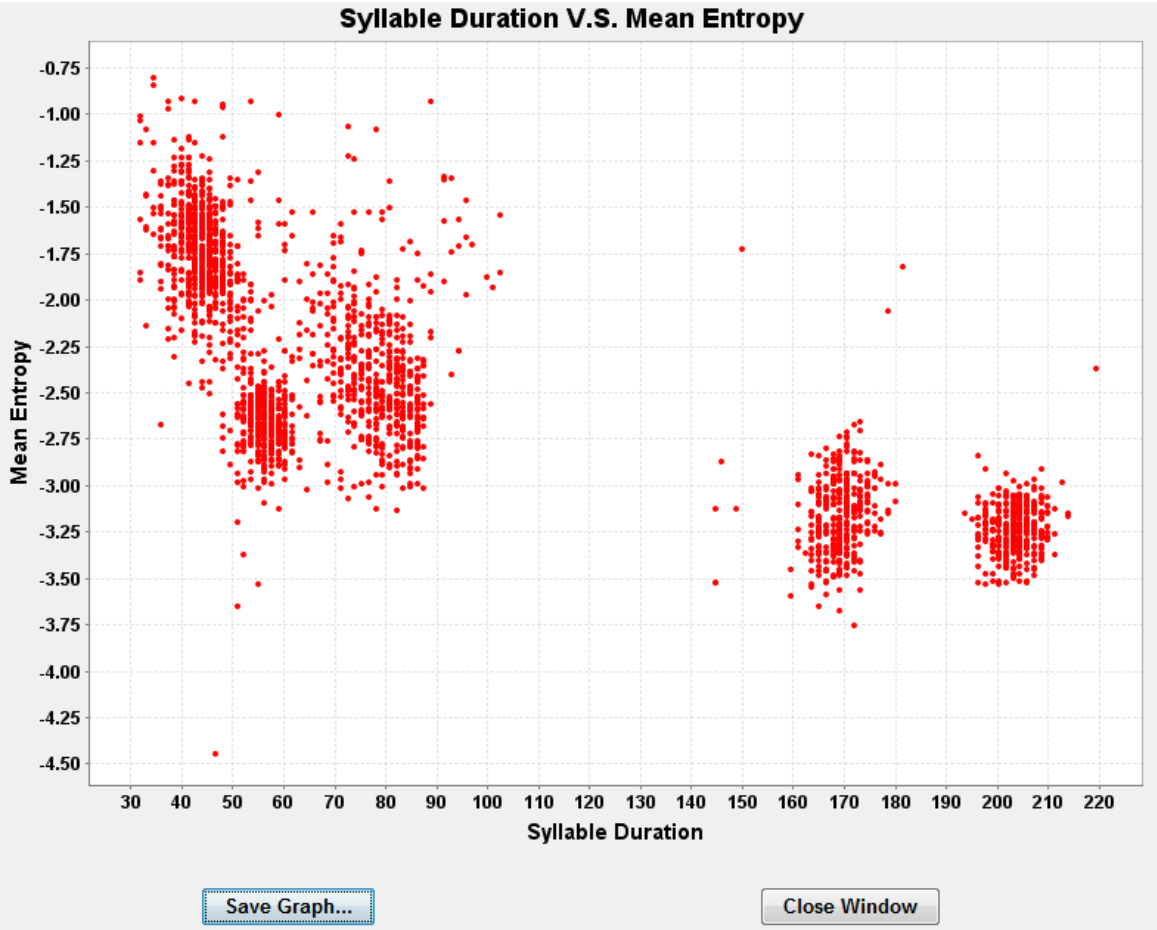
Syllable Duration V.S. Mean Entropy
Mean Entropy
-0.75
-1.00
-1.25
-1.50
-1.75
-2.00
-2.25
-2.50
-2.75
-3.00
-3.25
-3.50
-3.75
-4.00
-4.25
-4.50
30 40 50 60 70 80 90 100 110 120 130 140 150 160 170 180 190 200 210 220
Syllable Duration
Save Graph...
Close Window

Fig. 33

View Sequenced Files

"Sequenced Files" module (**Fig. 34**) allows one to view previously simulated files and recall the results.

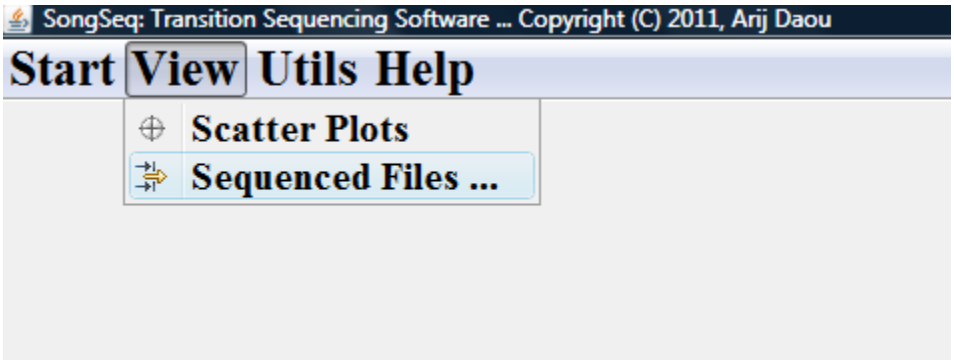

Fig. 34

The user is asked to browse for the template spreadsheet (through "Browse" button) and the target files spreadsheets (through "Add Target(s)" button) that had been run previously through SongSeq (**Fig. 35**).

