Welcome to VOIS3D!
Beginning your VOIS3D Session

General:

1. You use VOIS3D by completing 4 "steps:"
          1) Loading data (page 1)
          2) Normalization (Spectral, page 2; Duration, page 4)
          3) Defining Grouping Constraints (page 6)
          4) Overlap and Plot (page 8)
2. Each step has a box on the screen.
3. VOIS3D guides you through the process of calculating and graphing your overlap data by "walking you through" each step. When you first begin a VOIS3D session, you will see that the "Step 1" dialogue box is pink, and other boxes are blue. VOIS3D will only activate steps you can take given what it knows about your data; step boxes will be activated for you in succession when the requirements pertaining to a prior step have been satisfied. Active steps have pink boxes; inactive ones have blue boxes. VOIS3D will not allow you to activate a step whose specifications depend on a preceding step that has not yet been completed.

Saving your sessions:

You may save a session by using the "File" pulldown menu. Select "Save session." You will be directed to name a session file, which will automatically received the filename extension .vois. VOIS3D saves all of your data from your working session (e.g., any normalized and raw values, grouping information, logfiles, and graphs). If you later open your saved .vois file (pulldown menu "File", "opening an existing VOIS3D session"), you can click on any step's box to reactivate it and change the specifications it contains.

The Log:

VOIS3D keeps a record (or log) of all of its procedures. You may toggle between "show log" / "hide log" (bottom left of screen) to view or close the record of VOIS3D's actions regarding your data, including timestamps (included for your convenience in keeping records of each session). The log is saved every time you save a VOIS3D session. New actions will be concatenated to existing logfiles, so that you may keep a cumulative, running record of each VOIS3D project. You can copy the contents of the log for use in other applications (e.g., they may be pasted into a Word document). You may not, however, type into the log window.

Loading Data

(Step 1)

Data should be stored in a tab delimited text file. Practically speaking, it is best to avoid special characters in naming files, however, VOIS3D will allow names with words separated by spaces ("my data.txt"), underscores ("my_data.txt"), or using mixed cases ("MyData.txt"). The first row (i.e. line) of your datafile must contain the header names (also known as "column headers") that you wish VOIS3D to use to refer to the data and factor groupings.

Header names:

Header names cannot begin with a numeric character, but rather should begin with an alpha character(a-z or A-Z). For example, the following are legitimate header names:

- "male"
- "F1_formal"
- "F1"
- "duration"

The VOIS3D program will not allow you to load a data file with illegal header names. For example, if you have a header named "2001Session", or a column header called "1M," VOIS3D will give you a warning message and then fail to load the specified file. In

such instances, it is necessary for the user to manually open the data file (using any text editor) and rename the appropriate header(s).

In each field of the Step 1 action area, use the pulldown menus to select the header in your datafile that contains the data VOIS3D requests: Subject (column containing identification codes or names of your speakers), Vowel (the code by which you distinguish vowel qualities, e.g., "vowel" or "V" are common header names), Spectral 1 (values for the first formant, F1), Spectral 2 (values for the second formant, F2), Temporal segment (segment duration values), Temporal phrase (phrase duration values). **VOIS3D requires that you have at least measures for F1 and F2. Subject* and Vowel fields are also required. Temporal segment and Temporal phrase are optional.** If your dataset does not contain duration values, you must tell VOIS3D so by selecting "not available" in the pulldown menu. **You must have at least one type of duration measure to do three-dimensional overlap calculations or graphs.** If your data does not have Subject or Vowel information, you must return to your text editor and modify your dataset to have dummy variables for these fields. In this way, every parameter must be accounted for.

Note: If you do not have "Subjects" but do have other speaker groupings that you wish to compare, you may use these as your "Subject" groupings. For example, if you have measured one speaker using separate guises (for example, different languages or registers), and your dataset has a column header ("Guise") with multiple levels for this factor listed for each token in the column below (e.g. "Guise1," "Guise2,") you may treat this speaker as if s/he were two speakers. Select "Guise" under the "Subject" pulldown. As a second example, sometimes you might want to compare productions within a speaker across different sessions (for example, before and after training on a set of stimuli), and your dataset has a column header ("Session") with multiple levels for this factor listed for each token in the column below (e.g. "Session1," "Session2,") you may treat this speaker as if s/he were two speakers. Select "Session" under the "Subject" pulldown.

