ETAnalysis Manual

EYE TRACKING DATA ANALYSIS TOOL

MANUAL VERSION 4.14
February 2025

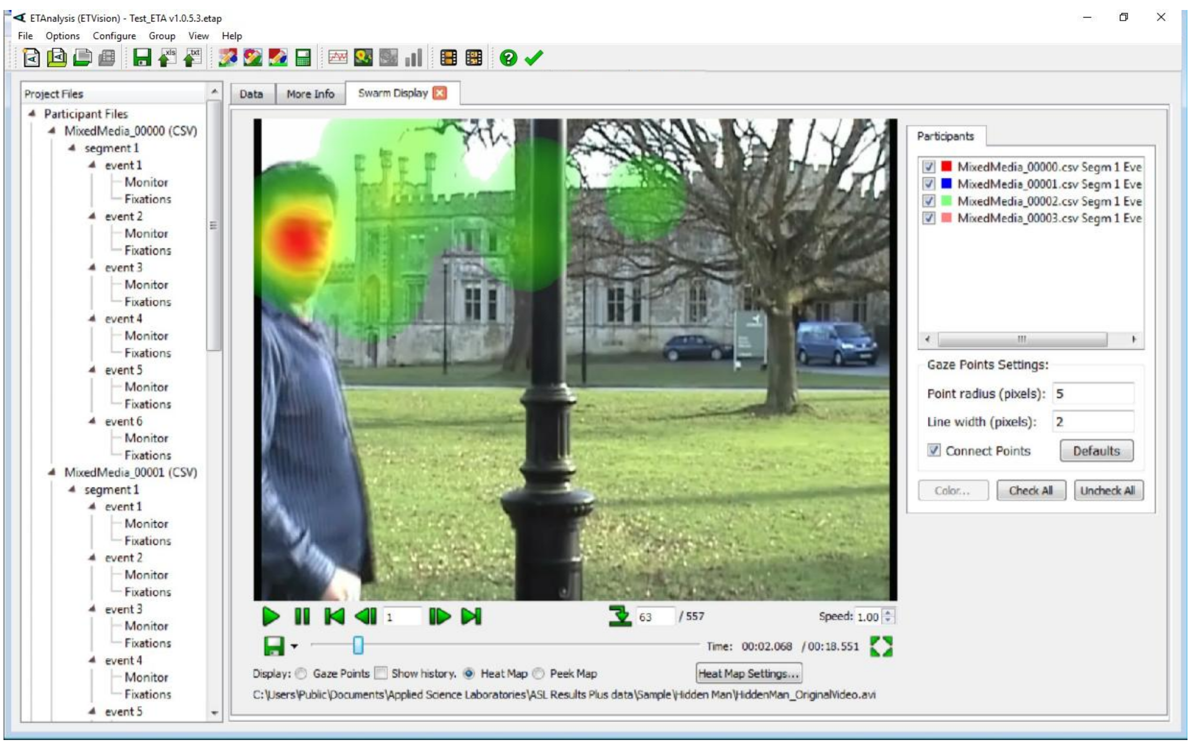

ARGUS SCIENCE

techsupport@argusscience.com
Web site: www.argusscience.com

Table of Contents

1 Introduction

1.1 Basic Features

Argus Science *ETAnalysis* is designed to help process and analyze data collected with eye trackers made by Argus Science. It can be used to:

- examine and plot raw data;
- associate scene images with sections of gaze data;
- define areas of interest on images;
- associate videos with sections of gaze data;
- define moving areas of interest on videos;
- reduce gaze data to fixations;
- reduce gaze data to "dwells" (periods of continuous gaze on one area of interest);
- detect saccades between identified fixations and compute saccade parameters;
- display data graphically;
- time plots;
- X/Y scan plots superimposed on scene image;
- heat map plots on scene image;
- compute various statistics that relate fixations or dwells to areas of interest and produce corresponding bar plots;
- combine results across trials or subjects by averaging statistical data from each, or by pooling the original data;
- create swarm display showing gaze from multiple trials or subjects overlaid on a single background or video;
- export results in Excel or ASCII text format for further custom analyses.

The Argus *ETAnalysis* application is project based. A project includes multiple data files, scene video files (if applicable), stimulus files (backgrounds and/or videos presented to participants) and all of the computations requested by the user. The project is organized by sections of data called "events", defined by start and end conditions specified by the user. A tree diagram, in the left panel of the main program window, shows the project hierarchy, and a context menu available by right clicking each node lists all operations that can be performed on that node and its sub-nodes.

Argus *ETAnalysis* can analyze ".csv", ".eyd", and ".ehd" file data recorded by Argus products.

1.2 Optional Features

1.2.1 Stimulus Tracking

Stimulus Tracking ("StimTrack" or ST) is an optional feature that can greatly enhance analysis capabilities when a head mounted eye tracker has been used to record gaze with respect to a head mounted scene camera (when there is no external head tracker and *ET3Space* cannot be used).

Automated analysis of this type of data is traditionally difficult because objects that are stationary in the environment are moving images on the head mounted scene camera. The digital data specifies point of gaze on the scene camera field of view, but not with respect objects or surfaces in the environment.

Stimulus Tracking allows users of a head-mounted eye tracker with only a head mounted scene camera (no external head tracker) to analyze data of participants looking at a computer monitor as efficiently as if the data came from a stationary (table mounted) eye tracker. Stimulus Tracking can be used to:

- define backgrounds or videos associated with participant trials,
- define screen capture videos recorded with participant data sessions,
- define areas of interest within these stimulus backgrounds or videos to share across multiple subjects or trials,
- automatically track the computer monitor through a participant's scene video,
- analyze gaze within the scene monitor and within areas of interest defined in stimuli,
- create swarm display showing gaze from multiple trials or subjects overlaid on a single background or video that was presented on the computer monitor,
- perform all the options listed in the previous section to view, combine, and export gaze results.

The *Stimulus Tracking* feature is described in detail in section 17 of this manual.

1.2.2 AI Object Training

ETAnalysis can use an artificial intelligence (AI) model to recognize and detect objects that appear in the head mounted scene camera image and to create moving areas of interest (MAOIs) in the video. Argus Science *ETVision* eyetrackers can also use an AI model to perform a similar task in real-time. Argus Science provides an AI model trained to recognize a list of 80 standard objects. An optional add-on to the Argus Science *ETAnalysis* application enables users to create custom AI models that will recognize objects specific to their requirements and environment. This feature is described briefly in section 18 of this manual, and in more detail in a separate *ETVisionAI* manual.

1.2.3 ET3Space (*ETVision* feature)

An optional addition to the Argus Science *ETVision* eyetracker, called *ET3Space*, uses head position and orientation data from one of several third party motion tracking devices to compute gaze with respect to a room fixed coordinate system and to determine point-of-gaze on multiple surfaces in the environment. This optional feature is described in detail in a separate manual. Note that *ET3Space* is an *ETvision* option, not an *ETAnalysis* option. *ETAnalysis* is always equipped to process *ET3Space* data. Some additional information about *ET3Space* is provided in section 20 of this manual, and *ET3Space* is described in detail in a separate manual.

1.3 Hardware Requirements

The basic *ETAnalysis* program will run on any PC running Windows 10 or Windows 11, although many functions will run faster and more efficiently on computers with more powerful processors and more memory. This is especially true when processing very large data files or a large number of data files in an *ETAnalysis* project.

Use of the *AI object detection* feature available in *ETAnalysis* (section 16.7.9) requires a PC with at least 16 GB RAM and equipped with an Nvidia GPU, model 3060 or above. The GPU driver must be version 31.xxx or above (must support Cuda). If a GPU driver supporting Cuda is not detected, the AI object detection feature will not be enabled.

An optional *AI Object Training* feature (section 18) requires an additional license and has the same minimum PC requirements listed in the preceding paragraph. For use of this feature it also recommended that there be at least 200 GB of free space available on the C:\ drive.

2 Project structure

The project structure is represented by a tree diagram on the main program window left panel. "Nodes" are added to the tree as data and analysis results are added to the project. Each node "holds" a section of data or analysis computation result. At the top level of the tree is a node called "Data Files", with sub-nodes that are the original data files recorded by the eye tracker.

When data is recorded by some Argus eye trackers the user can start and pause recording on a single file as many times as desired. Data files are therefore divided into "segments" of continuous data (between each record and pause). In *ETAnalysis*, these data file "segments" form sub-nodes under the data file node. A file may have only one segment or multiple segments. (".csv" type files, used by some eyetracker models, always include only a single segment). Note that the number of segments in a data file is not determined by *ETAnalysis,* but rather was determined as the data was recorded.

ETAnalysis can further sub-divide each data segment into "events" defined by some beginning and end criteria, and these "events" form sub-nodes under the each segment node. Each segment must have at least one event sub-node. The default event is the entire segment, but the user may specify criteria to divide a segment into multiple events. If the data is part of an experiment design, events usually correspond to experiment "trials". The data file, segment, and event nodes all represent sections of originally recorded data.

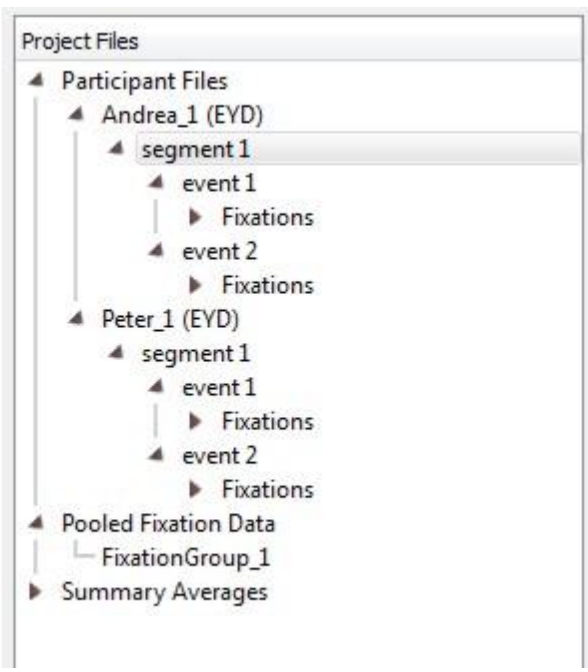

Sub-nodes under each event are all created by data processing in *ETAnalysis* program. Gaze Data in an event can be reduced to a set of fixations, forming a sub-node under the event. Fixation sets can be further processed to match fixations with "areas of interest" on the scene, forming "fixation sequence" and "dwell" nodes under the fixation node. Various statistics can be computed from the fixation sequence and dwell data to form additional sub-nodes.

In the example shown at left, the project contains two files, "Peter_1.eyd", and "Andrea-1.eyd". Only one data segment was recorded on each file, but each segment has been divided, by *ETAnalysis*, into 2 events. Fixations sets as well as Fixation Sequence and Dwell statistics have been computed for all events.

To examine effects across different events (or trials), it is necessary to combine data from some of the nodes that are at the ends of these tree branches. Data gathered from multiple fixation nodes are grouped under a top-level node called "Pooled Fixation Data". Data gathered from groups of statistics nodes (the very ends of the Data File node branches) are grouped under another top-level node called "Summary Averages".

A right arrow symbol (▶) at a node (see "Summary Averages" node, or "Fixations" nodes on the tree diagram example) indicates that there are sub-nodes below it which can be expanded (made visible) by left clicking on the right arrow symbol. After expanding, the arrow symbol on the node will point diagonally towards lower right (◢). This symbol can be clicked to collapse the node.

Right clicking on any node brings up a context menu with a list of operations that can be performed on data in that node. Almost all nodes have a data display in the right panel of the main *ETAnalysis* window, which shows a listing of the data at that node. In each case there is also a “More Info” tab on the right panel, which provides various additional information about the contents of the node. The highest-level nodes (“Participant Files”, “Pooled Fixation Data”, and “Summary Averages”) are the only exceptions. These contain no data, but only serve to define the category of data in their branches.

ETAnalysis can analyze data from csv data files, eyd data files, and ehd data files recorded from Argus Science products. A single project, however, is intended to include files of a single type. These may be csv, eyd, or ehd.

3 Project Management

3.1 Project Size

While there is no set maximum number of events or files, as more events and files are added to a project it will become progressively more difficult to "see" the whole project on the tree diagram, and the program may begin to perform some operations more slowly. To the extent that experiment design allows, it is generally better to divide work into multiple small projects rather than one very large project. Note that different projects can be used to perform different tasks using the same data.

3.2 Saving and Backing up projects

In addition to automatically saving the current project before closing, the program will automatically save the current project at regular intervals. The user can also manually save the project to the same path and name by selecting **File→Save Project**, or to a different path and file name by selecting **File→Save Project As.**

Note that when the program saves automatically, or if **File→Save Project** is selected, it always saves to the same file name, and if this file becomes corrupted it is still possible to lose work.

If doing procedures in *ETAnalysis* that would take a long time or be difficult to recreate, it is strongly suggested that manual backups also be made at regular intervals. Such backups are easily created by using **File→Save Project As**, and using a different name each time (E.g., a sequential number can be added to end of the project name for each save). Once it is verified that recent backups can be loaded successfully, older backups can be safely deleted if desired.

Creating background configurations and area of interest sets are the tasks that most often require investment of significant time and work. These are saved as part of the project file, but can also be exported as independent files which are then available for use in other projects. In cases where significant effort is invested in configuring backgrounds and creating AOIs, it further recommended that the export feature be used periodically to save AOI sets and configured backgrounds (in addition to the periodic project backups previously described). See manual sections on background configuration and AOI creation for instructions on exporting these as saved files.

Be sure that projects are always saved to "public" locations where all users who will need to run the project have operating system "permission" to read and write without administrative privileges.

3.3 Data, Background Image and Video File storage

The first time a project is saved with a particular name, a folder is created, at the specified path location, with the name of the project. This will be referred to as the "project folder". A file with the same name and a ".etap" extension is created in the folder, and is the "project file". The project folder also contains a "privatedata" folder that is normally invisible in Windows Explorer. (It can be seen by selecting the "Show hidden files and folders" in Windows Explorer "Folder Options").

Gaze data files ("Participant" files), image and video files used in a project can be located any where on the computer running *ETAnalysis*. These files are not copied into the project file, but rather the project file stores a pointer to them. Although not a requirement, it is often very beneficial to have a single path location for all such files used in a project. This makes it less likely that a file being used in a project is inadvertently moved or deleted, and also makes it easier to move or copy the project to a different PC.

A recommended practice is to create another folder with a name that refers to the project. Before opening participant data, image, or video files in the project, copy them to this folder and open these copies in the project. If desired, the project folder can be used for this purpose. Be sure to always use locations where all users who will need to run the project have operating system "permission" to read and write without administrative privileges.

Text (csv) data files created by *ETVision* eye trackers always have a companion xml configuration file. The configuration file will always have the same name as the data file, but will have an xml extension. When moving or copying files be sure to move or copy this companion configuration file along with the data file.

3.4 Copying a project to a different PC

To copy a project to a different PC it is necessary to copy the project folder with all its contents, and it is also necessary to copy all participant data files, image files, and video files used by the project. As mentioned above, this is most convenient if they are all in a single known folder, or at least a small number of known locations.

When the copied project is first opened on the new PC, the project may not initially be able to find the ancillary files (participant, image, and video files). This is because the project records the absolute (rather than relative) file locations. If the ancillary files are not in exactly the same path location as on the PC from which the project was copied, the project will not initially find them. The following warning message may appear.

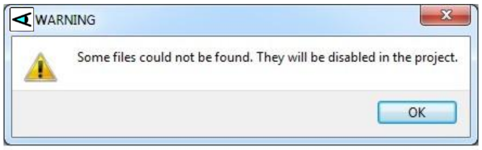

Click OK. A **Required Files** dialog will appear with some files labeled "missing" in the "Status" column.

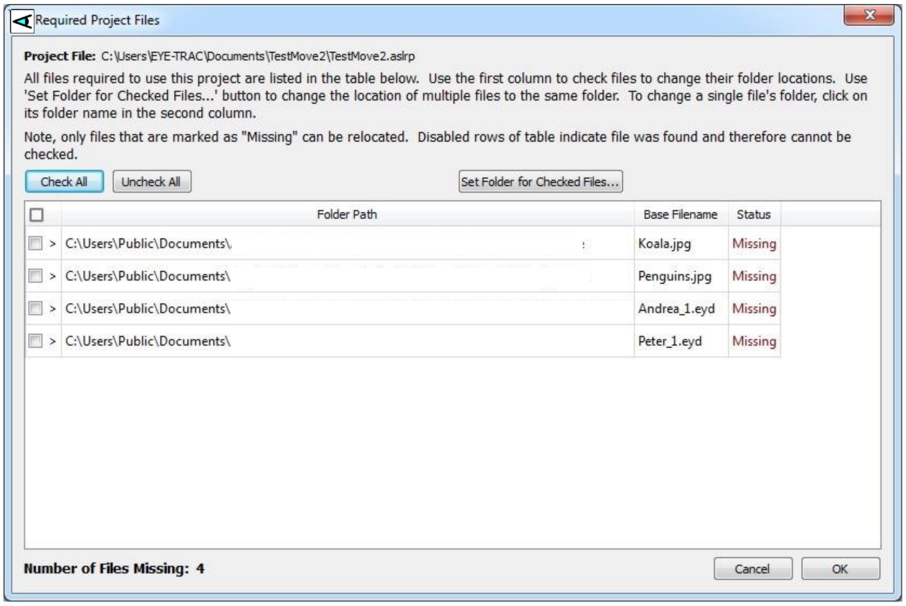

Check mark a group of missing files that are in a common folder (Check all if they are all in the same folder), and click the "Set Folder for Checked Files…" button.