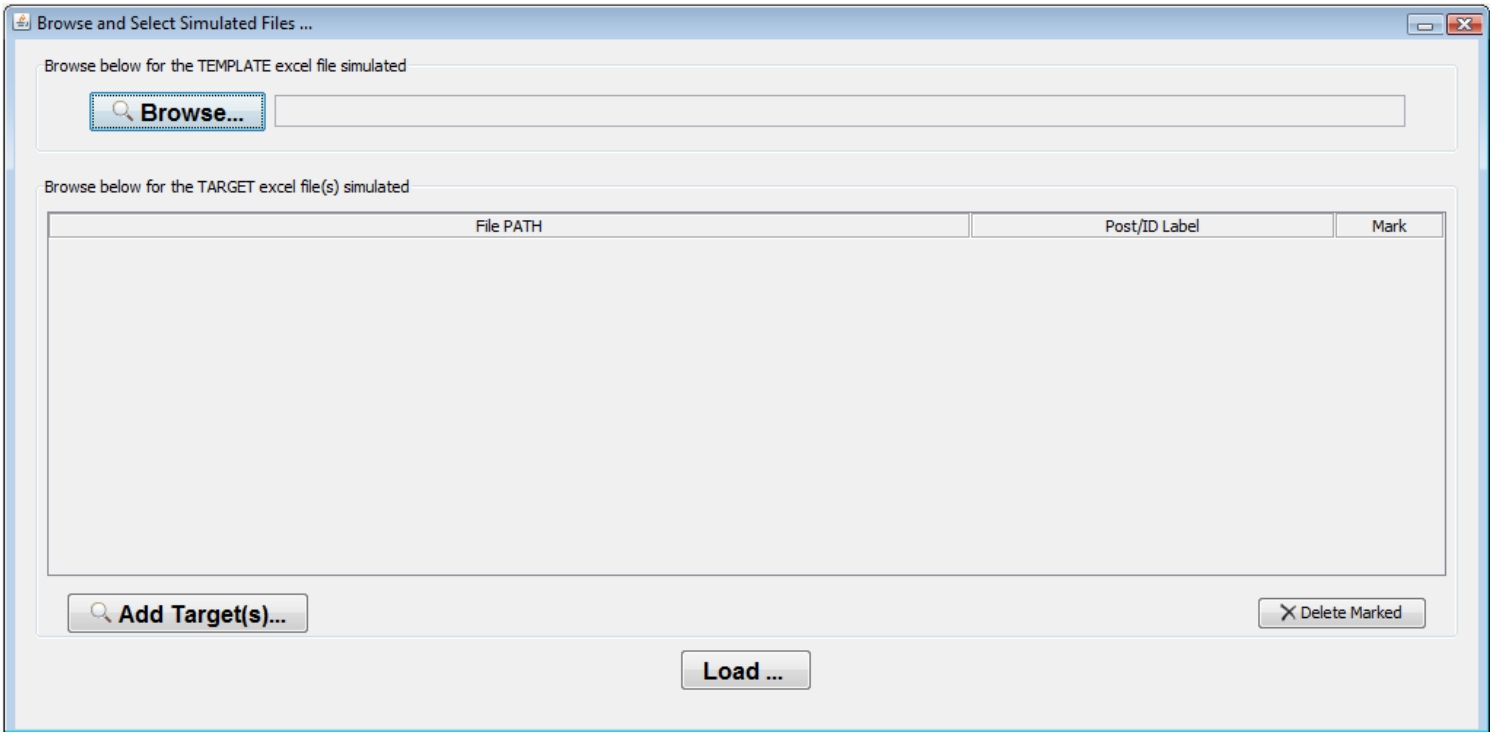

Fig. 35

Fig. 36 shows the window after loading the template and the target spreadsheets. Notice that SongSeq remembers the target's name or ID (second column of the table) as the user entered it previously when the targets were simulated (**Fig. 10**).

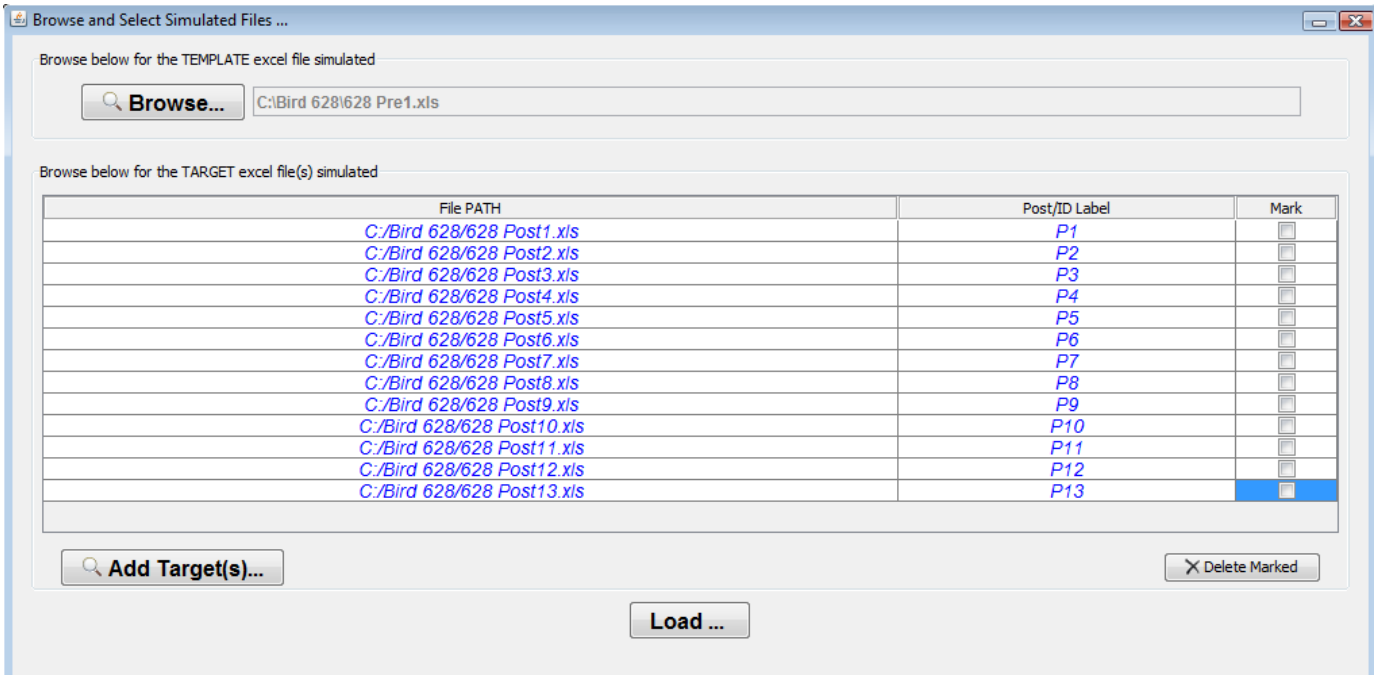

Fig. 36

Clicking on the "Load" button opens a window (**Fig. 37**) that is similar to **Fig. 29**, where the template scatter plot is shown along with the painted boxes, the syllable names and their associated colors, and typical transitions. Also shown on the lower right corner is a drop down menu showing a list of names of all the target files uploaded. Choosing one target file and hitting "View" loads the results for that target file chosen.

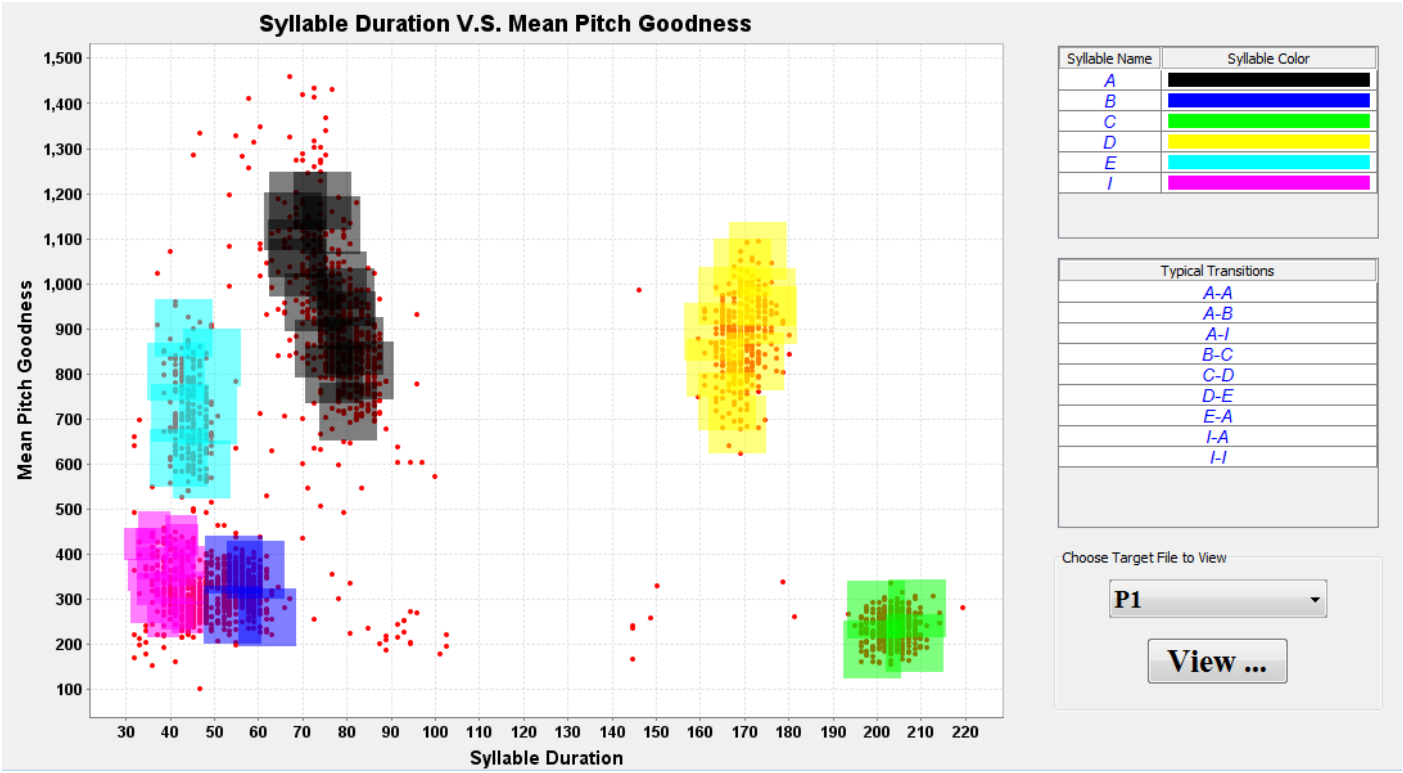

Fig. 37

Fig. 38 shows a sample window for the results that are displayed (in this case for target file "P1"). Three tabbed panes are displayed, showing the simulated scores, and the transitional probabilities (in bar chart and pie chart formats).