Note, header names are case sensitive which means that the program will treat "Duration" and "duration" as two separate fields. You can use an underscore to have a header with more than one word in it. Any spaces will automatically be converted to underscore by VOIS. Avoid having duplicate header names. VOIS3D will not recognize the second instance and data will not load due to the perception of mismatched number of data columns compared to number of head fields.

Once all information for this step has been provided, click "continue" to proceed to the next step. Note that sometimes, if VOIS3D needs to perform calculations, you may need to wait a while for the next step to be activated. Please be patient: Do not try to enter data in the next action area until the box for that step has turned pink.

Spectral Normalization

(step 2)

In this step, you will normalize your data, if desired. If no normalization is necessary, please select "data already normalized" in all pulldown menus, and then click "continue." Skip to step 3.

If normalization is necessary, you may perform normalization without leaving the program. VOIS3D contains algorithms for normalization of spectral and temporal data. The techniques for formant (i.e., spectral) normalization are:

A. Z-score spectral normalization
B. Logarithmic normalization
C. "Known extremes" normalization

A detailed technical explanation of each algorithm may be found in:
Wassink, A. B. (in press) An Analytic Geometric Method for Quantifying Vowel Overlap, *Journal of the Acoustical Society of America*. You may obtain a PDF version of the paper from the author by emailing wassink@u.washington.edu

The following paragraphs outline each normalization technique. First, a description is given of the type of study for which each might be used. Next comes a description of the handling of your data. Please pay particular attention to the data structures required by each of these algorithms. Not all normalization techniques are appropriate for all types of data. VOIS3D will NOT check that the requirements for appropriate data structure are met before applying an algorithm to your numerical data. You will find a discussion of the specific data structures required by each technique below in the sections entitled, "data structure requirements."

Of course, it is possible to complete a VOIS3D session without normalizing your data. This is desired, for example, when your data have already been normalized before being input to VOIS3D. Simply select "already normalized" in the pull-down menus. Please note that overlap calculations are not possible with raw data. It is, however, possible to generate a graphical plot of raw data.

A. Z-score Spectral Normalization

Used for:
Within-subject or between-subject normalization

Data structure requirements:
-Data must be additionally coded for vowel category (=vowel target) and subject

-Calculation of Grand Mean requires that all vowel qualities must be known and represented (note: a vowel category mean is calculated, making it unnecessary for each vowel category to be represented in equal numbers of tokens). This method is not appropriate if the dataset is missing values for some of the language's vowel categories.

Method:
For each subject and then for each token, the algorithm is applied as follows (note: F1 is used in this example, but the procedure is the same for F2):

1. Calculate a vowel category mean as follows: For each vowel category, calculate the average over all values for F1. For example, if John is the current subject and we have 5 tokens for the vowel category /i/ with values {290, 290, 295, 300, 277}, then John would have an F1 /i/ AVERAGE of 290Hz. In addition to /i/, we have data for John's /a/ and /u/. Calculate vowel category means for each vowel in John's dataset in the manner described for /i/.

2. Using all available vowel category F1 means, calculate a single F1 GRAND MEAN for the subject. The F1 GRAND MEAN is the result of summing all vowel category means and dividing by the number of categories. For example, if John has F1 means of 290Hz for /i/, 452Hz for /o/, and 320Hz for /u/, the result is an F1 GRAND MEAN of 354Hz ((290+452+320) / 3 = 1062/3 = 354).

3. Find a single F1 STANDARD DEVIATION for the subject. Given all vowel category F1 AVERAGES, calculate the standard deviation between category means.

4. Transform each original F1 value using the following formula: (F1 original value - F1 GRAND MEAN) / F1 STANDARD DEVIATION. This yields a normalized F1 value for each token.

5. Repeat steps 1-4 for F2

Research Scenarios:
This method is moderately restrictive in terms of data structure requirements, but it allows for a large number of possible numerical comparisons.

Scenario 1: You may desire to determine the extent of spectral separation or overlap between two adjacent vowels in one language. You have data for several speakers of this language. You have also measured data for each vowel in the system, sufficient to provide you with a means of mapping out the entire system. However, your speakers are of different genders or ages.