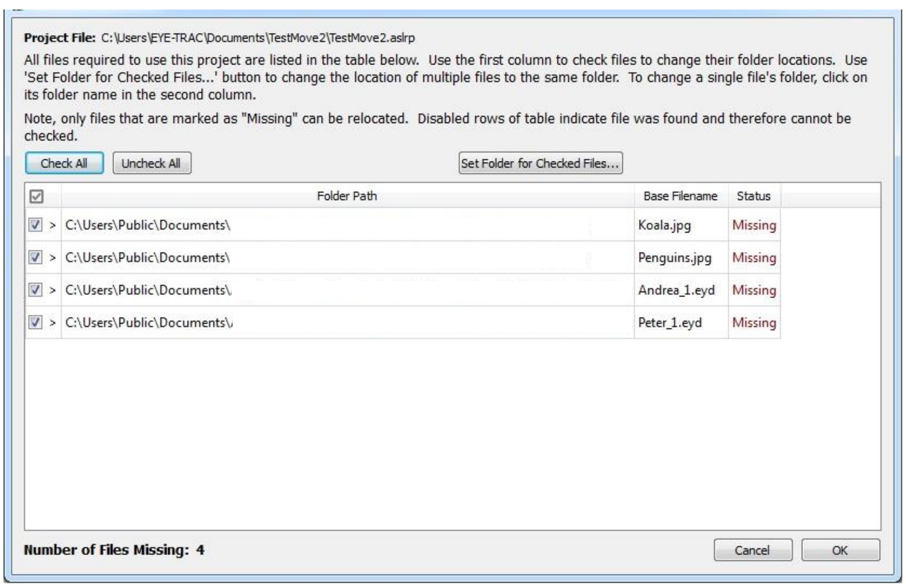

Browse to the containing folder and click "Choose". The "status" of those files should now be labeled "Relocated". Repeat for other files if necessary.

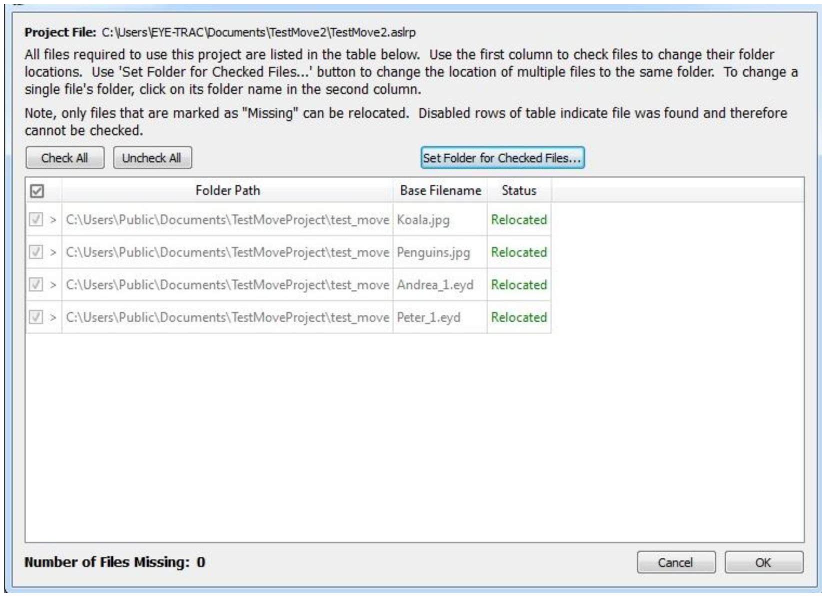

When all files are labeled “Found” or “Relocated”, click OK. The project should be ready for use.

If the initial error message does not appear, select Manage Project Files from the Files menu to bring up the **Required Files** dialog, and make sure that all files are “Found” or “Relocated” before proceeding.

3.5 Copying a project to a different location on the same PC

Copy the project folder with all its contents to the new location. If all of the participant data files, image files, and video files used by the project have not moved, the project should open and operate normally. If the missing files warning appears proceed as described in the previous section. Even if the warning does not appear it is prudent to bring up the **Required Files** dialog, as previously described, and make sure that all files are “Found” or “Relocated” before proceeding.

4 User Interface and Creating a Project

There are four main sections of the user interface as labeled in the following picture: 1) Menu bar, 2) Toolbar buttons, 3) Project Tree, 4) Graphics/Data Display Area. These areas are customized according to the current project type and what features are selected to be made available when starting a project.

4.1 Creating or opening a Project

When the program is started it may automatically open the last project that was saved ("Restore last project on start up" is an option that enabled or disabled under the "Options" menu). To open a new project select **File→New Project** from the menu bar, or click the equivalent shortcut button (Hovering the mouse over a short cut button displays text describing its function). This will cause the current project to be saved and will bring up the "Open or Create Project" dialog.

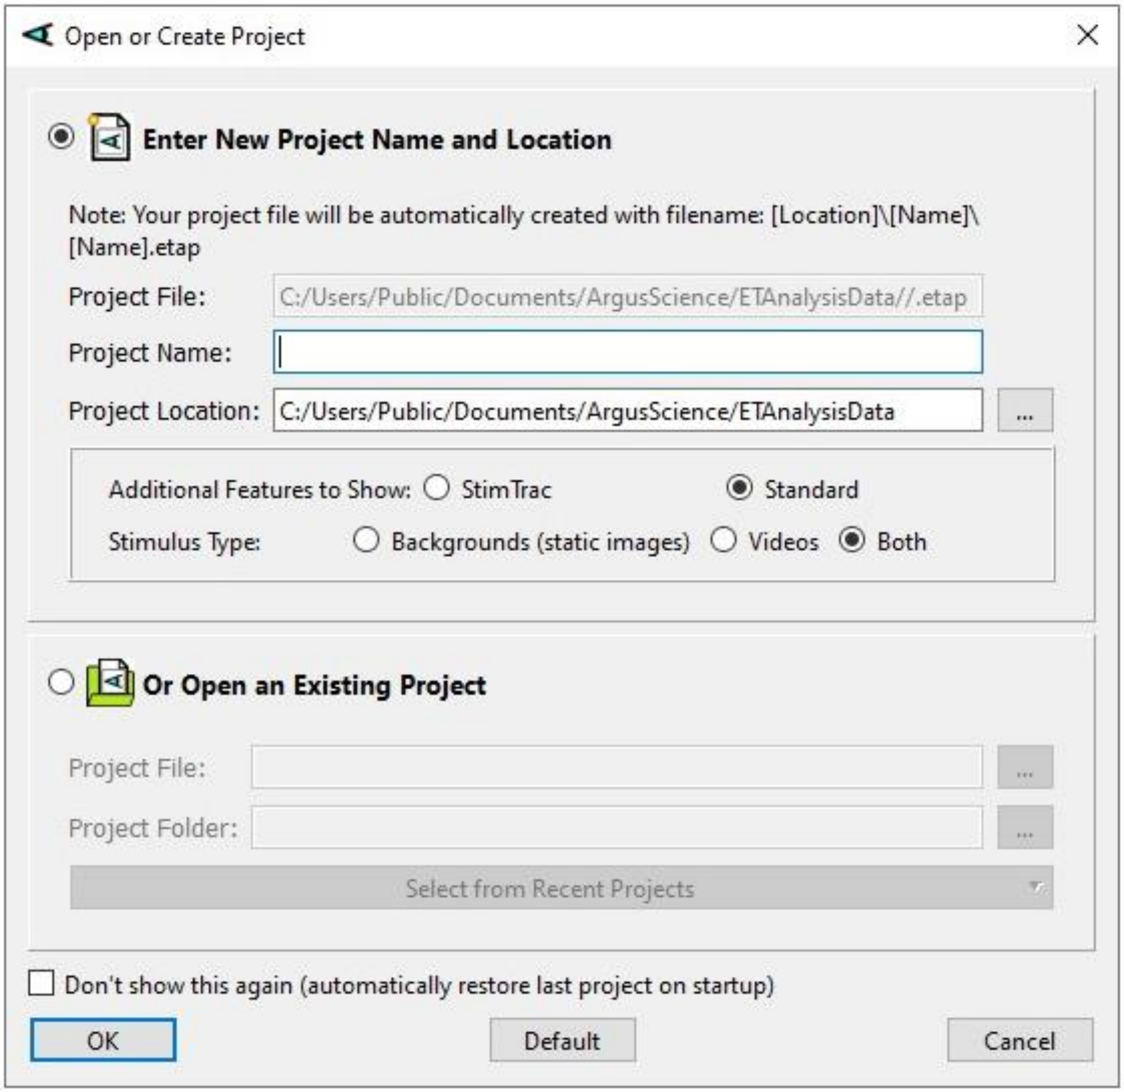

If opening the program for the first time, if no "last project" is detected, or if the "Restore last project on start up" option has not been enabled, the "Open or Create Project" dialog will appear automatically. To open a new project set the radio button to the "Enter New Project Name and Location" group box and follow the instructions below.

Set the "Additional Features" radio button

The "StimTrac" radio button will be grayed out and inactive if the program is not equipped with a valid license for the option. If "StimTrac" is available and will be used, set the radio button appropriately.

Set Stimulus Type

If the Project uses *ET3Space* data or data collected with a remote optics type system, data can be analyzed with respect to stationary scene images, or videos, or both. If only static images or only videos will be used, it is suggested that the radio button be set to one of these. In this case the various pull down menus used once the project is opened will show only choices that apply to that stimulus type. If the radio button is set to "Both", all menu choices will be available. (Note: the only disadvantage to "Both" is that menus may be cluttered with items that are not applicable if only one type of scene image will be used).

Data collected by head mounted eye trackers not using the *ET3Space* feature is usually analyzed only with respect to scene video data recorded from the head mounted scene camera. (See section 20 for brief explanation of *ET3Space*.)

Select the location for the project file
Use the browser button next to the "Project Location" item to select the computer directory that will hold the project file. The location will usually default to C:\Users\Public\Public Documents\ArgusScience\ETAnalysisData. Any location can be selected so long as all users who will need to access the project will have operating system permission to read and write files to that location.

Enter a project name
Type in a Project Name. The program will add an "etap" extension. The project will be stored in a folder with the same name. This project folder will consist of a .etap file and a hidden folder in which internal project data is stored.

The default location of project files is a subfolder under the Public Users Documents folder, since this folder is shared and accessible by all users of the computer. This default location can be changed but it is recommended that a location be used that will be accessible by all users of the computer, without requiring "Administrative privileges".

Open the project
Once all project type selections have been made and a project location and name have been specified, click OK to open the project. Under the folder specified by Project Location, the program will create another folder with the project name. This folder will contain the project file, with the project name and an "etap" extension, as well as other subfolders created by the program.

The project will open to a window with a menu bar and shortcut bar, and two blank panes separated by a vertical boundary, like that shown below. Hovering the mouse over a short cut button displays text describing its function.

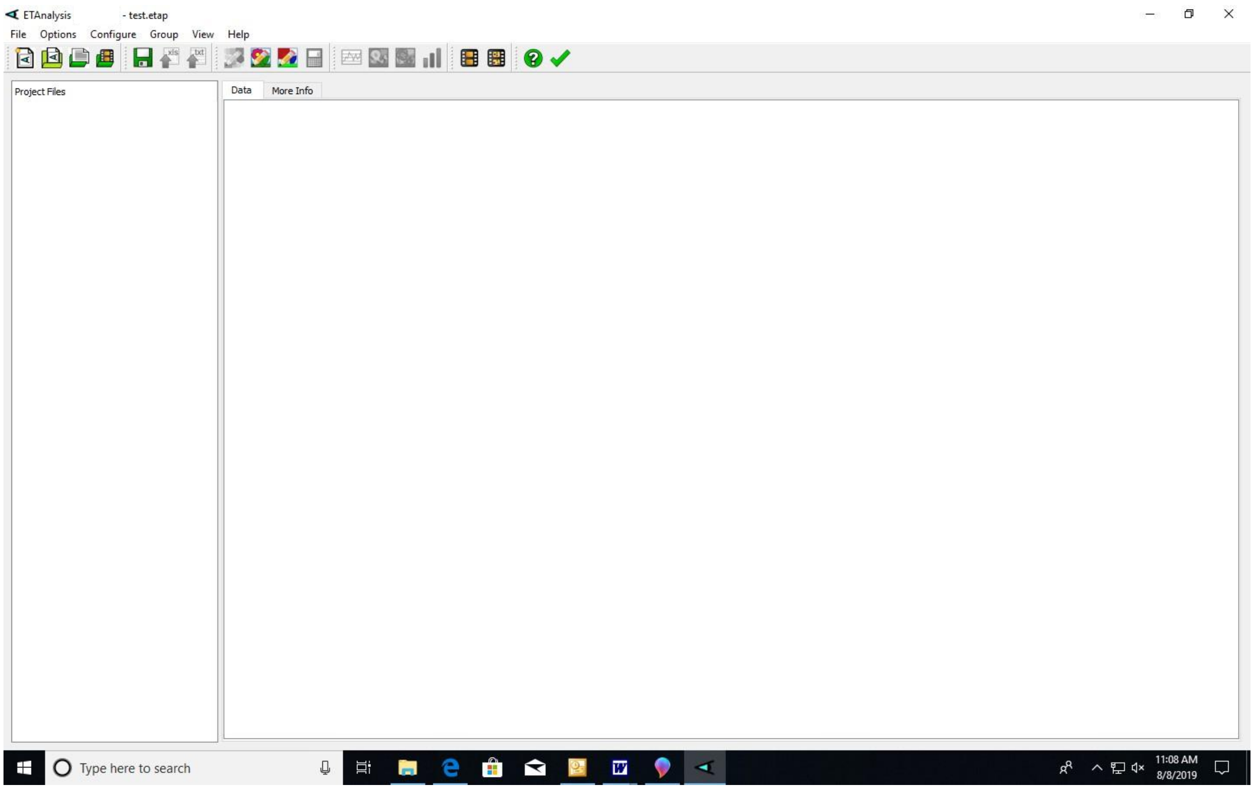

The project options selected can be examined by selecting **Options → Project Options**. Project Options can be changed to add features that are available and were not originally selected, but not to remove features that were selected. For example, if an *ETVision* project has been opened with stimulus type set to "Backgrounds (static images)", this can be changed to "Both". However, if the project was opened with stimulus type set to "Both", it cannot be changed to just allow static Backgrounds. To have a more restricted set of features, a new project must be opened.

Exiting the program will either cause the project to be saved automatically (if the "Save Project on Close" option has been enabled), or will pop a prompt to save changes. Projects can also be saved at any time by selecting **File→Save Project**, or saved with a new name by selected **File→Save Project As**. When the program is started on subsequent occasions it will open the last project saved if the "Restore last project on start up" option is enabled. **File→New Project** will automatically save the currently opened project and open a new blank project under the name and path specified by the user.

Open an existing project

From the "Open or Create Project" dialog, set the radio button to "or open an existing project". Use the browser buttons to select the project file and folder, or click the "Select from Recent Projects" button to select from the last 10 projects opened. If a different *ETAnalysis* project is already opened, select **File→Open Project** from the menu bar. Browse to the desired project file (file with ".etap" extension), and click Open. Alternately, if *ETAnalysis* has not yet been opened, use Windows Explorer to browse to the project file (*.etap), and double click the project file name.

4.2 Menu Bar Items

Items available under each menu in the menu bar (described in the remainder of this section) are customized to each project type. Items irrelevant to the project (e.g., configuring static AOIs in a SceneMap project) will not be visible. Here we describe all possible items, the majority of which are relevant to all projects.

4.2.1 File Menu

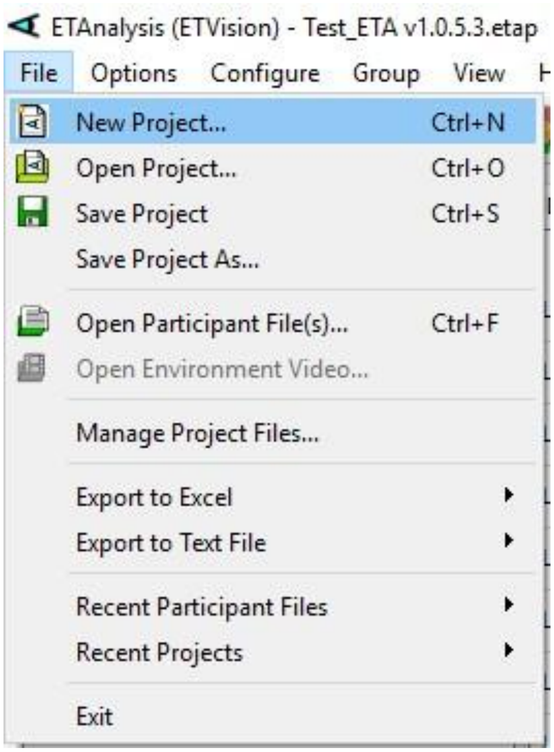

The file menu contains options for opening data files, videos (if applicable to the project), and projects. The "Recent" menu items allow opening a recently opened file by choosing between the last 10 data or project files opened within *ETAnalysis*. The "Manage Project Files" option can be used to relocate data or stimulus files which may have been moved or located on another machine (see Section 3.4 for details). The file menu also contains export options to export input data or output results to text or Excel files (see Section 19.2 for details).

4.2.2 Options Menu

The Options menu contains program features that are applicable to all projects opened within *ETAnalysis* (i.e., not project-dependent). The first 4 items in this menu can be "checked" or "unchecked" and the default states are shown in the following image (and highly recommended).