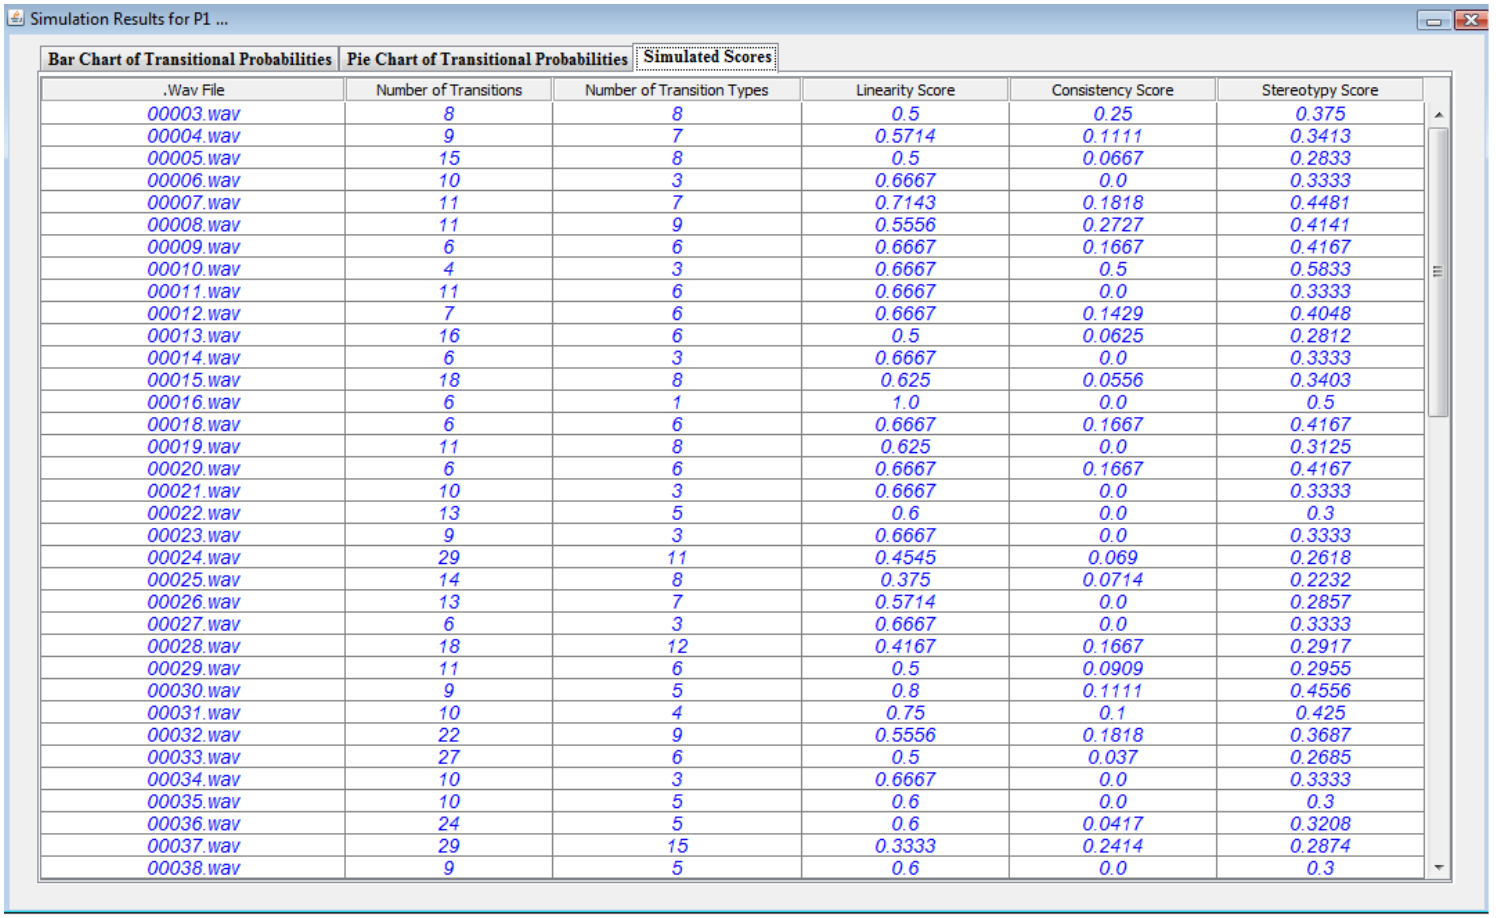

.Wav File	Number of Transitions	Number of Transition Types	Linearity Score	Consistency Score	Stereotypy Score
00003.wav	8	8	0.5	0.25	0.375
00004.wav	9	7	0.5714	0.1111	0.3413
00005.wav	15	8	0.5	0.0667	0.2833
00006.wav	10	3	0.6667	0.0	0.3333
00007.wav	11	7	0.7143	0.1818	0.4481
00008.wav	11	9	0.5556	0.2727	0.4141
00009.wav	6	6	0.6667	0.1667	0.4167
00010.wav	4	3	0.6667	0.5	0.5833
00011.wav	11	6	0.6667	0.0	0.3333
00012.wav	7	6	0.6667	0.1429	0.4048
00013.wav	16	6	0.5	0.0625	0.2812
00014.wav	6	3	0.6667	0.0	0.3333
00015.wav	18	8	0.625	0.0556	0.3403
00016.wav	6	1	1.0	0.0	0.5
00018.wav	6	6	0.6667	0.1667	0.4167
00019.wav	11	8	0.625	0.0	0.3125
00020.wav	6	6	0.6667	0.1667	0.4167
00021.wav	10	3	0.6667	0.0	0.3333
00022.wav	13	5	0.6	0.0	0.3
00023.wav	9	3	0.6667	0.0	0.3333
00024.wav	29	11	0.4545	0.069	0.2618
00025.wav	14	8	0.375	0.0714	0.2232
00026.wav	13	7	0.5714	0.0	0.2857
00027.wav	6	3	0.6667	0.0	0.3333
00028.wav	18	12	0.4167	0.1667	0.2917
00029.wav	11	6	0.5	0.0909	0.2955
00030.wav	9	5	0.8	0.1111	0.4556
00031.wav	10	4	0.75	0.1	0.425
00032.wav	22	9	0.5556	0.1818	0.3687
00033.wav	27	6	0.5	0.037	0.2685
00034.wav	10	3	0.6667	0.0	0.3333
00035.wav	10	5	0.6	0.0	0.3
00036.wav	24	5	0.6	0.0417	0.3208
00037.wav	29	15	0.3333	0.2414	0.2874
00038.wav	9	5	0.6	0.0	0.3

Fig. 38

Clustering Test

The "Clustering Test" module (**Fig. 39**) intends to help the user identify the syllable clusters on the scatter plot accurately and eliminate any obscurity regarding whether the cluster seen belongs to one syllable or more. This will be illustrated below.