Scenario 2: You may desire to investigate possible register effects on the realization of vowels in one language. You have paired data for several speakers of this language (i.e., for each speaker, you have collected comparable data in both register 1 and register 2). You have also measured data for each vowel in the system, sufficient to provide you with a means of mapping out the entire system. You may have VOIS3D caluclate the extent of overlap between V1 in register 1 and V1 in register 2.

B. Logarithmic Spectral Normalization

Used for:
Within-subject or between-subject normalization

Data structure requirements:
-Data must be additionally coded for vowel category (=vowel target) and subject
-Calculation of Grand Mean requires that all vowel qualities must be known and represented (note: a vowel category mean is calculated, making it unecessary for each vowel category to be represented in equal numbers of tokens). This method is not appropriate if the dataset is missing values for some of the language's vowel categories (i.e., the data enable an unbiased estimation of the center of the system).

Method:
This is an implementation of the method pioneered by (Nearey, 1977). This method is appropriate when vowels are evenly distributed within the perimeters of the overall vowel space.

Using an adaptation of the z-score normalization method described above, VOIS3D follows the following procedure to generate normalized values for each token for each subject (note: F1 is used in this example, but the procedure is the same for F2):

1. For all F1 values, convert values to log(base 10).

2. For each vowel category, calculate the MEAN of the log values. For example, Spanish has 5 vowels. Given a Spanish-speaking subject named Patricia, we would calculate an F1 AVERAGE for Patricia's system across the averages of the log values for her /i/, /e/, /a/, /o/, and /u/.

3. Based on the vowel category AVERAGES calculated in 2 above, calculate a single F1 GRAND MEAN by summing each vowel category AVERAGE and dividing by the number of vowel categories.

4. Transform the log of each F1 value using the following equation: (Log of each F1 value - F1 GRAND MEAN).

5. Repeat steps 1-4 for F2

Research Scenarios:

This method provides an auditory-scaling for numerical data. This method is the most restrictive in terms of data structure requirements, but it allows for the greatest number of possible numerical comparisons. You might wish to use this method when your set of comparisons will include pairs of vowels in both the front and back subsystems of overall vowel space.

C. Known Extremes Spectral Normalization

Used for:

Within-subject normalization

Data structure requirements:

-Data must be additionally coded for vowel category (=vowel target)
-Calculation of Grand Mean requires data for vowels that describe the minima and maxima for each spectral dimension. This method is appropriate for asymmetrical vowel systems, or if the dataset is missing values for some of the speaker's vowel categories.

Method:

This method is described fully in Shirai (2004). For each subject, the algorithm performs the following operations:

1. Calculate the log(base 10) for the observed F1 value of each token.

2. (Check on how VOIS3D does this) Calculate the F1 average for all available vowel categories. Locate the vowel category that contains the maximum value in this spectral dimension (e.g., F1). This is the MAXIMUM F1.

3. Locate the vowel category that contains the minimum value in this this spectral dimension (e.g., F1). This is the MINIMUM F1.

4. Values calculated in steps 2 and 3 are taken to represent the EXTREMA of the subject's vowel space. Using these two extremes (one minimum and one maximum for each spectral dimension, here F1=2 extremes), calculate a F1 GRAND MEAN.

5. Transform each token using the formula: ((LOGToken(norm) - corresponding GRAND MEAN)).

6. Repeat steps 1-5 for F2

Research Scenarios:

You desire to calculate overlap between vowels in different morphological categories (e.g., such as those in Japanese). You have data for several speakers, but will confine your comparisons to within each speaker. Within each speaker, you have paired data. For example, each vowel category is represented in a content morpheme and a function morpheme. You desire to see whether function morphemes are spectrally reduced (i.e., more central in the overall acoustic vowel space) relative to content morphemes. In your language, vowels are distributed unevenly in the overall acoustic space. For example, an F2 GRAND MEAN calculated across all vowel categories would be skewed towards /i/, if the language has, say 3 front vowels (/i/ being the most extreme, point-vowel), but only 1 back vowel, /u/. In this method, you will see a dialogue box asking you which two vowels define the edges of the height and front-back dimensions of your acoustic vowel space, and then procedures will be automatically followed to calculate GRAND MEANS on the basis of this information.