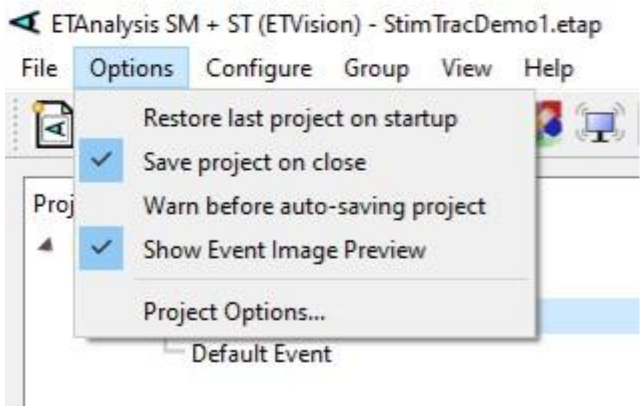

If "Restore last project on startup" is checked, when launching *ETAnalysis* the previously opened project will open. If "Save project on close" is not checked, there will be a prompt to save changes when closing *ETAnalysis*. Keep in mind, if this option is unchecked, changes made may still be saved, since the program automatically saves the project after significant changes are made. If "Warn before auto-saving project" is checked, a warning will appear each time one of these saves is made. "Show Event Preview Image" is a convenience feature that can be turned off by unchecking. This may be useful in large projects if the preview feature is slowing down the project. When this option is checked, the "More Info" tab of any event in the project (see Section 6.1 for description of events), will show a small preview image of the stimulus or scene video corresponding to this event. The "Project Options" item in the Options menu shows the options selected when creating the current project and, if applicable, allows additional options to be activated. Note that once optional features have been included, those features cannot be turned off. The ability to add features is helpful if, for example, just Background or just Video stimulus types have been selected and later it turns out that the other stimulus type is also needed.

4.2.3 Configure Menu

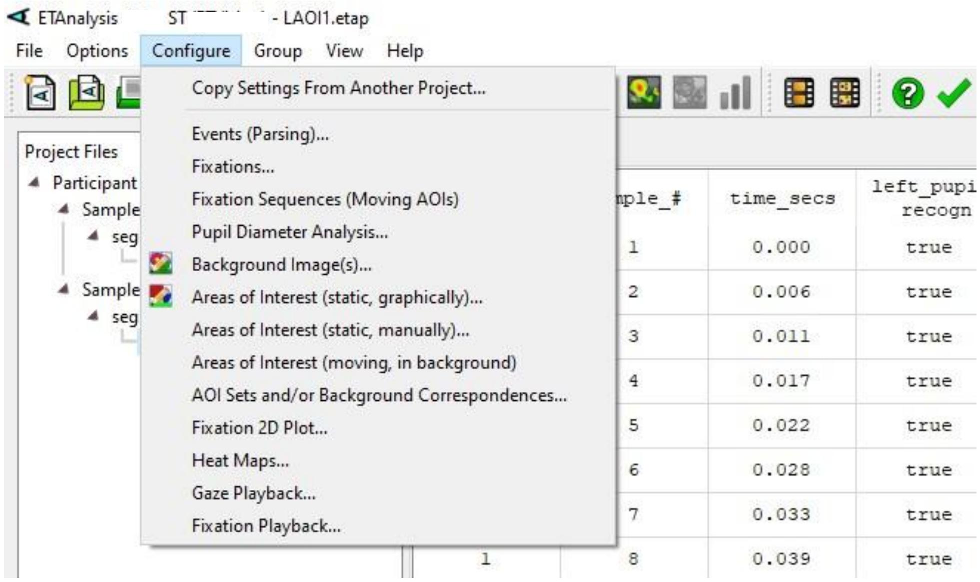

The Configure menu is used to configure settings associated with the entire project. These settings are stored in the project's ".etap" file and can be copied between projects via the "Copy Settings From Another Project" menu item. These settings include parameter settings for calculations such as fixation detection as well as background image configurations and AOIs defined within these background images. These options are described in more detail later in this manual. If only a subset of these options are displayed, it is because not all options are applicable to all project types. Options not applicable to the project (e.g., configuring static AOIs is not applicable to SceneMap projects) are hidden.

4.2.4 Group Menu

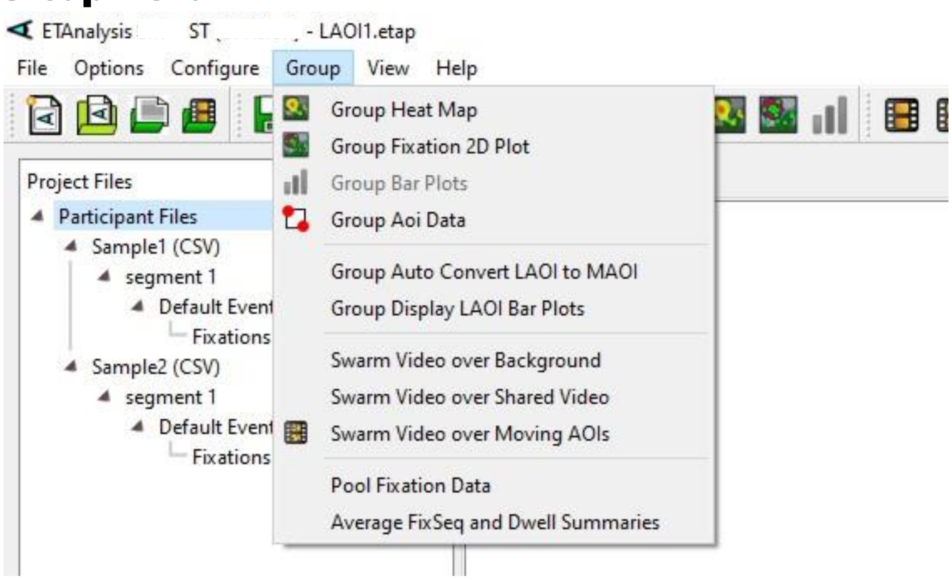

The Group menu allows access to features that combine data from multiple participants or events. These items are described in detail in section 15.

4.2.5 View Menu

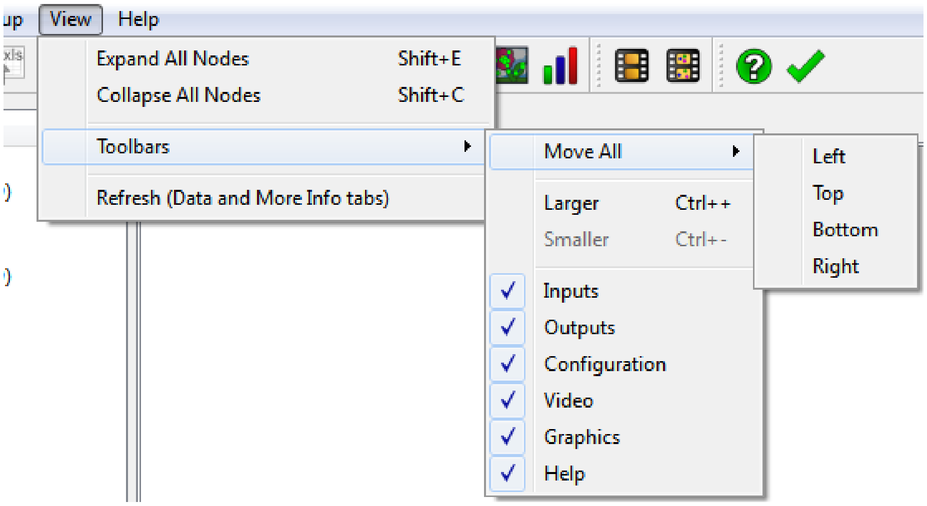

From the View menu, it is possible to expand or collapse all nodes in the project tree (see Section 2), customize the view of toolbars or refresh the display. The toolbar buttons are described in Section 4.3; but, in short, these buttons are organized into groups – Inputs, Outputs, Configuration, Video, Graphics, Help – and these groups of buttons can be clicked and dragged to move to different sections of the interface or to detach them from the interface. They can also be moved via the View menu "Move All" submenu. This menu can also be used to increase or decrease the size of the toolbar button icons via the "Larger" and "Smaller" menu items or by clicking Ctrl and "+" to enlarge or "-" to shrink the buttons. Toolbar sections can be turned off (hidden) by unchecking them in this View menu or by right-clicking anywhere within the toolbar area.

4.2.6 Help Menu

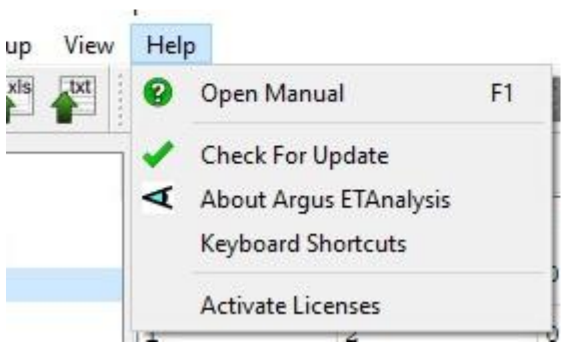

"Open Manual" will open the ETAnalysis manual that was current when ETAnalysis was installed.

The "Check For Update" feature can be used, if the computer is connected to the Internet, to determine if the current version of *ETAnalysis* is out of date. It is highly recommended to check for updates frequently, especially if encountering a problem. *ETAnalysis* does not update automatically.

Clicking "About Argus ETAnalysis" will show which version is currently running, the latest release version (if the current version is out of date and the PC is connected to the internet), and which licenses have been activated or are running as trial versions. (If the license is installed but not activated, a trial version is running.) If a license for a particular feature has not been installed, please contact argus@argusscience.com for a trial version. If the latest version is running, and all licenses are installed and activated, the "About Argus ETAnalysis" dialog should look like the one in the following image (except that the version number may be different, and the AI and ST license status may be different).

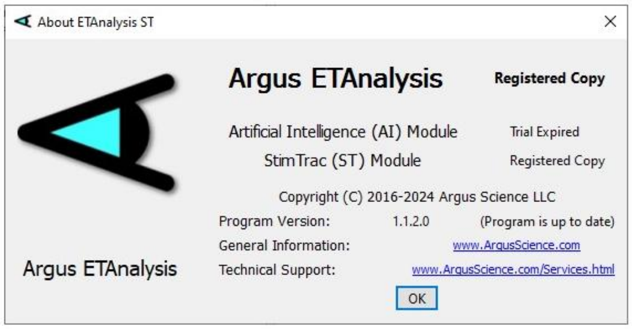

"Keyboard Shortcuts" will pop up a list of keyboard shortcuts, and "Activate Licenses" will pop up a dialog showing current license status, with buttons to initiate activation of new licenses.

4.3 Toolbar Buttons

Based on the project type and selected stimulus options, there will be a subset of the toolbar buttons shown above. Hovering the mouse over an individual toolbar button, shows a description of that button. Most project options can be accessed via a toolbar button. Many items are accessible by either right-clicking a node in the project tree (see Section 2) or by clicking a toolbar button with the desired node selected. If a toolbar button is used to access a particular option – e.g., mapping a SceneMap environment – the default project settings are used to perform that task; if the same task is accessed via right-clicking a tree node, an intermediate dialog will appear allowing changes to the settings for that task (if additional settings are relevant to that task). The following table shows each toolbar button, its corresponding task and the section of this document that describes that task in more detail (if applicable).

Button	Task Description	Manual Section
	Start a New ETAnalysis Project	4.1
	Open ETAnalysis Project	4.1

	Open Participant File	5
	Open Shared Stimulus Video or SceneMap Environment Video	16.3
	Save ETAnalysis Project	3.2
	Export current node data to Excel	19.2
	Export current node data to Text file	19.2
	Track Computer Monitor (+ST only)	17.3
	Configure Moving Areas of Interest	16.7
	Configure Static Background Image	7
	Configure Areas of Interest in Static Background	8
	Compute All Available items: Fixations, Sequences, Dwells and Pupil Diameter Analysis	10, 11, 12, 13
	Play Video for current node	16.9
	Play Swarm (Group) Video	16.10.1
	Plot current node data against time	14.1
	Show Heat Map for current node or group	14.2.1
	Show 2D Fixation plot for current node or group	14.2.2
	Show AOI Bar Plots for current node or group	14.3.1
	Open Manual (pdf)	4.2.6
	Check whether the version is up to date (if not, a link will be provided to update to the latest version)	19.5

4.4 Project tree

The project structure is represented by a tree diagram on the main program window left panel. "Nodes" are added to the tree as data and analysis results are added to the project. This has already been discussed in more detail in the "Project Structure" section (section 2).

4.5 Display Area

The Display Area is used to view raw data from the eye tracker files, interactively configure backgrounds and Areas of Interest (AOIs), and display graphical results, including video playback. This area consists of multiple tabs. The "Data" and "More Info" tabs are always present and display information for the currently selected node. If one of the topmost "Participant Files" or "Environments" nodes is selected, the "More Info" tab will show all Project Settings. This information may be helpful but is primarily intended for technical support purposes. Typically, all Project Settings can be viewed via the "Configure" menu as well, where they can also be edited.

Many tasks within *ETAnalysis* will result in an additional tab being opened in the Display Area. These tasks include configuration tasks (e.g., configuring backgrounds or AOIs) as well as display of graphical results (e.g., bar plots, fixation plots over backgrounds, video playback). These additional tasks can always be exited by clicking the small red cross button on the tab itself as shown in the following picture (red arrow).

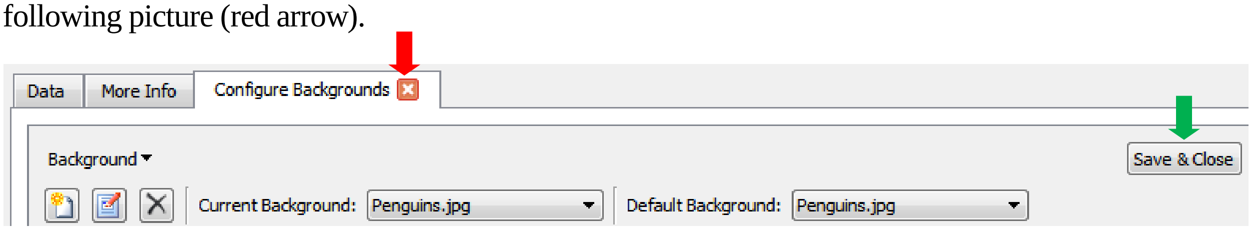

If the task involves configuration of stimuli or areas of interest, there will also be a "Save & Close" button located at the top right of the tab window (green arrow in previous image); as the name states, clicking this button will save all changes before closing the tab. Clicking the small red cross button (red arrow) will cancel editing in most cases, but will also warn if changes have not been saved as shown in the following image.

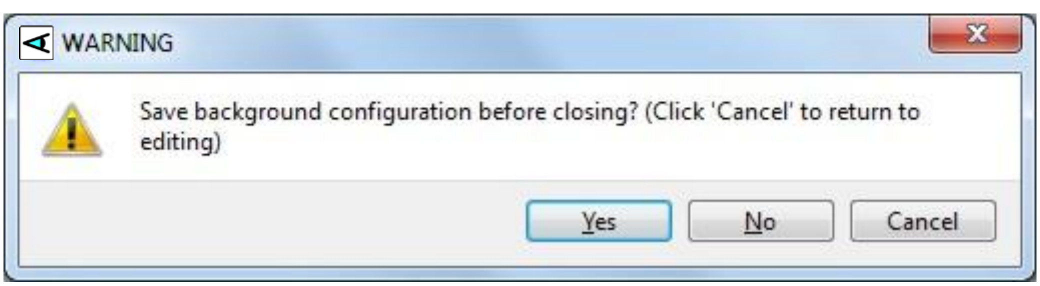

Similarly, trying to close the entire *ETAnalysis* application when one of these configuration tabs is open will result in a prompt to save if there are any unsaved changes.

In most cases, when an additional tab is opened, the rest of the application will be disabled. In these cases, the tab must be closed before any other tasks can be performed. Clicking in a disabled area will produce a window that offers to close the tab if "Yes" is selected.

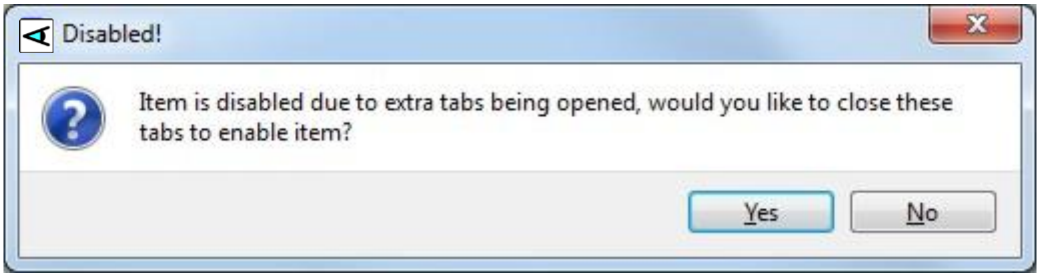
Disabled!
Item is disabled due to extra tabs being opened, would you like to close these tabs to enable item?
Yes
No

5 Opening Participant (gaze data) files

To open a data file in the project, Click **File→Open Participant File(s)** or , and browse to a data file recorded with an Argus Science eye tracker.

Highlight one or more files in the browser (hold down the <Cntrl> key to select multiple files) and click "Open". If using *ETAnalysis* for the first time, it is suggested that a single file or a small number of files be opened.