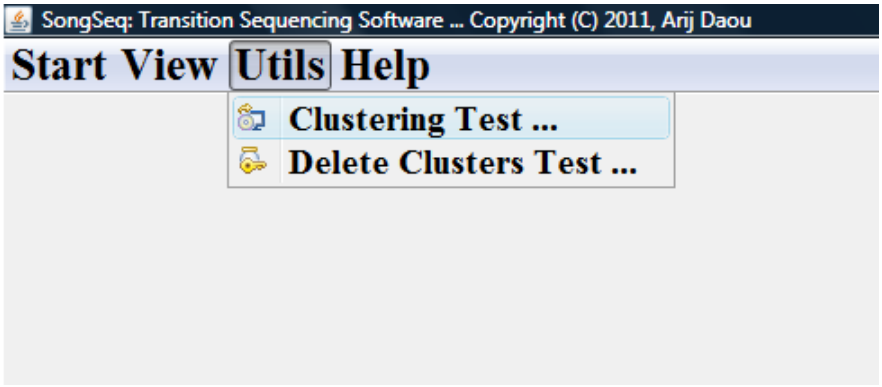

Fig. 39

A window (**Fig. 40**) asks the user to browse for an SA+ spreadsheet and chooses the first pair of acoustic features.

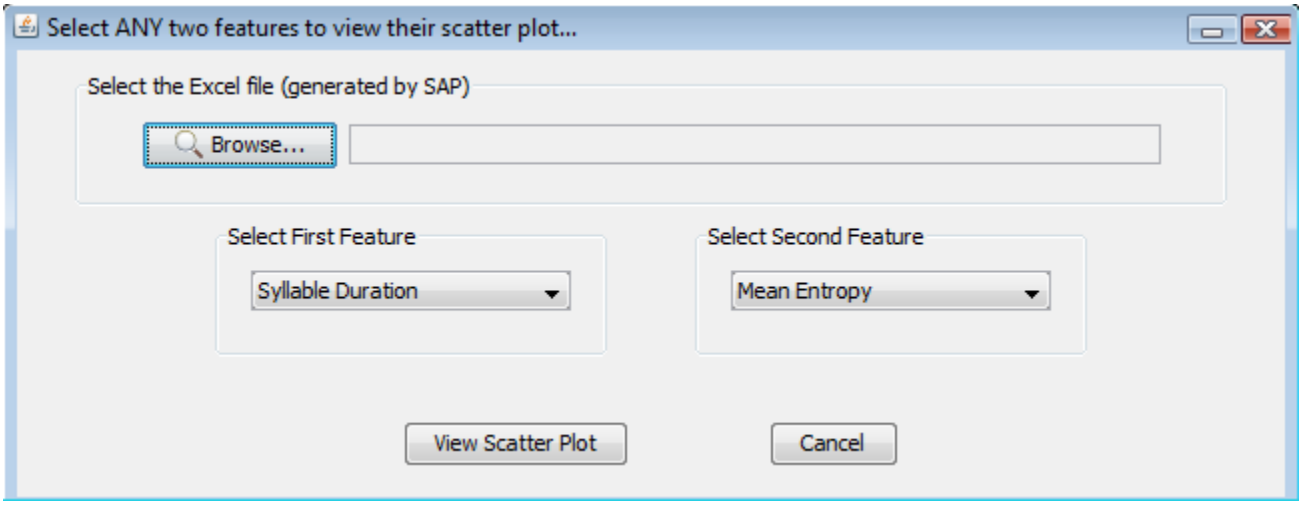

Fig. 40

This loads a window (**Fig. 41**) similar to **Fig. 4**. In this window, the user paints the syllables as descried earlier and then clicks on "Move Next" (**Fig. 42**).