Duration Normalization

(step 2, continued)

VOIS3D contains algorithms for normalization of spectral and temporal data. The techniques for duration (i.e., temporal) normalization are:

A. Z-score duration normalization
B. Phrase duration normalization

A detailed technical explanation of each algorithm may be found in:
Wassink, A. B. (in press) An Analytic Geometric Method for Quantifying Vowel Overlap, *Journal of the Acoustical Society of America.* You may obtain a PDF version of the paper from the author by emailing wassink@u.washington.edu

The following paragraphs outline each normalization technique. First, a description is given of the type of study for which each might be used. Next comes a description of the handling of your data. Please pay particular attention to the data structures required by each of these algorithms. Not all normalization techniques are appropriate for all types of data. VOIS3D will NOT check that the requirements for appropriate data structure are met before applying an algorithm to your numerical data. You will find a discussion of the specific data structures required by each technique below in the sections entitled, "data structure requirements."

Of course, it is possible to complete a VOIS3D session without normalizing your data. This is desired, for example, when your data have already been normalized before being input to VOIS3D. Simply select "already normalized" in the pull-down menus. Please note that overlap calculations are not possible with raw data. It is, however, possible to generate a graphical plot of raw data.

A. Z-score Duration Normalization

Used for:
Within-speaker normalization of segmental duration

Data structure requirements:
-Data must be additionally coded for vowel category (=vowel target) and subject
-Data must be coded for segment durations (either ms or seconds is fine, so long as the same units are used for all speakers)
-Calculation of Grand Mean requires that all vowel qualities must be known and represented (note: a vowel category mean is calculated, making it unnecessary for each vowel category to be represented in equal numbers of tokens). This method is not appropriate if the duration vowels are word durations or phrase durations.

Method:
1. Calculate the duration mean for each vowel category.

2. Calculate a single duration GRAND MEAN by summing each vowel category duration mean and dividing by the number of vowel categories.

3. Calculate a single STANDARD DEVIATION based on the duration mean of each vowel category.

4. Transform each original duration value with the formula: (Original token duration - GRAND MEAN duration) / STANDARD DEVIATION

Research Scenarios:
Scenario 1: Researcher desires to compare the durations of phonologically long vs. phonologically short vowels within a speaker. Vowel data are collected in similar consonantal contexts.

Scenario 2: Researcher desires to compare the durations of vowels produced in different consonantal contexts, speech registers. Context is systematically varied (and data are coded for context) so that a similar set of vowels is collected under each context or register.

B. Phrase duration normalization

Used for:
Within-speaker normalization of larger durational units than the segment (i.e., intonational phrase or sentence). Data are initially normalized within each subject, and then by-subject data are pooled in VOIS3D step 3 for subsequent analyses.

Data structure requirements:
-Data must be additionally coded for vowel category (=vowel target) and subject
-Data must be coded for "phrase" durations (either ms or seconds is fine, so long as the same units are used for all speakers)
-In addition to "phrase" durations, data must be coded for token durations. Thus, datafile must contain *two* columns of duration information. The user will be asked to select the column that contains each type of duration information.

Method:
1. Find the duration mean for each vowel category.

2. Transform each original duration value using the formula: (original token duration / token phrase duration) x duration mean of the corresponding vowel category.

Research Scenarios:
Scenario 1: Researcher has collected data under several different conditions known to affect duration (but has determined that pooling data is an appropriate step). For example, data were collected from a speaker who varied his or her speech rate across the course of a reading list elicitation session.

Scenario 2: Researcher desires to normalize temporal information for between-speaker comparisons, where speakers vary in their speaking rates, although data were collected under the same conditions of formality. The object of the investigation is the relationship between vowel duration and the length of the overall syllable. This may be useful, for example, in examining duration reduction in stressed vs. unstressed syllables.

Once all information for this step has been provided, click "continue" to proceed to the next step. Note that sometimes, if VOIS3D needs to perform calculations, you may need to wait a while for the next step to be activated.

Defining Grouping Constraints (add and delete)

(Step 3)

Groupings and Types

The purpose of Step 3 is to define your desired vowel comparisons, based on whatever constraints you so choose. After your data have been loaded (Step 1) and you have specified the headers that correspond to your dependent variables, VOIS3D allows you to view the independent variable/factor groupings (e.g., subject code, gender, vowel targets, etc.) that were detected in your data file. The 'Categories' listbox contains all of the eligible factor group headers that remain after the principal factors of Subject and Vowel were isolated in Step 1. Clicking on any of these categories will result in that category's types appearing in the "category types" listbox, below. You may select from among these 'Category Types' to define your grouping parameters, as follows:

1. Tell VOIS3D how many groupings you wish to compare. By default, VOIS3D defines two groupings (which is the minimum required). These default groupings appear in the "Select a grouping to modify" listbox. If you wish to have more than two groupings, create more groupings by clicking "Create new grouping" button. Alternatively, if you wish to delete a grouping, click the "Delete grouping" button. Each grouping is automatically numbered in the listbox, for your convenience in determining how many groupings you have defined.