Each file will appear as node in the tree diagram, under a "Participant Files" node. The ▶ arrow symbol next to an entry in the diagram expands sub-levels of the tree diagram. Upon clicking this and expanding the node, the ▶ symbol will then change to a ◢ symbol; clicking the ◢ symbol will collapse or close the sub-levels. If an arrow symbol is not present, the node does not contain any sub-levels.

Initially, each file entry will have a list of data segments (periods of continuously recorded data) at the first sub-level of the tree. At the next sublevel, each segment will have a single "Default Event" consisting of the entire segment. Dividing segments into multiple "events" is explained in the next manual section.

Left clicking to highlight a node on the tree diagram will cause the data described by that node to be displayed in the right pane "Data" tab of the program window. The "More Info" tab will show additional information about that section of data.

In the example below, the project contains two eyd type files, Andrea_1.eyd, and Peter_1.eyd. Each file has a single segment, and since the files have just been added to the project the segments have not yet been divided into multiple events. The "Default Event" is simply the entire segment. In this case there is no difference between the "Segment" and "Default Event" underneath it.

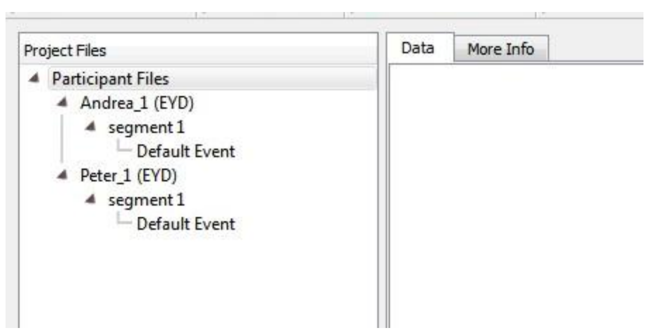

The contents of the right pane depend upon node selected in the tree diagram. When the "Participant Files" node is highlighted, the "Data" tab is empty, and the "More Info" tab contains some general information about current project default selections. When a file name is highlighted, the "Data" tab will show the list of segments recorded on that file and "More Info" will show the system configuration and related information recorded in the data file header. In the case of eyd or ehd type files this includes all eye tracker configuration information, subject calibration information, and in the case of ehd files, environment configuration information.

When a segment or event node is selected, the Data tab will contain a list of all the data in the segment or event. The list is in tabular form with each row representing a data sample, and the different data items in separate columns. In the example below, the Segment 1 node, under the "Andrea_1" file is selected. There is a data column for each data item recorded. The data items recorded to the file are generally selectable on the Eye Tracker Interface, when the data is recorded, and will vary somewhat depending on system type and configuration. The data items that can be recorded are described in the Eye Tracker manual for the system type being used. However these will almost always include sample number, time, pupil diameter, and horizontal and vertical gaze coordinates.

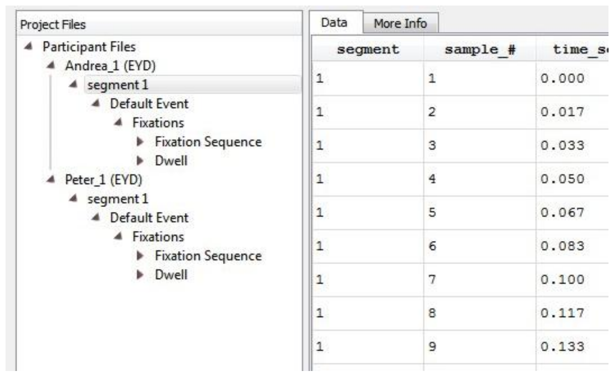

The "More Info" tab will show information about what determined the beginning and end of the Segment or Event and will include some summary information about that section of data (for example, the start time, stop time, duration, and number of records in the data section).

Right clicking a node on the tree diagram will open a context menu with additional actions that are available to further process or display the data defined by that node. These actions include dividing ("parsing") the original data into additional events, computing fixations, etc. These actions are explained in subsequent sections. An important principle to note is that most actions taken by right clicking a particular node apply to all data under that node. For example right clicking a Segment node will lead to actions that can be applied to that entire segment. This may include several events, but not other segments. Right clicking an event node will lead to actions on only that event, and so forth.

Note that data files created with *ETVision* eyetrackers can be either text files (csv extension) or binary files (eyd or ehd extension). See the *ETVision* manual for an explanation of the two data file types. Text (csv) data files will always have a companion xml configuration file, automatically created in the same directory folder by *ETVisio*n, and having the same name as the data file, but an xml extension. Although the xml configuration files do not appear on the project tree, *ETAnalysis* uses information from these files to implement or enhance some features. Be sure to keep both the csv data file and companion xml file in the same directory folder.

6 Parsing Events

6.1 Definition of Events

Data files on *ETVision* eye tracker systems can be recorded either as "csv" type text files, or as binary "eyd" or "ehd" files (see *ETVision* manual for more details). When data is recorded as an "eyd", or "ehd" type binary file, recording can be started and paused multiple times on a single data file. Continuous data, between a start and pause, is referred to as a "data segment". An "eyd", or "ehd" data file can, therefore, have multiple data segments. A "csv" type data file can have only a single "segment". When data recording on a csv file is stopped, the file is automatically closed and recording can resume only on a new file.

When a data file is opened in *ETAnalysis*, the tree diagram at the left of the main window shows the file name and also shows all of the data segments on that file as sub nodes under the file name.

ETAnalysis can further divide each data "segment" into to sub-segments based on various conditions involving time, XDAT values, or mark flags. These sub-segments are called Events and the process is called Event Parsing. All subsequent processing, such as finding fixations, etc., is done on the basis of Events. An entire data segment may be an "event", a single sub-section of the segment may be an "event", or each data "segment" may be split into multiple "events".

Initially, the tree diagram shows one event, labeled "Default Event", as a sub-node to each segment. The "Default Event" is the entire segment.

The command to parse events is available from the context menu as shown below. The context menu is invoked by right clicking either the file name or an individual segment on the tree diagram. (If invoked by right clicking the file name, it will apply to all segments in the file; if invoked by right clicking a segment, it will apply only to that segment). If "Delete → Delete Events" is selected from this menu, or if event parsing fails, the segments under the selected node will revert to the "Default Event".

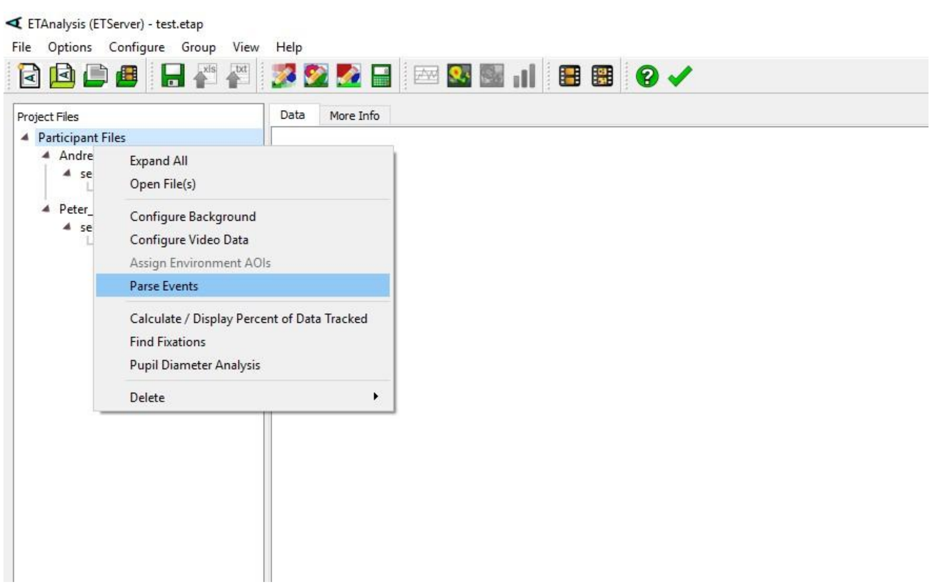

The user is then presented with the **Configure Events** dialog shown below

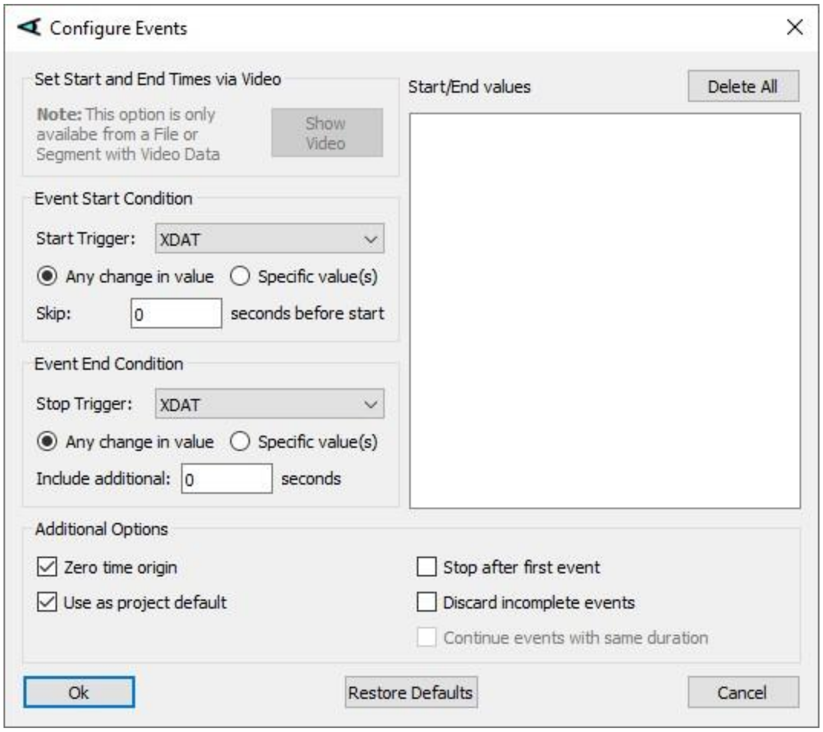

Selections on the dialog will be explained in the following sections.

6.2 Event Start condition

The start condition can be one of the following:

None. The event starts immediately with the first data record in the segment.

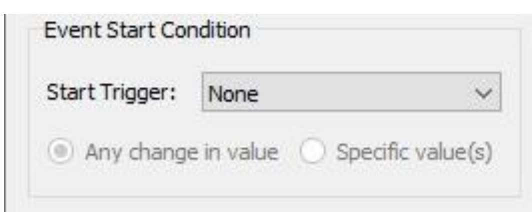

XDAT. The event starts on the first record that has an XDAT value contained in the user-defined list of Start values; or, if the user has set the radio button to "Any change in value", on the first record with an XDAT value different from the previous record. Note: the XDAT value on the very first field is always considered a "change" and will trigger an event if "Any change" has been selected. On the example below, events are set to start when XDAT changes to 1 and when XDAT changes to 2, or to 3.

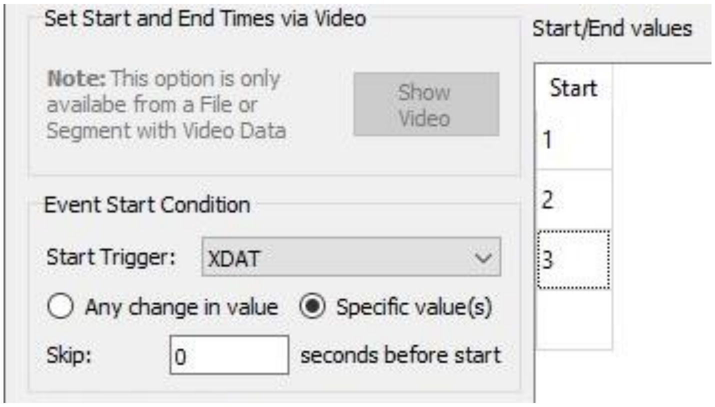

Mark_Flag. The same as XDAT, except that the event starts on the first record that contains one of the specified Mark Flags. If "Any change" radio button is set, the event will start on the first record containing any mark flag.

Time. The event starts at the specified time, entered as number of seconds from the beginning of the segment. The user can specify more than one start time value to create several events. In the example shown below, the first event would start 10.5 seconds after the beginning of the data segment, the next event would start 20 seconds after the beginning of the segment, and third event would start 30 seconds after the beginning of the segment. *It is important to note that this can create overlapping events, if one event specifies an end time that is later than the start time for a subsequent event.* If multiple events are created using "Time" as the "Start Trigger", pay careful attention to the stop condition (described in the next section).

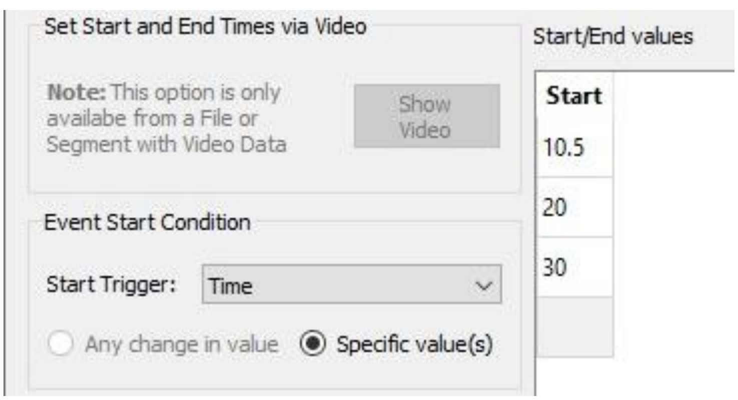

Skip … seconds before start. If *Start Trigger* is *XDAT or Mark_Flag*, the user can also specify additional interval that will be skipped before the event start. In other words, if the skip time is *t*, the event will start t seconds after the Start Trigger is encountered. Suppose, for example, that the Start Trigger is *XDAT*, and that *any change in value* and *Skip 5 seconds* are specified. Further suppose that the first change in XDAT is 7 seconds from the start of the segment. In this case the first event will start 7+5=12 seconds from the beginning of the data segment.

6.3 Event Stop condition

Event stop (or end) conditions are very similar to start conditions and contain the same choices, plus one additional choice called "Next Event Start". The Start and Stop Triggers can be different. Almost any combination of Start and Stop conditions can be used. The dialog grays out the combinations that are illegal. “Next Event Start” means that there is no explicit stop condition, and the event will end when the next Start condition is encountered.

Note that when the stop condition is a time value, it specifies event duration rather than a time from the beginning of the segment. This is different from the time start condition, which is time from the segment beginning.

Special cases:

1. In most cases the event continues until the stop condition is met, and any start conditions are ignored until the event has ended. One exception is when the stop condition is “Next Event Start”. The other exception is when the start condition is a time value. In this case, an event can start before the previous event end, and overlapping events are possible.

2. If the stop condition is never satisfied, the event continues to the end of the data segment.

3. If the *Stop Trigger* is *None*, the event continues to the end of the data segment.

4. If the *Stop Trigger* is *Time*, the value determines the event duration. In other words, stop time is measured from the event start rather than from the segment start.

5. If multiple XDAT or Mark_Flag values are specified as *Start Trigger*s and the *Stop Trigger* is Time, then a different duration may be specified for each *Start Trigger*. In the example shown below, all events starting with XDAT=1 will continue for 10 sec, events starting with XDAT=2 will continue for 20 sec and events starting with XDAT=3 will continue for 30 sec.

Start/End values

Start	**Duration (sec)**
1	10
2	20
3	30

6. If *Start Trigger* is XDAT or Mark_Flag and *Stop Trigger* is Time, the next event will start only after XDAT / Mark_Flag has changed. For example suppose that the Start Triggers are XDAT values 1 and 2 and Stop Triggers are Duration values 10 and 15 seconds. Further suppose that the first 20 sec in data segment have XDAT=1, the next 20 sec have XDAT=0, followed by 20 sec with XDAT=2. The first event will start with the first record of the data segment (XDAT = 1) and will end after 10 seconds. The remaining 10 sec of data with XDAT=1 will be ignored. The next event will start 20 seconds from the beginning of the segment, on the first record containing XDAT=2, and will continue for 15 sec.

6.4 Parse by Video

If there are no markers saved within the data file to aid in parsing data based on task or stimulus, it is possible to visually parse the data using the eye tracker scene video, if a video is available. This possibility is available if the project “Stimulus Type” is “Videos” or “Both”, and the scene video has been properly configured as described in section 16.

To use the Parse by Video feature, Click “Show Video” in Configure Events window (A), then select the start (B) and end frames (C) of each stimulus presentation or task. Use the “Play” button or the video slider to advance to the vicinity of the desired start or stop frame. Once in the vicinity of the desired frame, use the left and right arrow keys or frame advance buttons to conveniently find the proper frame. When done marking start and stop frames, close the video player and choose “Ok” in the Configure Events window.