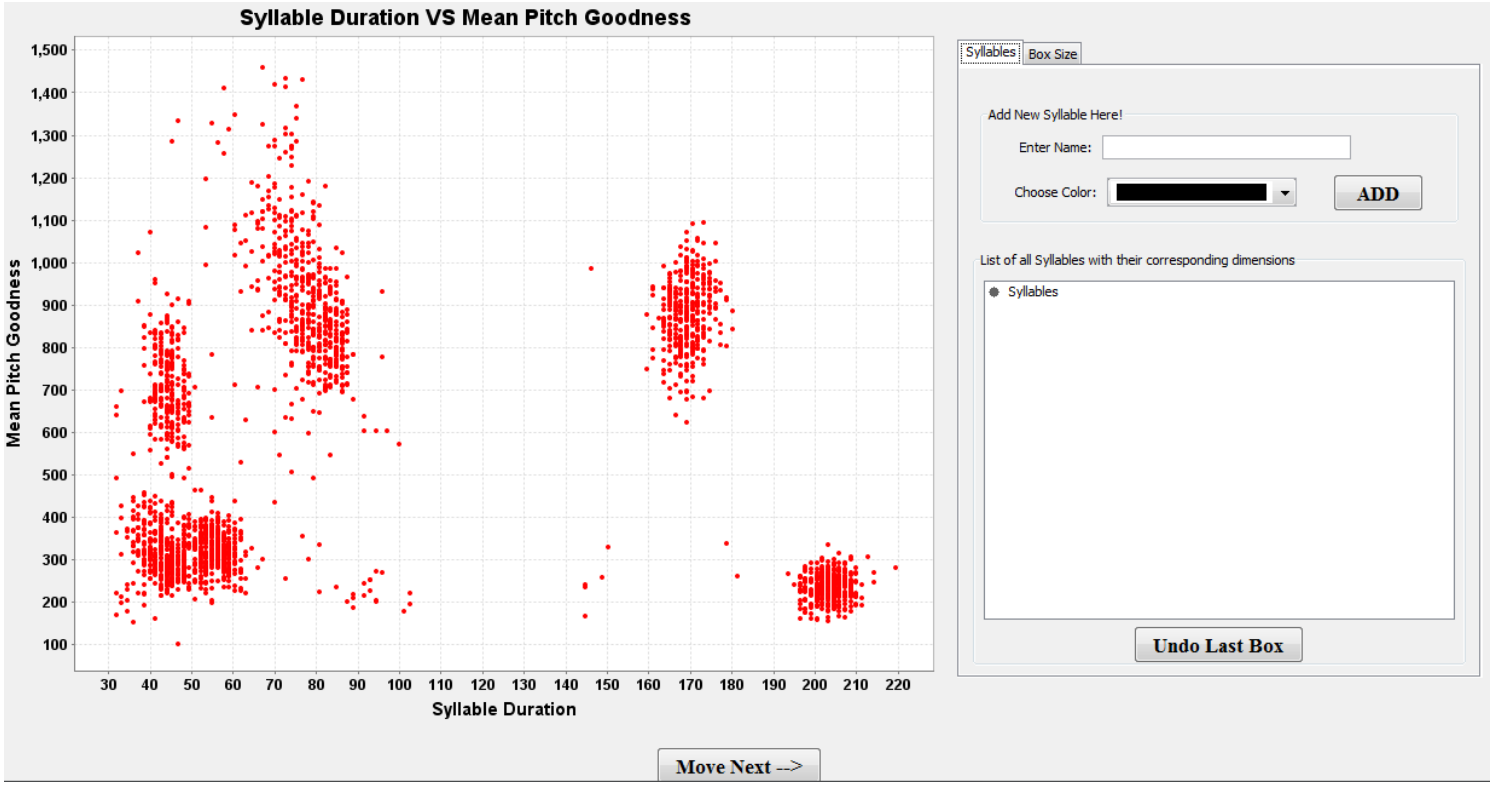

Fig. 41

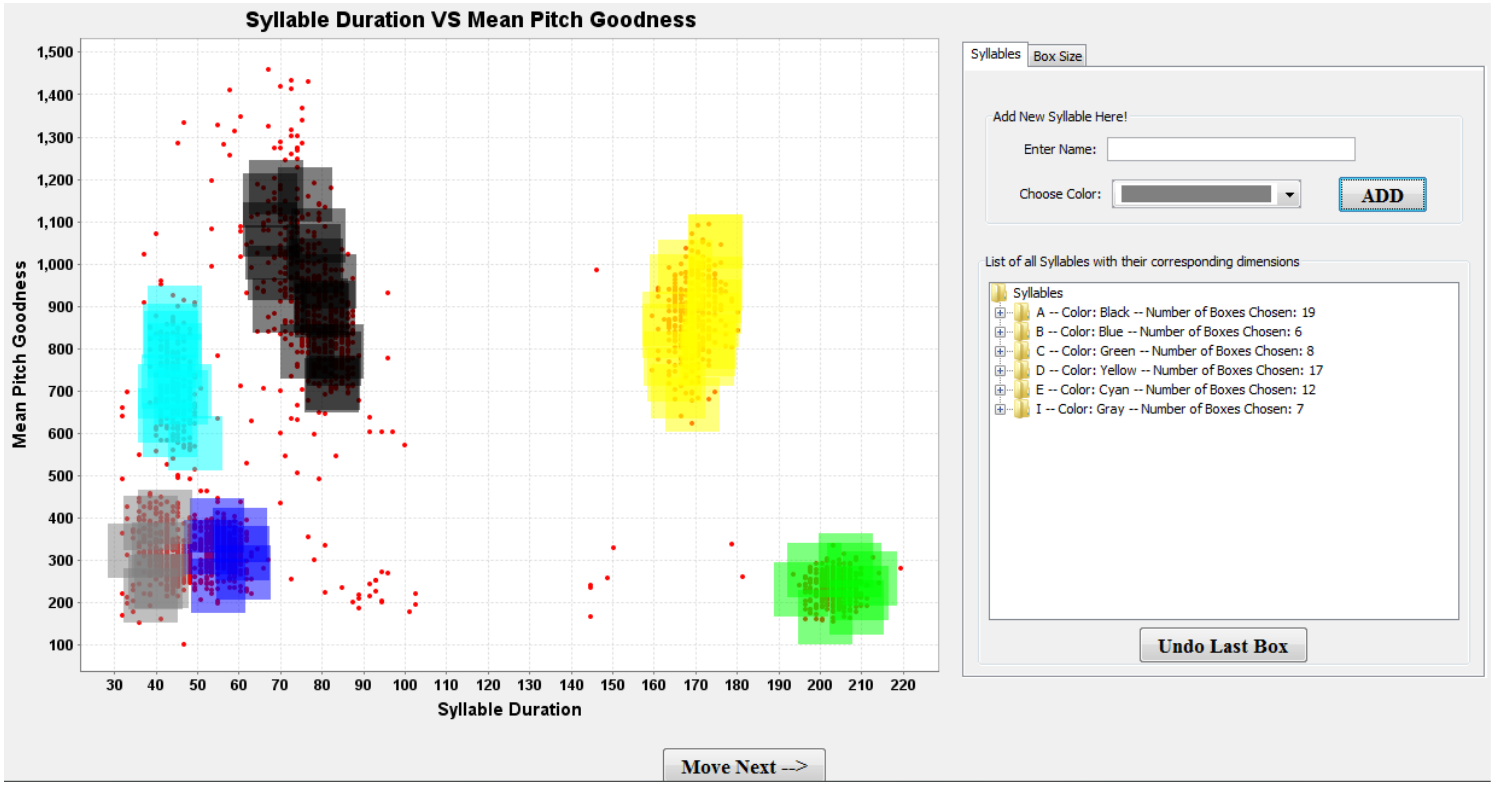

Fig. 42

The next window (**Fig. 43**) asks the user to choose a second pair of features. According to this choice, a new 2D scatter plot appears. The data points are now color coded according to the first step of identification (**Fig. 42**). This is illustrated in **Fig. 44** where the same first pair of features is chosen.

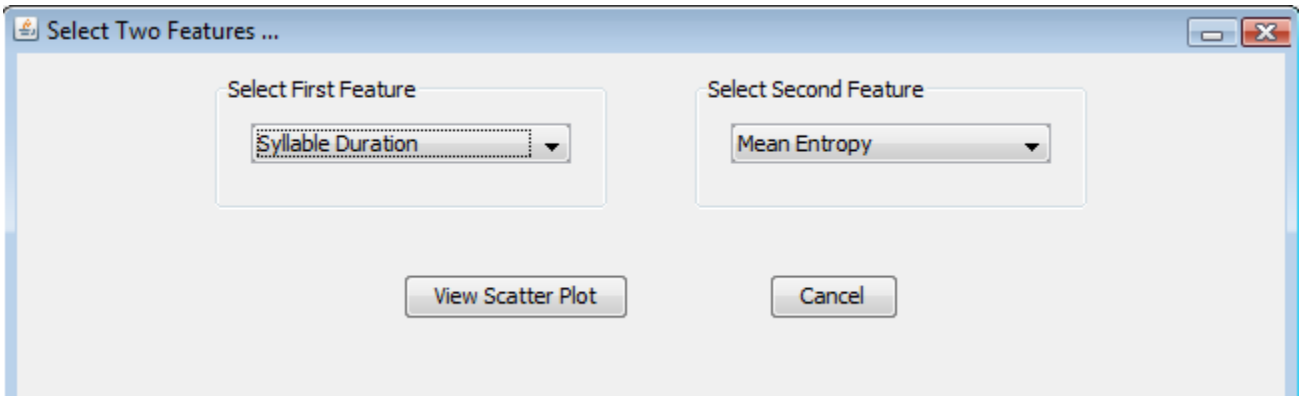

Fig. 43

Any data points that were within the black boxes defining the borders of syllable A in the scatter plot of **Fig. 42**, will be painted in the same color (black) in the scatter plot of **Fig. 44**.