2. Use "Change grouping name" to name each grouping. Each grouping will appear with the default text "(change my name)" until defined. Select the grouping you wish to name (its name will be highlighted), and type a new name in the "Change grouping name" dialog box. Press the Enter/Return key on your keyboard. The name will be listed in the "Select a grouping to modify" listbox. You may change a name as many times as you like before leaving this step.

3. Specify or Modify grouping constraints. Now you can define your groupings. Select one of the groupings you have just defined in step 2, above. Add category type values that define this grouping to the "Modify grouping constraints" listbox. To do this, you first select a "Category." The "Category types" listing will change to reflect the levels of your "Category" variable. Second, select one category type, and press the "Add" button. The category type name will appear in the "Modify grouping constraints" listbox. If you wish to add all levels of the Category to the grouping, click on "All>>." All category type names will appear in "Modify grouping constraints." Click on the second Category that defines an attribute for your grouping. Select one or more category types, as before. Continue until the grouping is specified as desired. You may change specifications as many times as you like before leaving this step. To remove a category type value, click "Delete." The type name will disappear from the "Modify grouping constraints" listbox.

4. Define any remaining groupings. Select the next Grouping to be defined from the "Select a grouping to modify" listbox. Define this new grouping following the procedures listed in 3, above. Repeat for each grouping until all have been specified.

For example…

Let's say you wish to conduct a within-speaker comparison of overlap between two vowels. Your factor groupings are listed under "Categories" as:

Your categories	Category types detected for this category:
Speaker	Speaker 1, Speaker 2, Speaker 3
Vowel	i, ih, a, u, uh
Language	English, Toda, French
Session	Casual, Wordlist, Minimal Pair, Reading passage

You wish to compare Speaker 1's productions of (i) and (ih), collapsing data for the Wordlist and Minimal Pair sessions.

(a) Name the two default groupings. It is not necessary to define additional groupings, since you wish to make only one two-way comparison. From the "Select a grouping to modify" listbox, select "1 (change my name)." Type "Speaker 1, vowel (i), Wordlist and Minimal Pair sessions" into the "Change grouping name" dialogue box. Press return/enter. The new grouping name will replace "(change my name)" above. Select "2 (change my name)." Type "Speaker 1, vowel (ih), Wordlist and Minimal Pair sessions" into the "Change grouping name" dialogue box. Press return/enter. Both groupings are now defined. (Unfortunately, IPA symbols cannot be used at this time).

(b) Select one of the groupings you have just created, e.g., "Speaker 1, vowel (ih), Wordlist and Minimal Pair sessions." Choose the levels (or types) of all category variables that define this grouping. Click on "Speaker" in the "Category" window. "Speaker 1, Speaker 2, Speaker 3" will appear in the "Category types" listbox. Select "Speaker 1" and click "Add." The "Modify grouping constraints" window will be updated to add this category type. Next, select the category "Vowel". Select "ih" from the list of types, and click "Add." Next, select the category "Language." If Speaker 1 is uniquely specified for one language, no selection is necessary for the "Language" category. However, if you have a "Speaker 1" for each language in your dataset, and the speaker you desire is the Speaker 1 from the Toda dataset, you must select the "Language" category, then select category type "Toda" and click "Add." Finally, you wish to combine tokens of (ih) across the wordlist and minimal pair conditions. Select the category "Session". Select "minimal pair"; click "Add". Select "Wordlist"; click "Add." Both category types will be listed in "Modify grouping constraints." The ordering of category types does not matter. You have now finished specifiying the grouping, "Speaker 1, vowel (ih), Wordlist and Minimal Pair sessions." Follow these steps to define "Speaker 1, vowel (i), Wordlist and Minimal Pair sessions" in the same fashion, with the exception that when you select the "Vowel" category, you will select "i" as the category type.