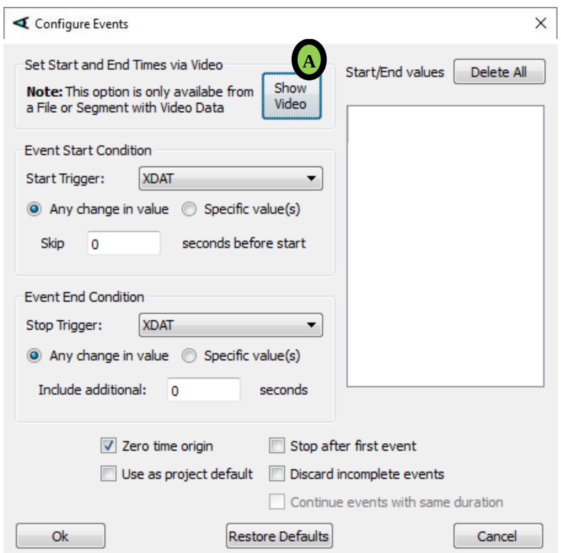

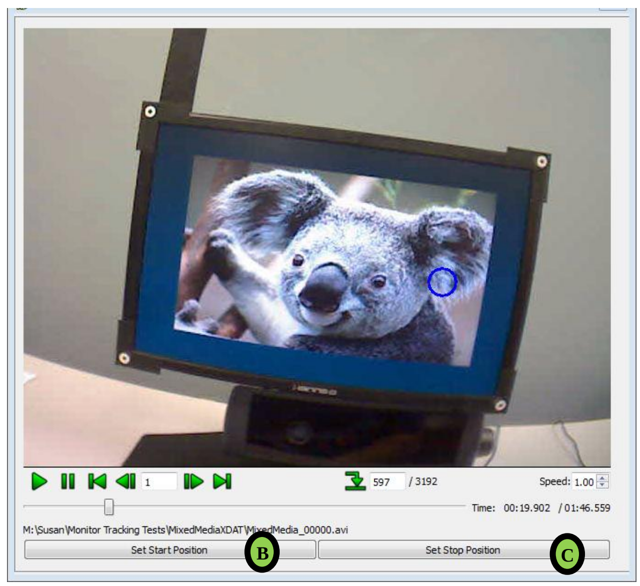

6.5 Additional options

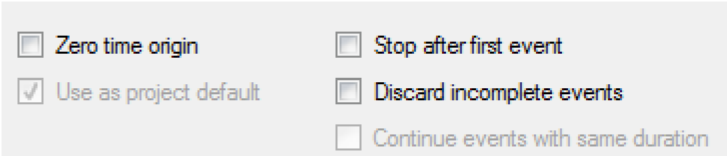

Zero time origin. By default, *Zero time origin* is selected and the time stamp of the first record of each event will be reset to zero. The time of each record in the event is calculated with respect to the first record of the event. To make each event record show a time stamp corresponding to time from the beginning of the data segment, uncheck this box.

Stop after first event. If this option is selected there will be not more than one event in the segment. After the end of the first event, the program will not start another event in the data segment.

Discard incomplete events. Do not create an event if stop condition is not met before the end of the segment. For example if an event stop condition is XDAT = 2, and after the event begins no record with XDAT=2 is encountered before the segment end, this event will not be created. If *Discard incomplete events* is not checked, this event will be created and will end at the end of the data segment.

Continue events with same duration. *Available only when both Start and Stop triggers set to Time*. When this option is selected, program will continue parsing events with the same duration as the last defined event until the end of the segment is reached.

Use as a project default. Use the current *Configure Event* dialog selections as the default for future event parsing operations.

7 Configure Static Background Images

Static backgrounds are applicable if static background images have been presented to participants and data was recorded with a remote or head restraint mounted system, or if data was recorded using *ET3Space*, or if a *Stimulus Tracking* ("Stim Trac") project. (See sections 17 and 20 for explanations of *Stimulus Track* and *ET3Space*, respectively). This type of analysis will be most appropriate when subjects looked at static images such as pictures, text, web page images, etc. If subjects looked at dynamic presentations, or if gaze was recorded with respect to a head mounted scene camera image, analysis with respect to moving images may be more appropriate, and this is discussed in a subsequent section.

In order to display 2-dimensional plots or heat maps it is necessary to configure one or more **Backgrounds**. A "Background" may be a blank screen, drawing or image that represents the scene viewed by the subject. If an image file is used, it can contain an image that was displayed on the presentation computer, a digital photograph of the scene (typical for head mounted optics), or a drawing that represents the scene that was viewed by the subject (for example, a sketch of an instrument panel that the subject viewed). If the image is a drawing of a physical scene that was viewed by the subject (I.e., an instrument panel), it should be to scale so that features in the drawing have the same spatial relation to each other as the real features. The program supports the following file formats: BMP, JPEG (JPG), GIF, or PNG.

In order to superimpose point of gaze on the image, the program needs to know how to translate the eye tracker coordinates to the pixel location on the image (called VGA coordinates). To define the transform we need to specify four points with known image locations and eye tracker coordinates. These are called "Attachment Points". It is best if the attachment points span a significant portion of the image. Ideally the points will be near the 4 corners of the image. These points should be easily identifiable landmarks in the image.

If *ET3Space* was used to acquire the data, there may be multiple "scene planes", and in this case 4 attachment points must be specified <u>for each scene plane</u>.

The eye tracker coordinates corresponding to the attachment points on an image file can be determined in advance. For example, if *ET3Space* was used, use the "pointer test" function, or measure to find the *ET3Space* coordinates associated with any point in the scene image. See *ET3Space* manual for details.

Any background that has been configured and made part of the project can be designated the "default background". If other image files have the same resolution and will have the same correspondence between image and eye tracker coordinates, the default background attachment points can be used for these image files as well without requiring the attachment point placement procedure for each file.

From the *ETAnalysis* main menu, select **Configure → Background Image(s)…**

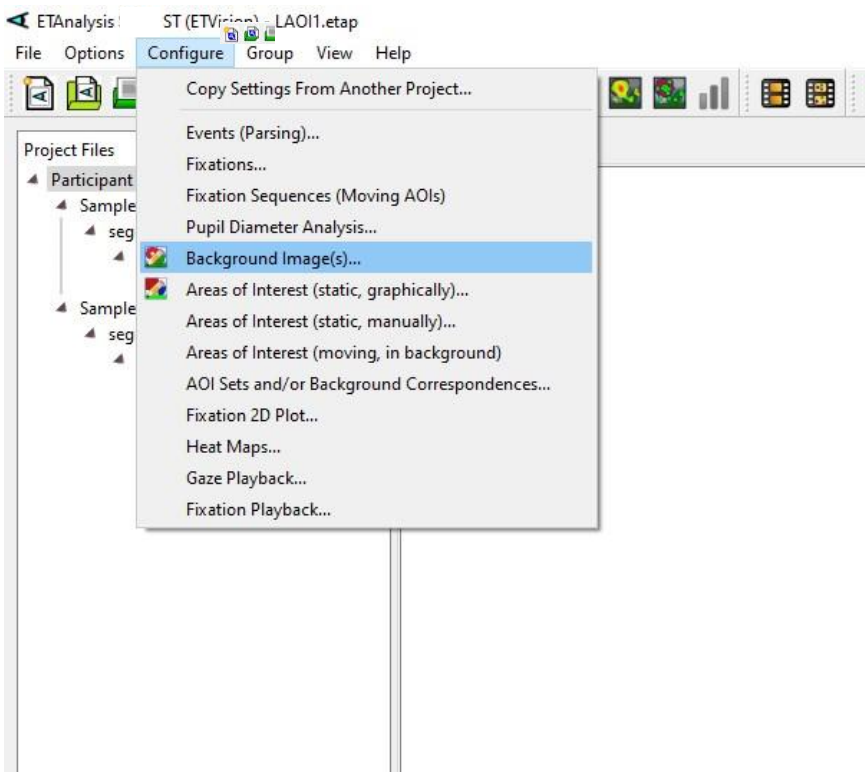

A tab labeled Configure Backgrounds will appear in the right pane of the *ETAnalysis* program window. From the upper left corner of the Configure Backgrounds tab, left click the pull down menu labeled **Background** and select **Create Single Background**.

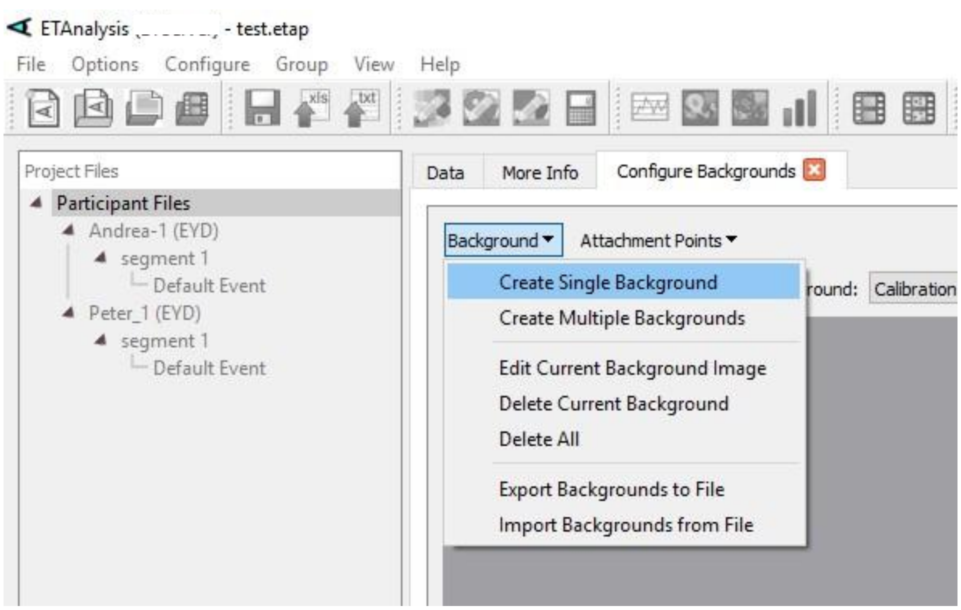

or click the Add New Background button 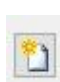 . A Create Background dialog will appear.

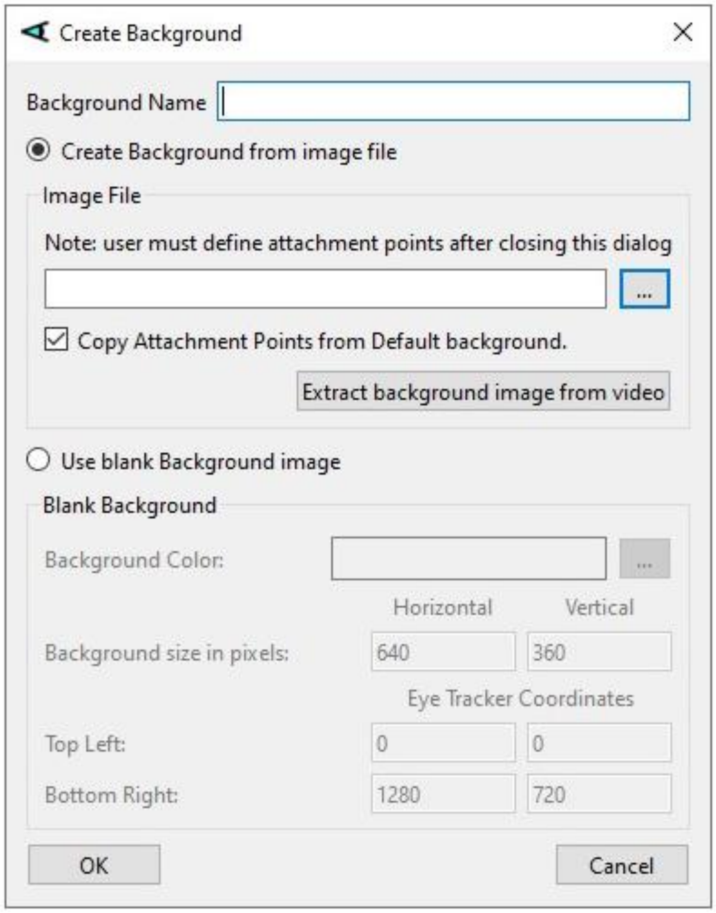

Blank Background Image

If plots are to be superimposed on a blank screen, set the radio button to “Use Blank Background Image” and select the color and size of the desired image (defaults are white color and 640 x 480 pixel image size). Type any text in the “Background Name” box. This name will subsequently identify this image and associated configuration parameters.

Under “Eye Tracker coordinates” select the horizontal and vertical gaze coordinate values that will correspond to the 4 corners of the blank image. If not using *ET3Space*, the units are Assuming data was <u>not</u> collected using *ET3Space*, top left coordinates of (h = 0, v = 0) will usually be appropriate, and bottom right coordinates will usually be [h=1280, v=720] for *ETVision* data. These are the default settings. If using *ET3Space*, the logical coordinate space depends on the coordinate frame assigned to the scene plane and the physical size of the scene plane.

When the “OK” button is clicked, a blank image will appear with attachment points, corresponding to the 4 corners labeled “P1”, “P2”, “P3”, and“P4. Click the “Save & Close” to close the image window. This background image with attachment points is now part of the project, and will be available for superimposing scan plots and heat maps.

Create Background Image From File

If an image file (rather than a blank screed) is to be used, set the radio button to "Create Background from image file", click the browse button under "Image File", browse to the desired jpg, bmp, gif, or png file. If the project already has a default background with attachment points that will be correct for the new image as well, check the box labeled "Copy attachment points from default background". Attachment points are explained in more detail in the next section. Type in a "Background Name", and click "OK". The image from the selected file will appear. If "Copy attachment points from default background" was checked, the two attachment points will be shown. Just click "Save & Close" to add the configured background image to the project.

If "Copy attachment points from default background" was not selected an **Add/Edit Attachment Points** dialog will automatically appear. (If "Copy attachment points from default background" was selected by mistake, click Attachment Points (Add/Edit Attachment Points to bring up the dialog.)

Select four easily identifiable landmarks, near as possible to corners of the image, as previously discussed, and type in the "Eyetracker Coordinates" for each of these points. In the example, below, the selected points are corners of two of the cookie box images.

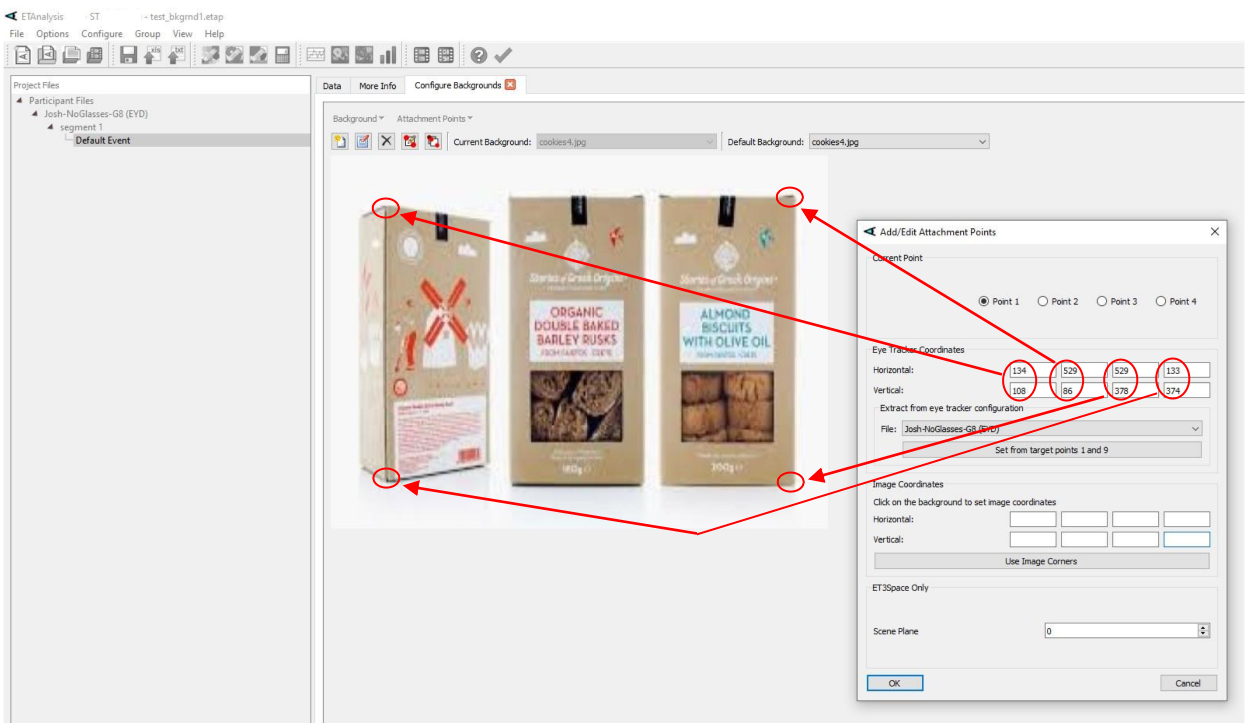

Set the radio button set to "Point 1", and use the mouse to click on the corresponding point in the image. A red dot with the label "P1" should appear at that point, and the pixel coordinates of the point will appear in the "Image Coordinates" section. Repeat for the other three selected points ("P2", "P3", and "P4"). Ignore the "Scene Plane no" unless the image is to be used with *ET3Space* type data.