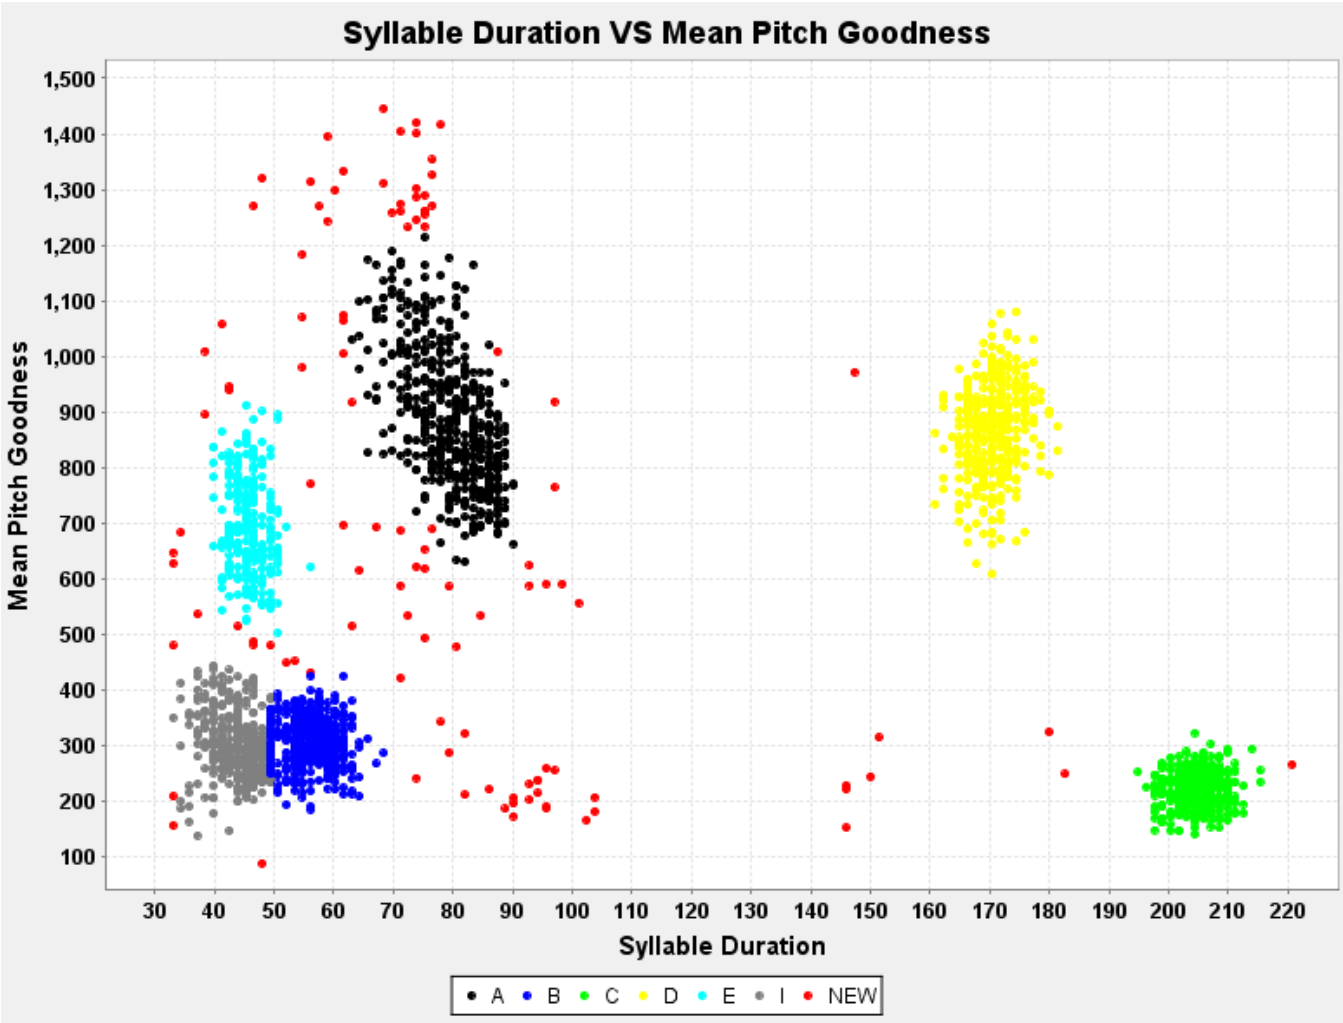

Fig. 44

If the choice of the second pair of features was Syllable Duration versus Mean Entropy, then the scatter plot would like that in **Fig. 45**. In this figure we see that the data points for B (blue) and I (gray) are pulled apart in the new 2D feature space, and it is clear that some data points previously misidentified as B actually cluster better with I (blue points in the bottom right of the gray cluster). Syllables A (black), C (green), and D (yellow) remain well isolated, but syllable E (cyan) has a similar mean entropy as the introductory notes (I) and therefore their corresponding clusters overlap. So, using this module we now have an idea that we were about to misidentify two clusters (B and I).

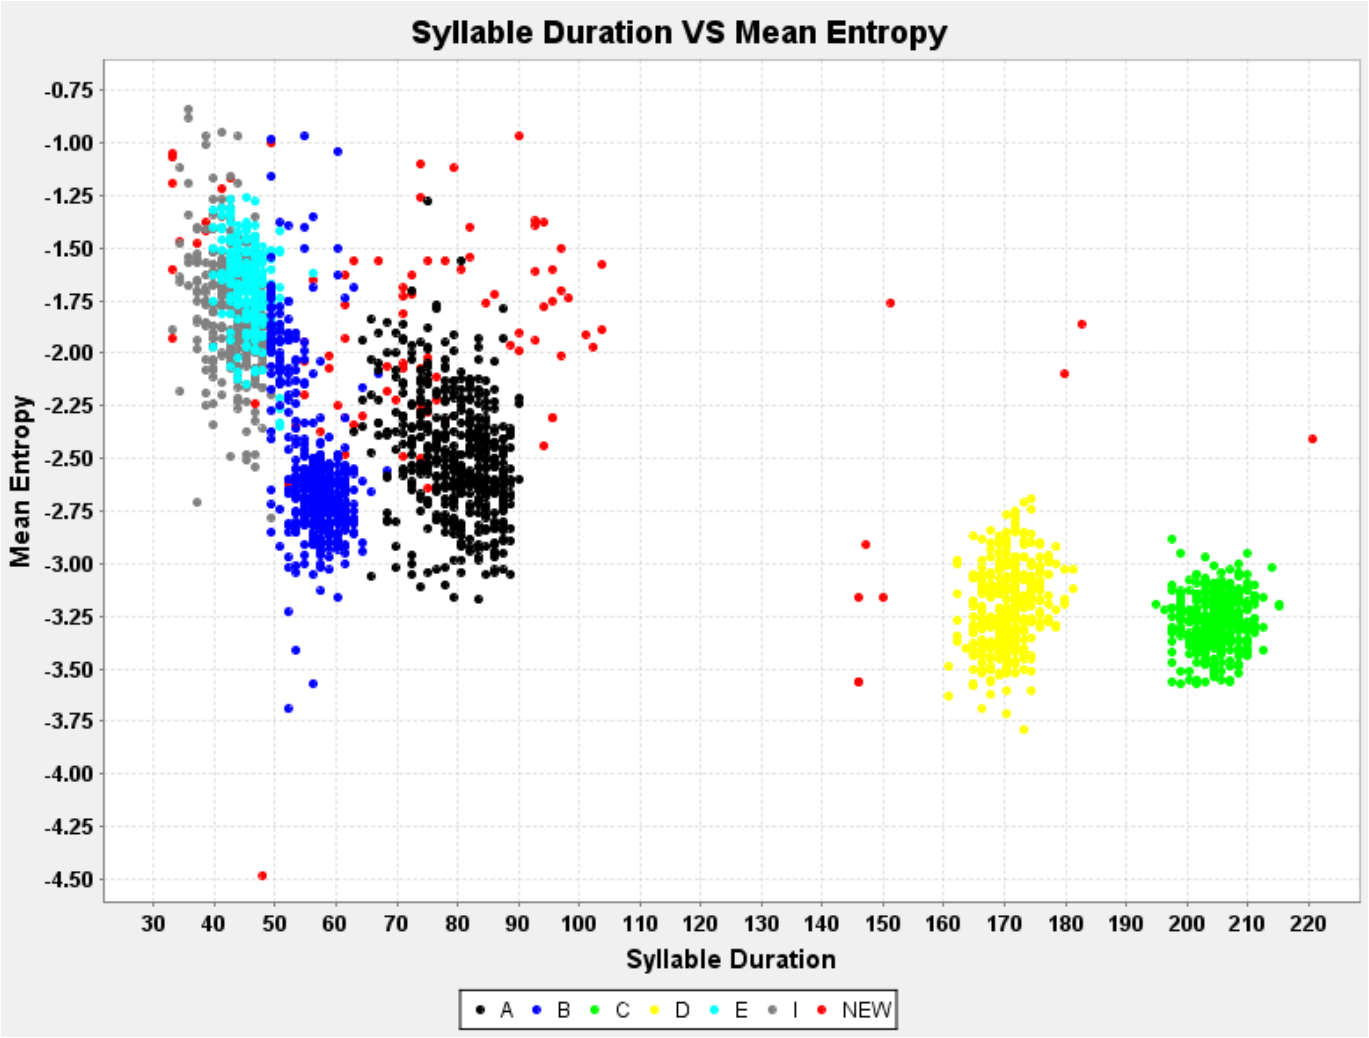

Fig. 45

SongSeq leaves the window of **Fig. 43** open allowing the user to choose as many pair of features as needed. **Figs 46** and **47** show another two examples with different combinations of features for the same example and painting done in **Fig. 42**.