You should now have two categories defined as follows (assuming you needed to specify a Language):

Grouping name:	
1 Speaker 1, vowel (ih), Wordlist and Minimal Pair sessions	2 Speaker 1, vowel (ih), Wordlist and Minimal Pair sessions
Modify grouping constraints listbox should say:	
Speaker 1	Speaker 1
ih	i
Toda	Toda
Minimal Pair	Minimal Pair
Wordlist	Wordlist

Note: In order to be a legitimate grouping, the category types in your file must be characters and not numbers. Thus, if you have a header entitled "Session" and the values beneath it's column are 1, 2, 2, 3, etc., it will NOT appear as possible Category that you may select for defining groupings; however, if the labels beneath the hypothetical column "Session" are "A", "B", "B", "C", then you can define vowel groups with the types found in your "Session" column.

VOIS3D's constraints step gives you a great deal of flexibility -- you can define comparisons narrowly (e.g., two groupings comparing two vowels within just one speaker), or broadly (e.g., two groupings comparing all tense (i) vowels produced by everyone in a group of speakers to all lax (ih) vowels for everyone in that same group of speakers). To define a broad grouping variable, for example, you might add more than one speaker to the "Modify grouping constraints" box, if this is appropriate for your data.

Once all information for this step has been provided, click "continue" to proceed to the next step. Note that in this step, VOIS3D must partition your data into new columns, defined over the types you have just specified. You may need to wait a while for the next step to be activated.

Overlap and Plot

(Step 4)

A. Overlap

In this step, which is the application of the Spectral Overlap metric, you choose the desired type of overlap calculation (2-dimensional, or 3-dimensional). From the pulldown menus, select the vowel groupings for which you desire to calculate overlap. You should see a list containing all groupings you defined in Step 3. Note: VOIS3D does not check to ensure that you have not chosen the same grouping twice. The overlap calculation will result in an overlap fraction of "1", indicating 100% overlap. You may return to the pulldown menu as many times as needed to calculate overlaps again.

This step takes time. Please be patient. A "progress" bar will appear on the screen while VOIS3D is working.

The overlap fraction will appear, as a decimal, beneath the appropriate calculation button, once calculations are finished. Values are rounded to 3 decimal places (e.g., .893=89% overlap, .001=0% overlap, etc.). Remember, values are saved in the log, so you can complete as many calculations as you like without worrying that you are losing the values printed to the screen once each calculation is completed.

B. Plot

Displaying your vowel groupings and vowel overlap
In this step, VOIS3D displays your data. VOIS3D can display:
1. Raw data (your original data in scatterplot format)
2. Normalized data (in scatterplot format, with units based upon statistical information)
3. "Wireframe" models of the 2-d ellipses or 3-d ellipsoids generated in Step 4A.
4. Filled area graphs (2-d) or volumes (3-d) of the ellipses or ellipsoids generated in Step 4A, including a shaded area representing any overlapping region defined by the two vowel groupings.

Note: you must again tell VOIS3D which vowels to display, using pulldown menus. Check the checkboxes to the left of each menu, to indicate that you wish to include that grouping in your display (VOIS3D can display a single grouping only). The distributions you define as Group A and Group B in this step will correspond to the symbols you choose below when plotting.

Note: you cannot display wireframe models or filled area graphs until you have calculated overlaps for the relevant vowel distributions. Do not try to skip Step 4A and plot data without first completing this step. These plot options require that mean and standard deviation information have been calculated for your data, and this information is calculated in Step 4A.

A separate figure window will appear to display your specified plot. This figure may be saved to a .jpg (or other file format), using the File menu, or via screen capture, if preferred.

Configuring your figure window
You may configure your graph's axes, title, and symbol markers. The bottom window on the screen allows you to specify Axis labels (X,Y for 2-dimensional figures; X, Y, Z for 3-dimensional ones), enter a title for your graph (e.g., "Speaker A, vowels /i/ and /e/"), and choose symbols and colors for representing vowel distributions. The rightmost listbox allows you to select a symbol and color for each group. Groups A and B are specified in Step 4A. You must choose a marker symbol and color for each Grouping with a "checkmark" in Step 4A, or VOIS3D will select one for you. You select a marker for each Grouping from the same pulldown menu (VOIS3D remembers which one is which).

Note: in the full version, three-dimensional displays may be rotated to be viewed from any angle. Rotation is not possible in the standalone version.

For further information, please contact:
wassink@u.washington.edu