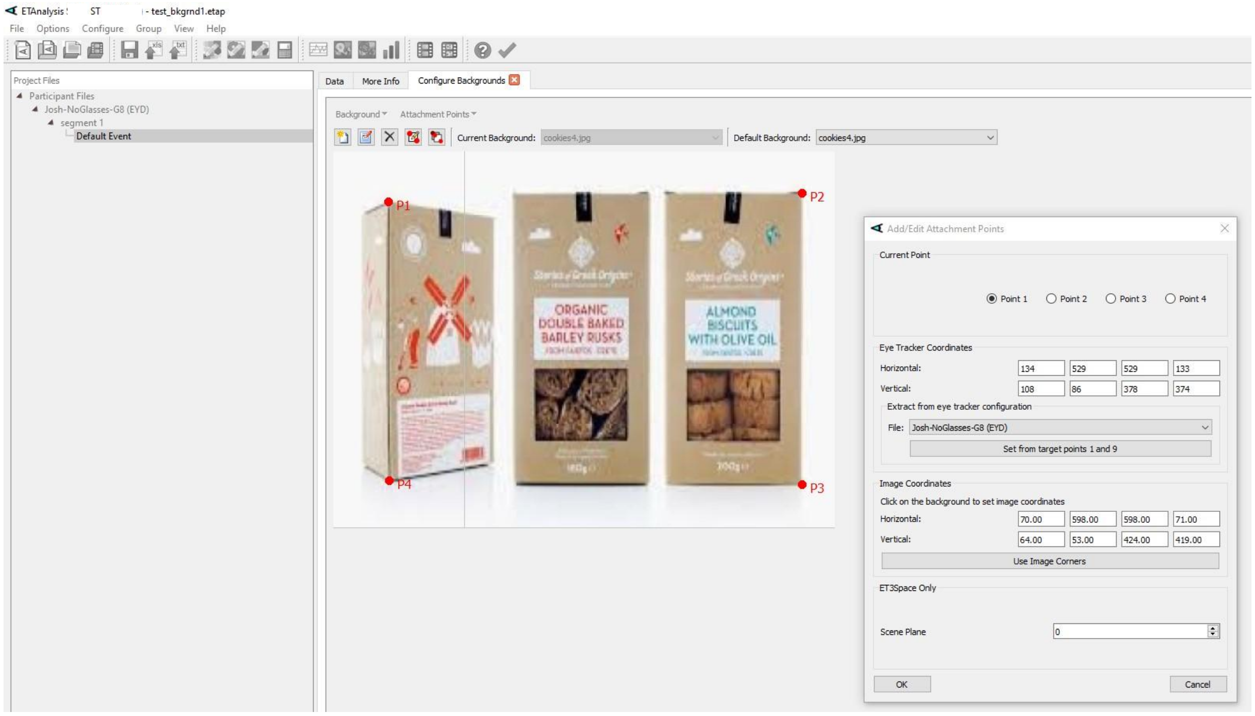

If the background is an image file that was viewed by the subject, and if the corners of the image area were visible on the subject display, then the image corners can be used as the landmarks for attachment points, as shown in the example image below.

Click OK to close the **Add/Edit Attachment Points** dialog.

Open and configure as many images as desired for use in the project. Be sure to give each a unique "Background Name". If multiple image files have the same resolution and were displayed to the subject in the same way, then it will not be necessary to find unique attachment points for each of them. Set attachment points on one of the set, as described above, and use that image as a "default" for attachment points on the others. This is described in more detail below, in section 7.1.

7.1 Using a default background image

If two image files have the same resolution and are displayed the same way (by the same computer application, etc.) then a given image pixel, say the 20th pixel from the left on the 20th row, will appear in the same spot on the monitor screen for both images when displayed to the subject. This pixel will therefore correspond to the same gaze coordinate for both images. The same information can used to transform between gaze coordinates and image coordinates for both images.

Assume that *Image_1.jpg, Image_2.jpg, Image_3.jpg, and Image_4.jpg* were all created the same way and were displayed to the subject the same way. Add *Image_1.jpg* to the project, give it a name, for example "Image1", and set attachment points as previously described. On the Configure Background Selection window, set "Default Background" (at the top of the window) to "Image1".

Now open the Create Background dialog**,** browse to *Image_2.jpg*, name it, check the box labeled "Copy Attachment Points from Default Background", and click OK. Image2 should appear with red dots labeled P1 and P2. They will be in the same position as they were on Image1, but since it is a different image they will probably not be on easily distinguished landmarks in the image. That is OK. The transformation between gaze and image coordinates will be correct. Leave "Image1" as the "default background" and repeat the procedure for the other two image files.

7.2 Images to be used with *ET3Space* data

Gaze data collected with the *ET3Space* feature can be on different scene surfaces. Every data sample contains a scene plane number, and a set of coordinates that represent gaze position on a reference frame attached to that surface. The gaze coordinates represent real distance units (inches or centimeters) along two coordinate axes that have an origin and orientation on the surface specified by the user (see *ET3Space* manual).

A background image to be used with *ET3Space* data may depict multiple scene planes, and separate attachment points must be defined for each. The background image may be a photo that includes multiple scene plane surfaces, or separate photos of each scene plane assembled into a single image with a picture editor. Alternately a graphics or drawing program may be used to create proportionately correct depictions of each surface. Perspective distortion (E.g., a photo taken from an angle) is handled correctly.

Each surface depicted must have four landmarks, preferably near corners of the surface, whose gaze coordinates are known. These will be used as attachment points. The gaze coordinates can be determined by measuring along the user defined coordinate axes on a given surface, or by using the eye tracker “Pointer Test” mode.

Open the image file in *ETAnalysis* as previously described. On the **Add/Edit Attachment Points** dialog, set the “Scene Plane” to 0 (at the bottom of the dialog window), and enter the gaze coordinates for the scene plane 0 attachment points under “Eye Tracker coordinates”. Click on each of the plane 0 attachment points on the image, and dots labeled “P1”, “P2”, "P3", and "P4" will appear as previously described.

If the image includes Scene Plane 1, now change the “Scene Plane”, on the **Add/Edit Attachment Points** dialog, to 1. Enter the gaze coordinates for the plane 1 attachment points under “Eye Tracker coordinates”. With the radio button on “P1”, use the mouse to click the first attachment point on the depiction of scene plane 1. Set the radio button to “P2” and click the second landmark on scene plane 1. Repeat for points 3 and 4. The labels on these points will appear as “P1,1” (plane 1, point1) and “P1,2” (plane1, point2), etc. The labels on the plane 0 attachment points will change to “P0,1” and “P0,2”. Repeat the procedure for any other scene planes depicted.

7.3 Exporting and Importing Background configurations

Background configuration information can be exported to an XML type file, for future import to other projects, or to protect against accidental loss. On the Configure Backgrounds tab, pull down the Background menu and select **Export**. Browse to the desired folder location, type in a file name and click **Save**. Background configuration information for all current backgrounds (all those listed under the “Current Background:” pull down menu on the Configure Backgrounds tab) will be saved to the specified file.

To import a set of configured backgrounds that was previously saved (exported), first open a Configure Backgrounds tab if not already opened (**Configure** → **Background Image(s)…**). Left click the pull down menu labeled **Background,** and select **Import**. Browse to the previously saved xml file and click **Open**. All backgrounds in the saved set will now be available under the **Current Background:** pull down menu.

Note that the saved xml file contains the scaling information for the set of backgrounds, created as described in the preceding sections, and has pointers to the original image files. It does not include the actual image files. For the import to work, the original image files must be at the same path location as when the xml file was created. If one of the image files is no longer in the same location, when that file is selected, on the **Current Background:** pull down menu, a warning message like the one shown below will appear. The message shows the path and name of the file that could not be found.

To restore the configured background image, click “Yes” to bring up a browser window, and browse to the current location of the image file.

To import configured backgrounds to an *ETAnalysis* project on another PC, the saved xml file, and all of the image files must be copied to that PC. Unless the image files are copied to the same path locations as on the original PC, it will be necessary to browse to each image file as described above.

8 Static Areas of Interest

Areas of interest (AOIs) are rectangular subsections of the scene surface defined by the user. In the case of *ET3Space* data there are multiple scene plane surfaces and areas of interest can be specified on each of them. Many of the statistics that can be produced by *ETAnalysis* relate fixations to AOIs.

Note that if data was collected with *ETVision* not using *ET3Space* or *StimTrac*, then gaze position is specified as a position on the head mounted scene camera image. Since environment objects being viewed by the participant usually move about on this image, it will usually be more appropriate to use Moving Areas of Interest (MAOIs) as discussed in section 16. Even if *ET3Space* or *StimTrac* have been used, if subjects were viewing a moving display such as a video presentation on a display monitor, then video analysis with Moving Areas of Interest may be appropriate (section 16).

AOIs are defined by top, bottom, left, and right boundaries, expressed in the scene reference frame. If the project has a background image corresponding to the scene, AOIs can be defined graphically on the background image. This is usually the easiest method. Alternately the boundaries can be entered manually (typed in). In this case the boundary coordinates can be determined using the Eye Tracker "Set Target Points" function; or in the case of Eye Head Integration data, by measuring or using the Eye Tracker "Pointer Test" function.

Each AOI can be given a name and collections of AOIs are organized in named sets. The project can contain multiple sets of AOIs. An event can be associated with any AOI set in the project for the purpose of computing fixation and dwell statistics.

AOI sets do not appear in the tree diagram, but all sets currently available to the project are listed in a drop down menu on the **Configure → Areas of Interest** dialog and all other dialogs on which an AOI set must be specified.

8.1 Defining Areas of Interest graphically

From *ETAnalysis* main menu select **Configure → Areas of Interest (static, graphically)**

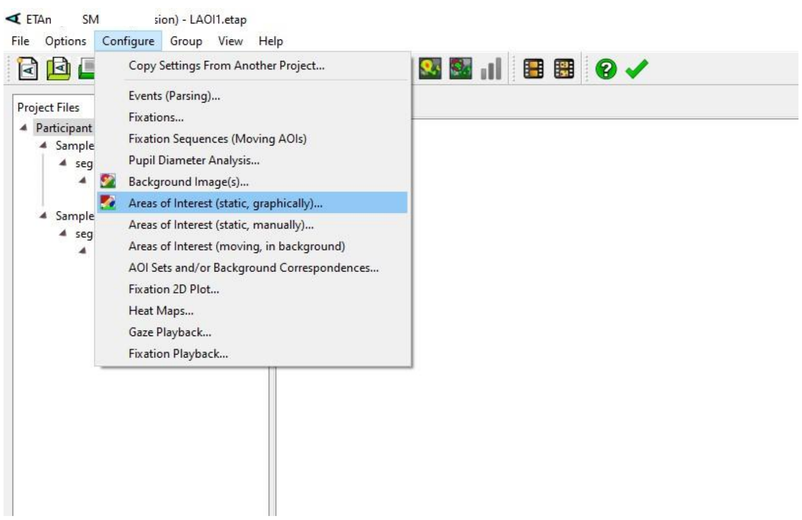

In the **AOI Configuration** dialog select a Background image and in the AOI Set selection combo box enter a name for the AOI set by typing it in the combo box. If this is the first set created in the project, the default name will be Aoi_Set_01. This name can be changed at any time.

From the AOI menu, select **Draw Area of Interest** (or use the shortcut <Ctrl>A). The mouse pointer will change to an "AOI" symbol, and remain in that form until the right mouse button is clicked.

8.1.1 Draw rectangular areas

To draw a rectangular AOI click to depress the "Draw rectangular AOI" button. 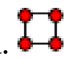

Holding down the left-mouse button, drag a rectangle of any desired size and release the mouse button when done. The area will appear as a shaded rectangle (drawn over the left penguin head, in the example below), and an AOI Properties window will appear. Replace the default name ("AOI 1" in the example below") if desired, by typing over the default name. If using Eye Head Integration data (ehd file), set the "Scene Plane:" to the scene plane number on which the AOI has been drawn. If not ehd data, leave the Scene Plane set to "0". The line thickness and color of the area outline will default to the values shown, and can changed if desired.

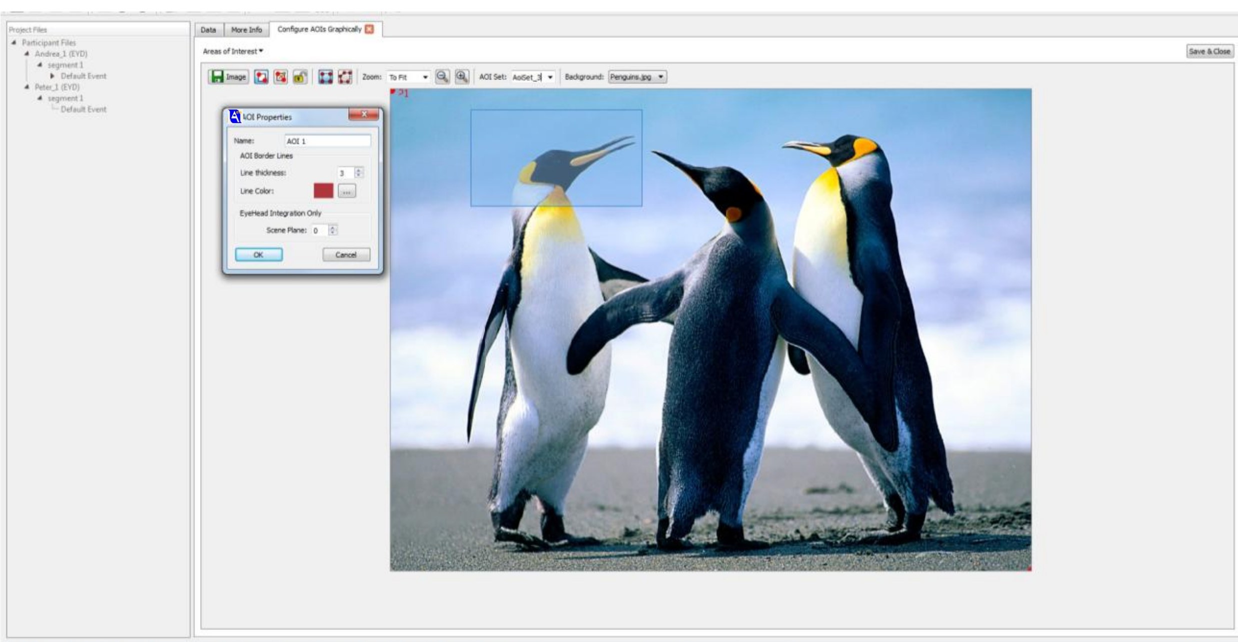

Click the "OK" button on the AOI Properties window. The AOI Properties window will close, and the AOI will be shown by a an outline of the specified color, with corner and side handles (small white squares) as shown below. A legend box at the upper right of the pane lists all areas currently created.

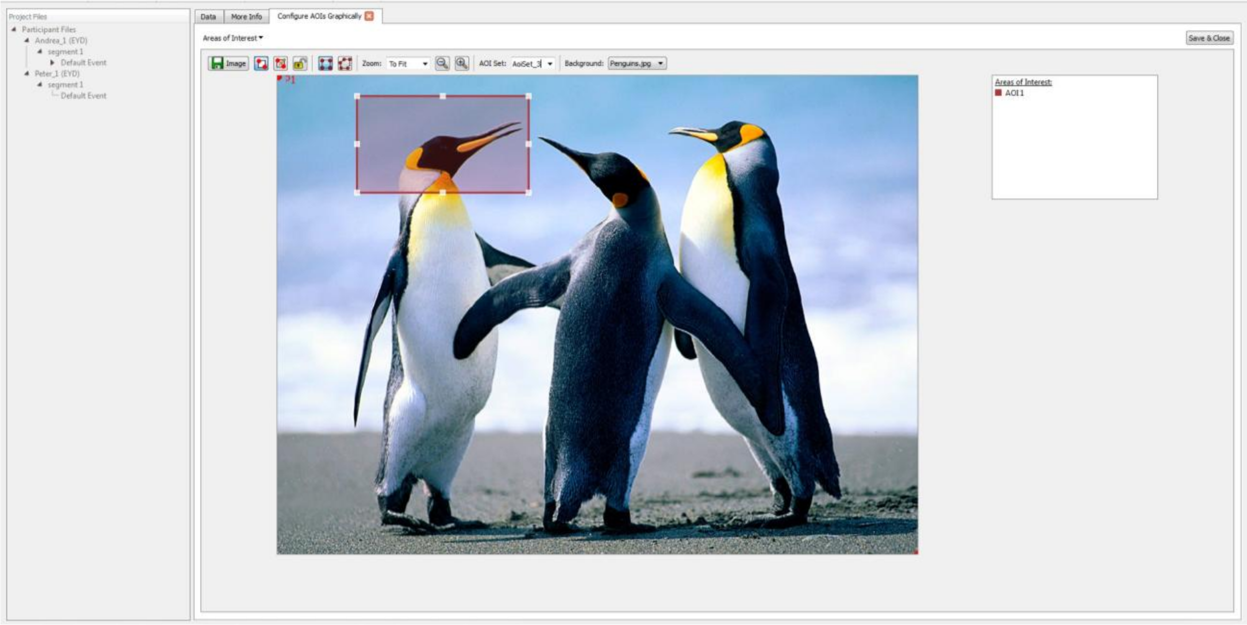

When the mouse arrow is hovered over one of the handles, it will change to a double arrow symbol and the AOI can be stretched or compressed in the indicated directions by holding down the left mouse button and dragging. When the mouse arrow is held inside the area, it becomes a hand symbol and the entire area can be moved by dragging with the left mouse button. Right clicking with the area pops up a menu that can be used to bring back the AOI Properties dialog, make a copy of the area, or delete the AOI.