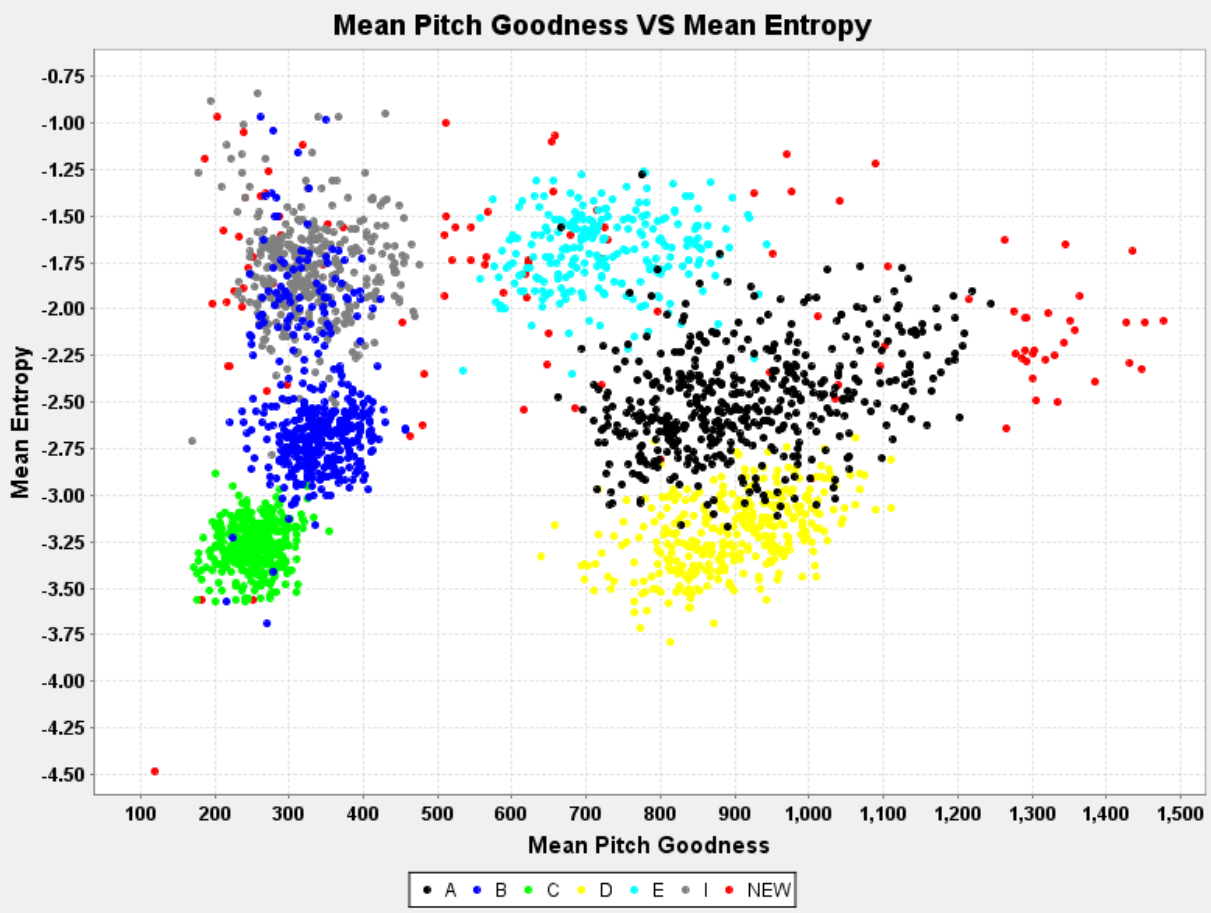
Mean Pitch Goodness VS Mean Entropy
Mean Entropy
-0.75
-1.00
-1.25
-1.50
-1.75
-2.00
-2.25
-2.50
-2.75
-3.00
-3.25
-3.50
-3.75
-4.00
-4.25
-4.50
100
200
300
400
500
600
700
800
900
1,000
1,100
1,200
1,300
1,400
1,500
Mean Pitch Goodness
A B C D E I NEW

Fig. 46

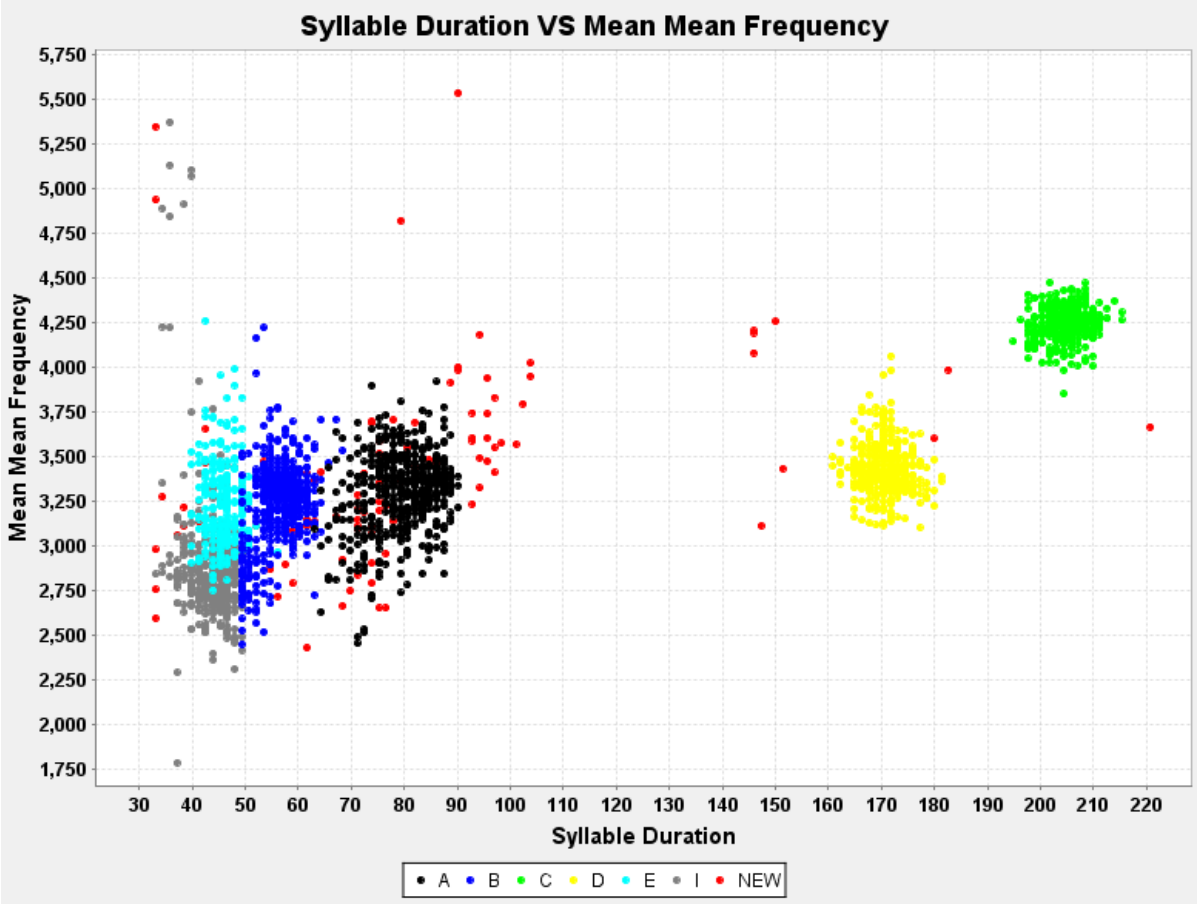
Syllable Duration VS Mean Mean Frequency
Mean Mean Frequency
5,750
5,500
5,250
5,000
4,750
4,500
4,250
4,000
3,750
3,500
3,250
3,000
2,750
2,500
2,250
2,000
1,750
30
40
50
60
70
80
90
100
110
120
130
140
150
160
170
180
190
200
210
220
Syllable Duration
A B C D E I NEW

Fig. 47

Delete Clusters Test

"Delete Clusters Test" (**Fig. 48**) is a module where the user can eliminates one or more clusters after designation. This helps clear the spectrogram and give a better visualization of the data points on the scatter plot.