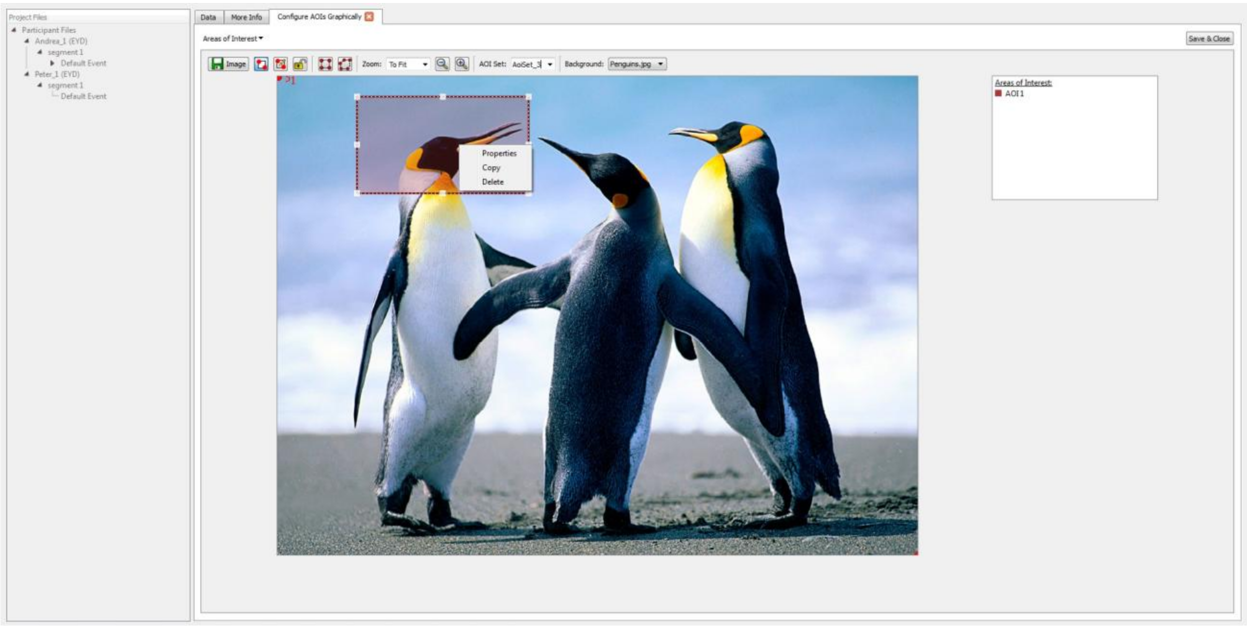

8.1.2 Draw Polygons

To draw a polygon, first click to depress the draw polygon button. 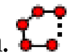

Left click on the image open an AOI Properties dialog and set the AOI name and other properties just as with rectangular AOIs. Note, however, that no AOI is draw until the AOI Properties dialog is closed by clicking "OK". At this point a triangle will appear at the spot on the image that was left clicked. The triangle (just above the middle penguin head in the example, below) will have a handle at each vertex.

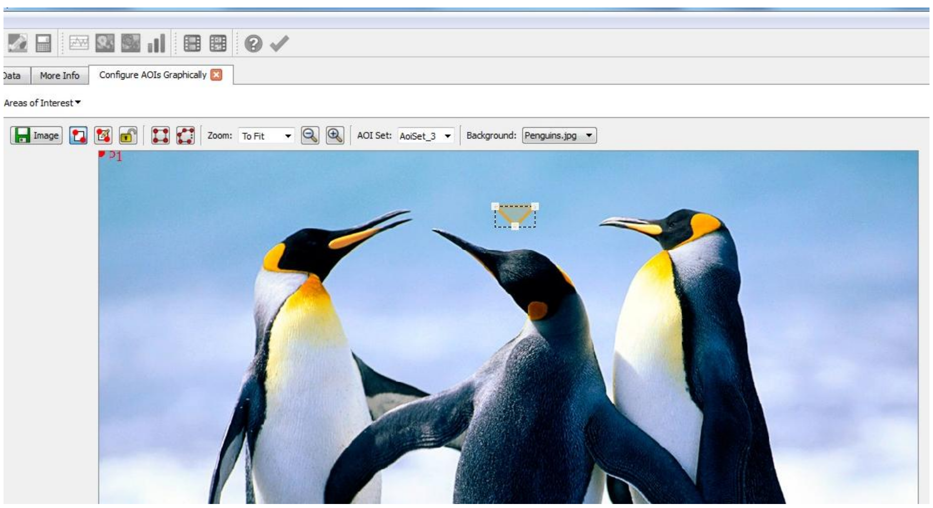

Left drag any of the handles (mouse arrow changes crossed arrows when on a handle) to move just that vertex, or place the mouse inside the object (mouse arrow changes to hand) and left drag the entire object. Left clicking inside the object will also cause a bounding box, formed with dashed lines to appear on the object. Left click on the bounding box to make it disappear.

Click on one of the lines to add a vertex (mouse arrow changes to gray cross when held in proper position to form a new vertex), and, holding down the left-mouse button, drag the newly created vertex to the desired position; repeat this process until all sides of the polygon are constructed. See the example sequence below.

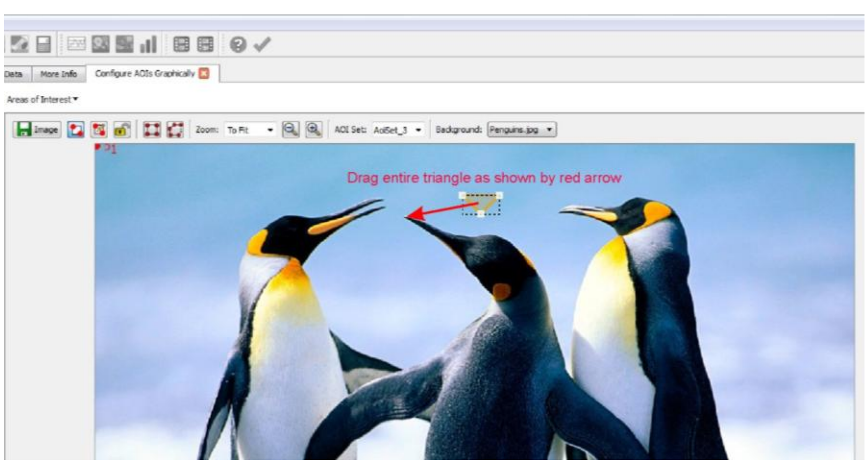

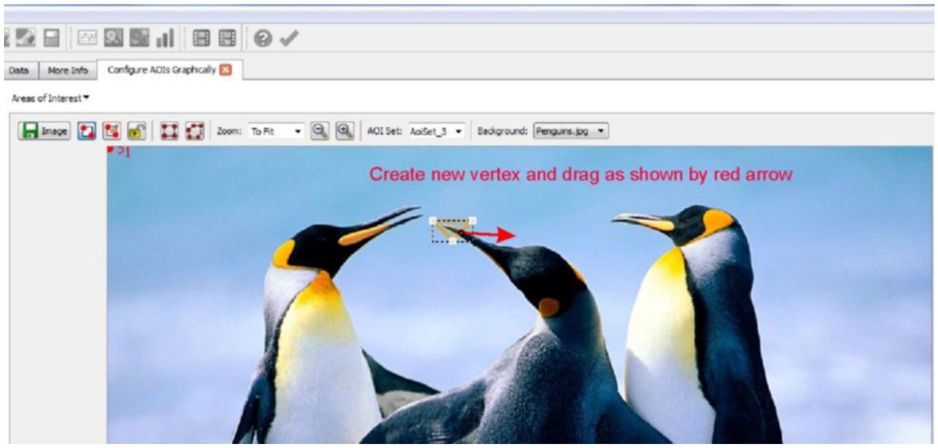

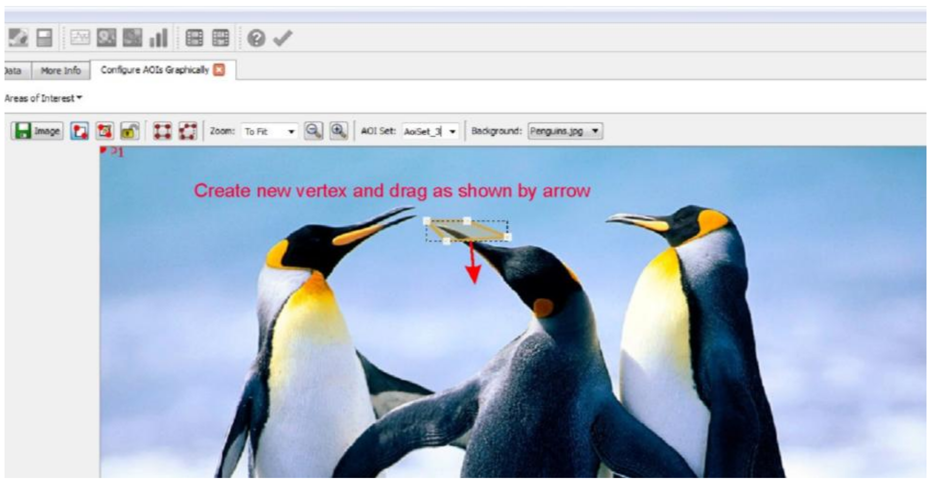

Continue in a similar fashion to obtain the result shown below.

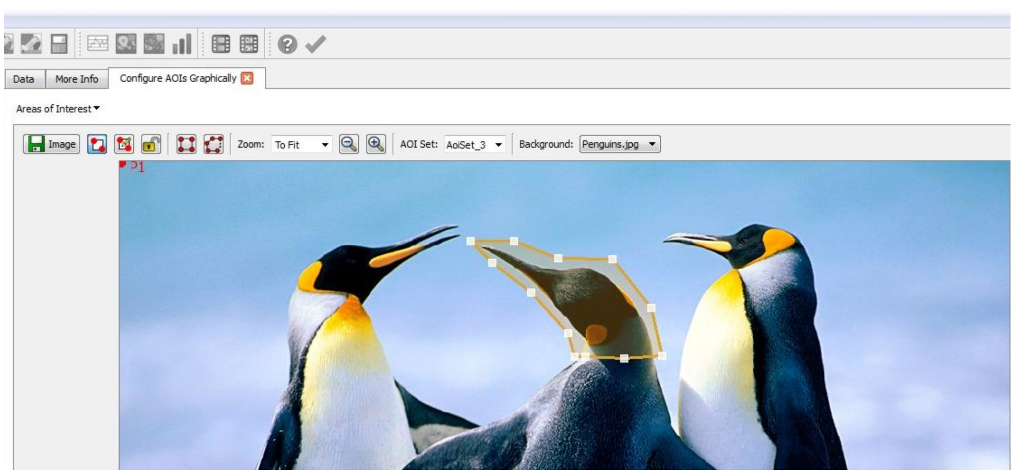

As with rectangular AOIs, right clicking with the area pops up a menu that can be used to bring back the AOI Properties dialog, make a copy of the area, or delete the AOI.

8.1.3 Saving AOI sets and other dialog window features

To create a new AOI as part of the current AOI set, simply depress either the rectangle or polygon button, and left click on the image beyond the boundary of any existing AOI. Then proceed as described in the previous sections.

Important: if are using multiple scene planes (typical for head mounted optics with *ET3Space* or *EyeHead Integration* feature) be sure to select the appropriate scene plane in the AOI Properties dialog.

To make a new AOI set, select Areas of Interest (Create Blank AOI Set, then proceed as described in the previous sections. To edit a different AOI set, previously created in the project, use the "AOI Set:" and "Background:" drop down menus to select the desired AOI set and background image. Selections in the "Areas of Interest" drop down menu can be used to delete all of the AOIs in a given set, or to delete the entire set.

To make a new AOI set by modifying an existing set, select **Configure→Areas of Interest →Copy AOI Set**. The currently selected AOI set will be copied as a new set and can be named and modified as desired.

AOI sets can be saved for export to other projects, and AOI sets that have been saved by other projects can be imported to the current project. To save the current AOI sets for potential export to other projects, select **Areas of Interest→Export AOI sets to File**. Use the resulting browser window to specify a path and file name for the AOI sets. All current AOI sets will be saved in the form of an XML file. To import a AOI sets saved in other projects, select **Areas of Interest→Import AOI sets from File,** and browse to the previously saved XML file. The AOI sets on the file will be imported and will replace any AOI sets already in the current project. Be sure to record the names and locations of these files or use names and locations that will be remembered.

A set of buttons at the top of the "Configure AOIs Graphically" window can be used to adjust image zoom setting, lock or unlock the window for further modifications of AOI modification, edit the attachment points, and save an image file of the current window in bmp, jpg, tiff, or png format. Hover the mouse over each button to see its function.

AOIs created graphically can be edited manually, as described in the next section. AOIs created manually, will be displayed on the Configure AOIs Graphically window when the appropriate "AOI set" is selected, and can be edited graphically

When finished creating and exporting AOI sets, click the "Save & Close" button at the upper right of the Configure AOIs Graphically window to save the contained AOI sets as part of the current project.

8.2 Defining Areas of Interest manually

To create AOIs manually, select **Configure→Areas of Interest (static, manually)** from the *ETAnalysis* main menu. A table entry dialog will appear as shown below.

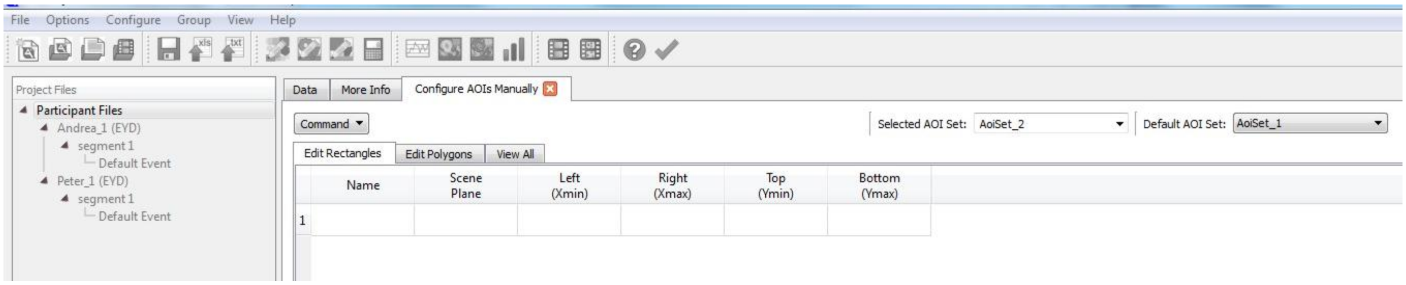

Select the Edit Rectangles tab for rectangular AOIs or the Edit Polygons tab for multi-sided AOIs. The View All tab shows coordinates for both rectangular and multi-sided AOIs, but this tab is only for viewing the coordinate values and cannot be used for editing.

Manual AOI creation or editing is usually done with rectangular areas. In this case, the left, right, top, and bottom boundary coordinates are specified. In the case of polygon AOIs, each horizontal and vertical vertex coordinate are specified. VnX is the horizontal coordinate for vertex n, and VnY is the vertical coordinate. If creating a new polygon AOI, only 6 vertices are available. Polygons with more than 6 sides, must be created graphically. However, if polygon with over 6 sides has been created graphically, all the vertices will listed and can be edited manually.

If creating a new AOI set, click the "Command" button and select "Create Blank Aoi set". Next to "Selected AOI Set" type in the desired name for the AOI set. In the appropriate columns type the AOI name, and boundary coordinates or vertex coordinates.

The boundary or vertex coordinates are floating point values. If the AOI set will be used for *ET3Space* data, be sure to include the scene plane number. The scene plane number must be an integer. As soon as an entry is made on one row, another row becomes available for new entry. To delete an AOI, highlight the row and select "Delete Selected AOIs" from the "Command" drop down menu. To edit an existing AOI set, select the set from the "Selected AOI Set" drop down menu. To close the dialog and save the current AOI sets click "Save & Close". To close the dialog without saving any changes made (since the dialog was opened), click the X in the red square next to the tab label.

Argus Eye Tracker gaze values have horizontal coordinates that increase as gaze moves to the right, and vertical coordinates that increase as gaze moves down. For rectangular areas therefore, on any given row, left boundary values must be less than right boundary values and top values must be less than bottom values. In the case of *ET3space* scene planes there may be occasionally be a scene plane for which it is not immediately obvious what is meant by "horizontal" and what is meant by "vertical". "Horizontal" always refers to the y axis and "vertical" to the z axis on *ET3Space* scene planes. See the *ET3Space* manual for a more detailed explanation of scene plane coordinate frames.

To copy the AOI set, or to save AOIs to a file or import previously saved AOIs from a file, click the "Command" pull down menu and select "Copy….", "Export AOIs…" or "Import AOIs…".

9 Event Correspondence with Background images and AOI sets

Once Backgrounds and static AOI sets have been defined (sections 7 and 8), each event in the project can be matched with a particular background and AOI set.

These correspondences are specified on a AOI Sets and/or Background Correspondences dialog. The dialog is available as a selection under the "Configure" menu, and also automatically appears when a Fixation Sequence computation is requested as described in the next section.