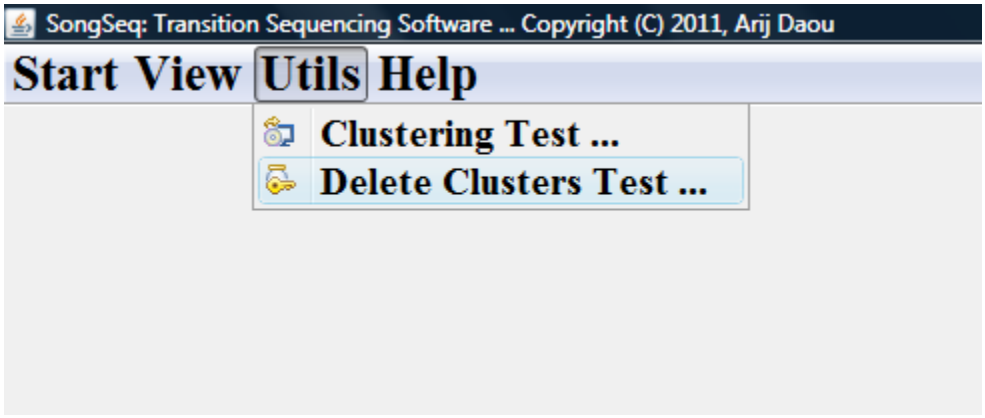

Fig. 48

This module opens initially windows identical to **Figs. 40**, **41**, and **42**. For this reason, we're omitting showing them here. Next, a window showing the color coded syllables appears along with the syllable's names on right hand side with checkboxes near them (**Fig. 49**).

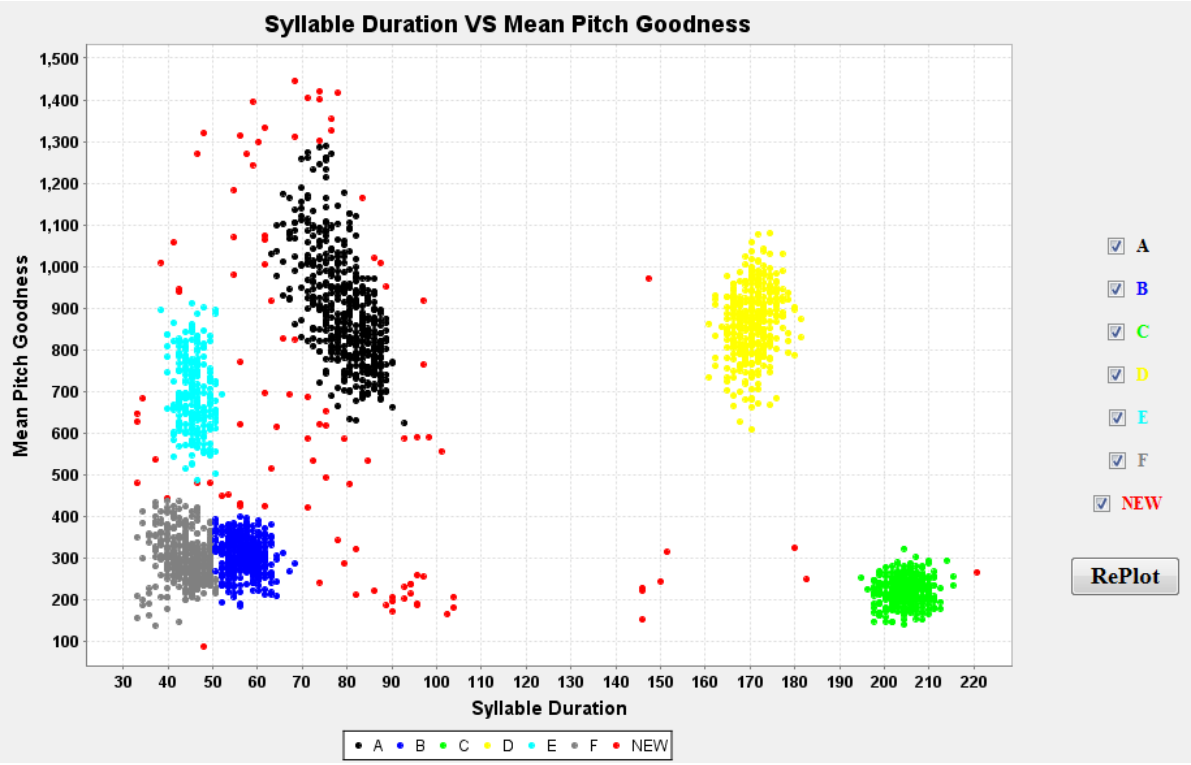

Fig. 49

A checked syllable means that the syllable is painted onto the scatter plot. To remove certain syllables from the scatter plot, uncheck the syllable name and click on “RePlot”. **Fig. 50** shows the scatter plot with syllables A, C, and E eliminated.

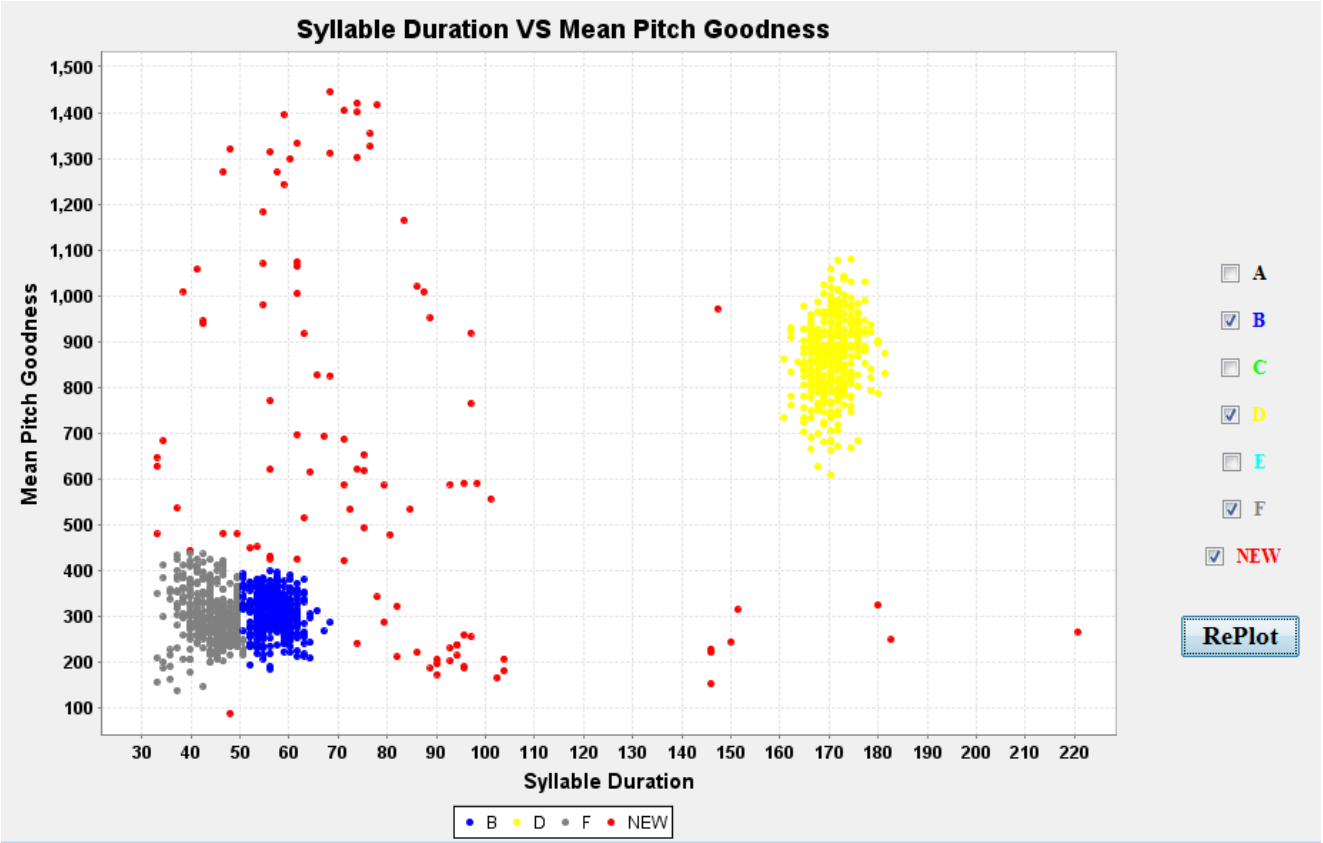

Fig. 50