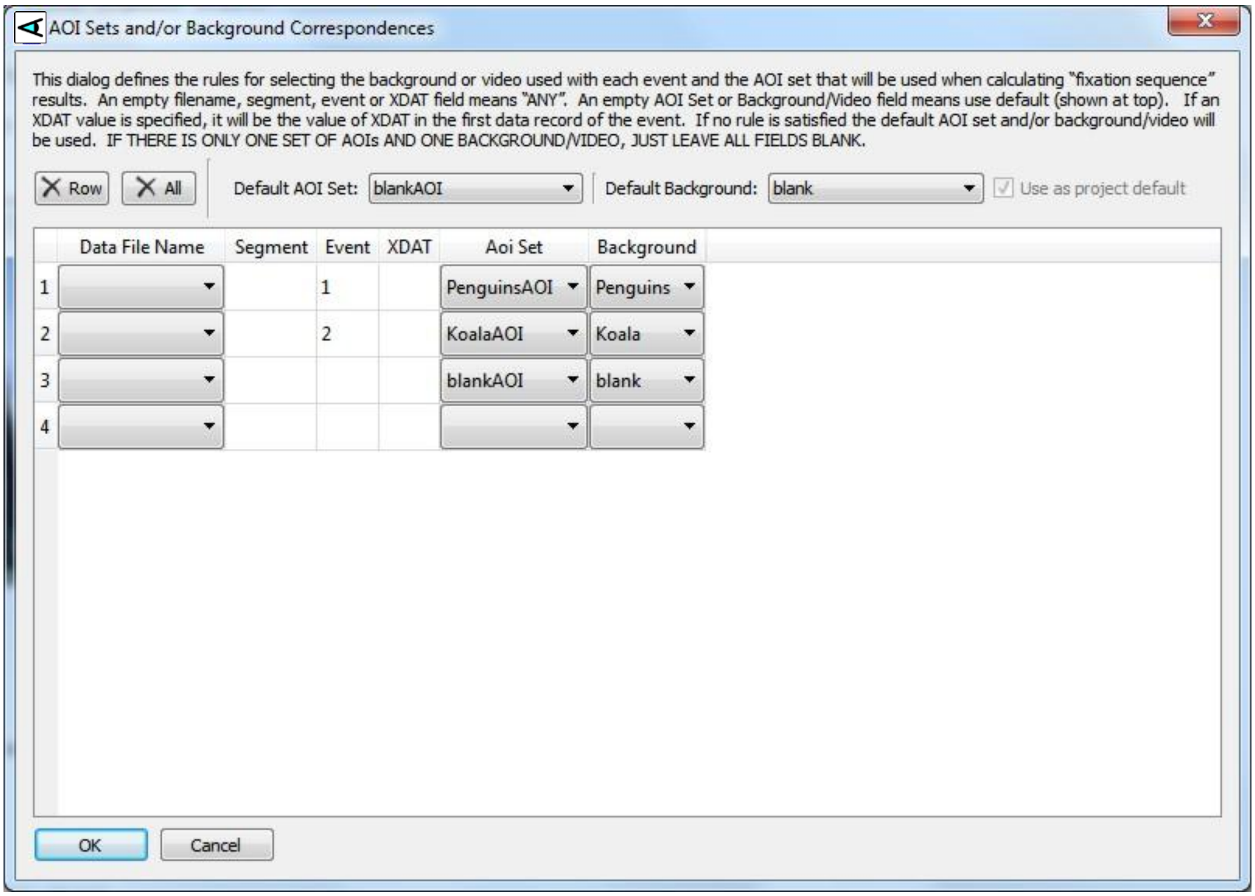

An AOI set can be explicitly assigned to each event, segment, or file. Alternately, the first XDAT value in each event can be used to specify the AOI set. In the example above, event 1, from any file and segment in the project, is assigned the "Penguin" background image and the "PenguinAOI" AOI set. Event 2, from any file and segment in the project, is assigned the "Koala" background image and the "KoalaAOI" AOI set. The "blank" background and aoi set are not assigned to any event.

The AOI Sets and/or Background Correspondences dialog defines the rules for selecting the background and AOI set that will be used with each event for calculating fixation sequence results. An empty field means "ANY". If an XDAT value is specified, it will be the value of XDAT in the first data record of the event. If no rule is satisfied a default assignment will be used.

There is an implicit logical AND between fields. For example, if both an XDAT value and an event number are specified on one row of the dialog, the AOI set specified on that row will be used only if that event number also has the specified XDAT value in its first field. Otherwise, a default assignment will be used.

If there is only one background and one set of AOIs, just leave all fields blank. Note that each Background and AOI set defined in the project must appear on the dialog, even if they will not correspond to any event.

In the example below, any event with an initial XDAT value of 1 will be associated with Penguins background and AOI set; while any event with an initial XDAT values of 2 will be associated with the Koala background and AOI set.

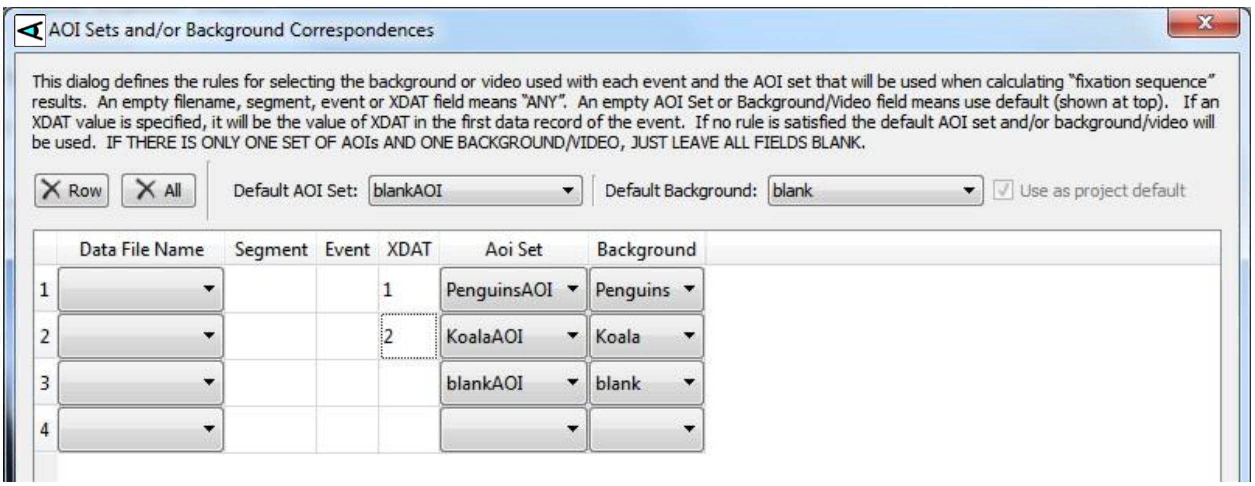

In the following example, Penguins will be associated with event 1 on the Andrea_1.eyd file, but with event 2 on the Peter_1.eyd file, etc.

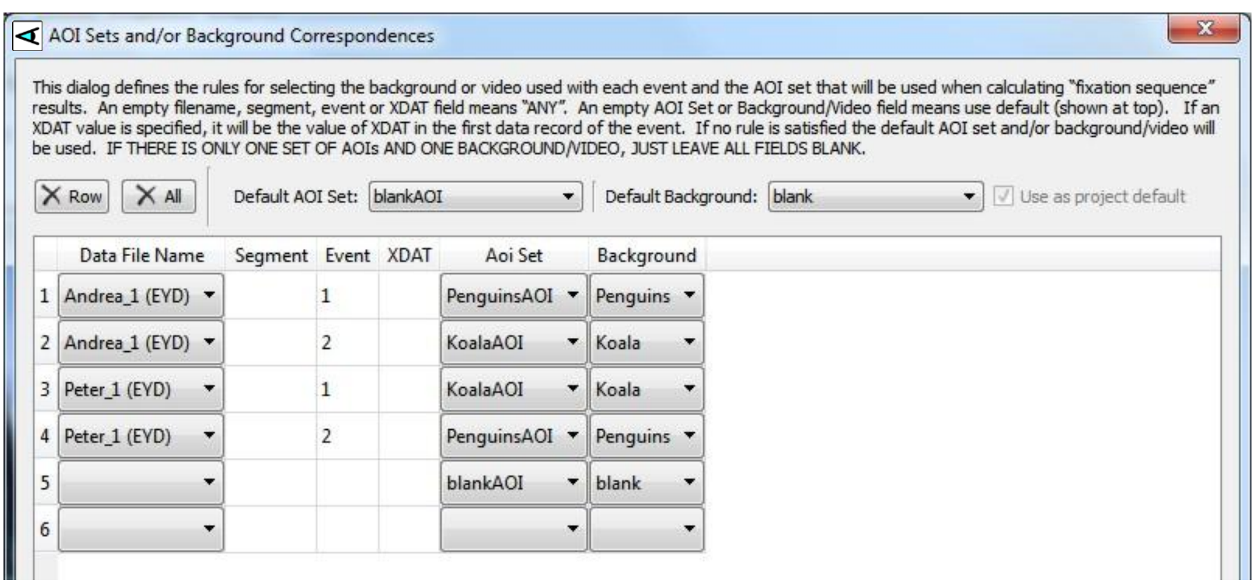

10 Fixation and Saccade analysis

During normal scanning of a visual scene, eye movement is characterized by a series of stops and very rapid jumps between stopping points. The stops, usually lasting more than 100 ms, are called fixations, and it is during these fixations that most visual information is acquired and processed. These rapid jumps between fixation points are called saccades. Saccades are conjugate eye movements (both eyes move together) that can range from 1 to 50 degrees of visual angle, and achieve velocities as high as 400-600 degrees per second. Very little visual information is acquired during saccades, primarily because of the very fast motion of the images across the retina, and an associated elevated visual threshold by the brain, just prior to and during a saccade, called visual image suppression.

The eyes are not completely stationary during fixations, but exhibit a variety of small involuntary motions, usually of less than one degree visual angle, called flicks (or micro saccades), drift, and tremor.

The eyes can smoothly track targets that are moving no more than about 30 deg/sec (faster for some people). These conjugate, slow tracking eye movements are usually called smooth pursuit and act to partially stabilize slowly moving targets on the retina. Similar slow conjugate eye movements called compensatory eye movements partially stabilize the visual field during either active or passive head or trunk motions.

10.1 Fixation Analysis

When gaze data is analyzed, a common practice is to first reduce the data to a set of fixations, since in most cases these are the periods when visual information was stable on the retina and available to be processed by the brain. The following subsections describe the algorithm used by *ETAnalysis* to compute periods of "fixation".

Video base eye trackers measure line of gaze with respect to the eye camera optics. In the case of systems with table mounted the data report gaze with respect to a stationary display surface (usually a display monitor located just above the eye camera optics). When the fixation algorithm is applied to this data, it finds periods of relatively stable gaze with respect to the stationary display. Note that sometimes, if the subject's head is moving, compensatory eye movements may produce these periods of stable gaze. However, the fixation algorithm does not make these distinctions. It just finds periods of stable gaze measurement.

ET3Space data (produced by an eye tracker with head mounted optics and an independent head tracking device) also specifies point of gaze on stationary surfaces, and when the fixation algorithm is applied to this data the situation is the same as that described in the previous paragraph.

In the case of head mounted optics used with just a head mounted scene camera (not using *ET3Space* feature) the gaze measurement is with respect to the subject's head (more specifically, with respect to the field of view of the head mounted scene camera). The fixation algorithm, when applied to this data finds periods of relatively stable line of gaze with respect to the head. Note that periods during which smooth pursuit or compensatory eye movements may be stabilizing gaze on some external target may not constitute periods of stable gaze data, since, during these periods, the eye is rotating with respect to the subject's head.

When *ET3Space* is used (see section 20), the fixations are computed with respect to the stationary surfaces defined as part of the *ET3space* environment. In this case it is possible for compensatory eye movements to stabilize gaze with respect to one of these surfaces even though the head is moving with respect to the surface. Such periods will be listed as fixations if gaze with respect to the surface is stable enough to meet fixation algorithm criteria.

In all cases described above, fixations are computed with respect to the scene image frame of reference, which may either be a stationary surface (E.g., *ET3Space*) or a head mounted scene camera image; and Areas of Interest on the scene do not need to be specified in order for fixations to be computed. Fixations computed in this way therefore appear as Fixation Nodes that are directly under an event node on the project tree diagram. This is the case addressed by the procedures described under the current section (section 10) of this manual.

As described later in this manual, *ETAnalysis* is also able to define "Moving Areas of Interest" which follow targets that move either with respect to a stationary display surface or with respect to a head mounted scene camera field of view. If moving areas of interest (MAOIs) are defined, then fixations may also be computed as periods during which the gaze remains relatively stable with respect to the boundaries of these areas. This may include periods during which smooth pursuit or compensatory eye movements stabilize gaze on a moving target. Note, however, that fixations defined in this way can only be computed if moving areas of interest have been defined, and can be detected only when gaze is within one of these areas. Such fixations appear as Fixation Nodes under Moving Area of Interest Nodes, and this type of analysis is discussed under section 16 of this manual. Note also that that if the moving area of interest boundaries do not accurately and stably track the edges of an image object, this can cause errors in fixation computations.

The same basic fixation algorithm, which is really an adjustable nonlinear filter, is used in all cases.

10.1.1 Origin of fixation algorithm

The fixation algorithm used in *ETAnalysis* derives from work done by Lambert, Monty, and Hall, (1974), further developed by Flagg (1977) and Karsh and Breitenbach (1983), and others over the years. The method falls in the category that Duchowski (2003) labels "dwell-time fixation detection" as opposed to "velocity-based saccade detection".

The original rational for a minimum fixation duration was that the latency in beginning a saccade to a new target that was probably a measure of the minimum time needed by the nervous system to process visual information meaningfully, and therefore the shortest sensible "snapshot". The shortest latencies were reported to be about 100ms, with latencies ranging up to about 300 ms (Alpern, 1969; Young, 1970, Yarbus, 1967). Looking at more recent data, saccadic latencies seem rarely to be less than 150 ms under most conditions, and are more typically over 200 ms, but "express saccades" can have latencies as short as 90 to 120 ms when the old fixation target disappears before the new target appears, or if the new targets are predictable (for example, see Darrien et.al. 2001, Fischer and Ramsperger, 1984).

The default "minimum" in the *ETAnalysis* fixation program is 100 ms. Note that if the data is collected at 180 fields per second, it corresponds to 19 samples. The 1-degree minimum change in gaze position required to define a new fixation is based, loosely, on the fact that miniature eye movements (tremor, drift, and micro-saccades) are generally smaller than 1 degree. Of course they are often significantly smaller than 1 degree, and the minimum could arguably be smaller. It is important to take into account the quality of the measurement. No matter what the underlying physiology, it is only possible to detect changes in fixation position that are larger than the measurement noise.

There is no firm definition for a fixation. It is less a physiological quantity than a method for categorizing sections of a data stream. Sensible selection of criteria depends on the experimental goal and the characteristics of the measurement a well as underlying physiology. There are quite a few different algorithms in the literature for detecting fixations, all of which represent logical strategies. Processing the same data with different algorithms or different parameters for a given algorithm, all of which may be justifiable, can easily result in a different number of fixations and different set of fixation start and stop times, and positions. This makes it important to report the method used.

The default parameters in the *ETAnalysis* fixation program are chosen with the hope that they will at least be reasonable in the majority of cases, but they were not intended to promote a particular definition for a fixation. The program is really an adjustable non-linear filter. It is useful to look at horizontal and vertical position time plots for some sections of data, and to superimpose the fixation start and stop points determined by the fixation program for that section of data. The researcher can verify that the program is doing a reasonable job of choosing periods that the researcher would regard as "fixations", or adjust the algorithm parameters as necessary.

References:

Alpern, M. Types of Movement, H. Davson (Ed.) The Eye (vol.3, 2nd ed), Academic Press, New York, 1969.

Young, L., Recording eye position, M. Clynes & M. Milsum (Eds.), Biomedical Engineering Systems, McGraw Hill, New York, 1970.

Yarbus, A.L., Eye Movements and Vision, Plenum Press, New York, 1967.

A.T. Duchowski, Eye Tracking Methodology Theory and Practice, Springer-Verlag, London, 2003.

Darrien, Herd, Starling , Rosenberg , and Morrison, An analysis of the dependence of saccadic latency on target position and target characteristics in human subjects, BMC Neuroscience 2001, 2:13

Fischer B, Ramsperger E: Human express saccades: extremely short reaction times of goal directed movements. Exp Brain Res 1984, 57:191-:195.

10.1.2 Fixation Algorithm Description

The fixation algorithm relies on three Criteria. The first is used to determine when a fixation starts; the second is used to determine whether subsequent data samples are part of the same fixation; and the third is used to determine which data samples should be averaged together to determine the final fixation coordinates.

To "start a fixation" the program looks for a specified period (Minimum Fixation Duration) during which gaze has a 95% confidence interval (twice the standard deviation) of no more than a specified amount (Threshold 1 value). The average horizontal and average vertical gaze position during this period is the "temporary fixation position". To end the fixation, it looks for a specified number of sequential gaze position samples, and also the average position of these samples, to be farther than a specified distance (Threshold 2) from the temporary fixation position. The final fixation position is the average position of all data samples between the beginning and end of the fixation. The exception is that any gaze coordinates that were farther than a specified value (Threshold 3) from the initial fixation position are not included in the average.

The reason that more than one sample must exceed threshold 2 in order to end a fixation is so that an extraneous spike in the data will not cause the fixation to end. The reason for Threshold 3, is to exclude such spikes from being included in the fixation position computation.

If a fixation has started and pupil recognition is lost (for example, due to a blink) for less than the time specified as "Maximum Pupil Loss", then this does not cause the fixation to end. The period of pupil loss is ignored, the gaze position on the first record for which the pupil is again recognized is compared to Threshold 2, and process continues as previously described. If pupil recognition is lost for a longer period, the fixation is considered to have ended at the beginning of the recognition loss period.

The initial default value for Minimum Fixation Duration and Threshold 1 are is 100 msec and 1 degree visual angle, respectively. The default for Threshold 2 is also 1 degree, and the default number of samples that must exceed Threshold 2 to end a fixation is 3. The default for Threshold 3 is 1.5 degrees. The default for Maximum Pupil Loss is 200 msec (the maximum duration of most blinks). All of these parameters can be adjusted by the user. The user can set adjusted default values for the project and can also adjust these parameters individually for every fixation set created.

When converting time periods to a number of samples, note that the period is the inter-sample period times one less than the number of samples. For example, the period between sample 1 and sample 7 is 6 sample periods. For a system operating at 60 Hz this is 100 msec.

10.1.3 Default Fixation Criteria

To view or adjust the current default values for fixation parameters, select "Fixations" from the Configure menu on the main menu bar. The screen shot below shows default values for *ETVision* data (180 Hz).

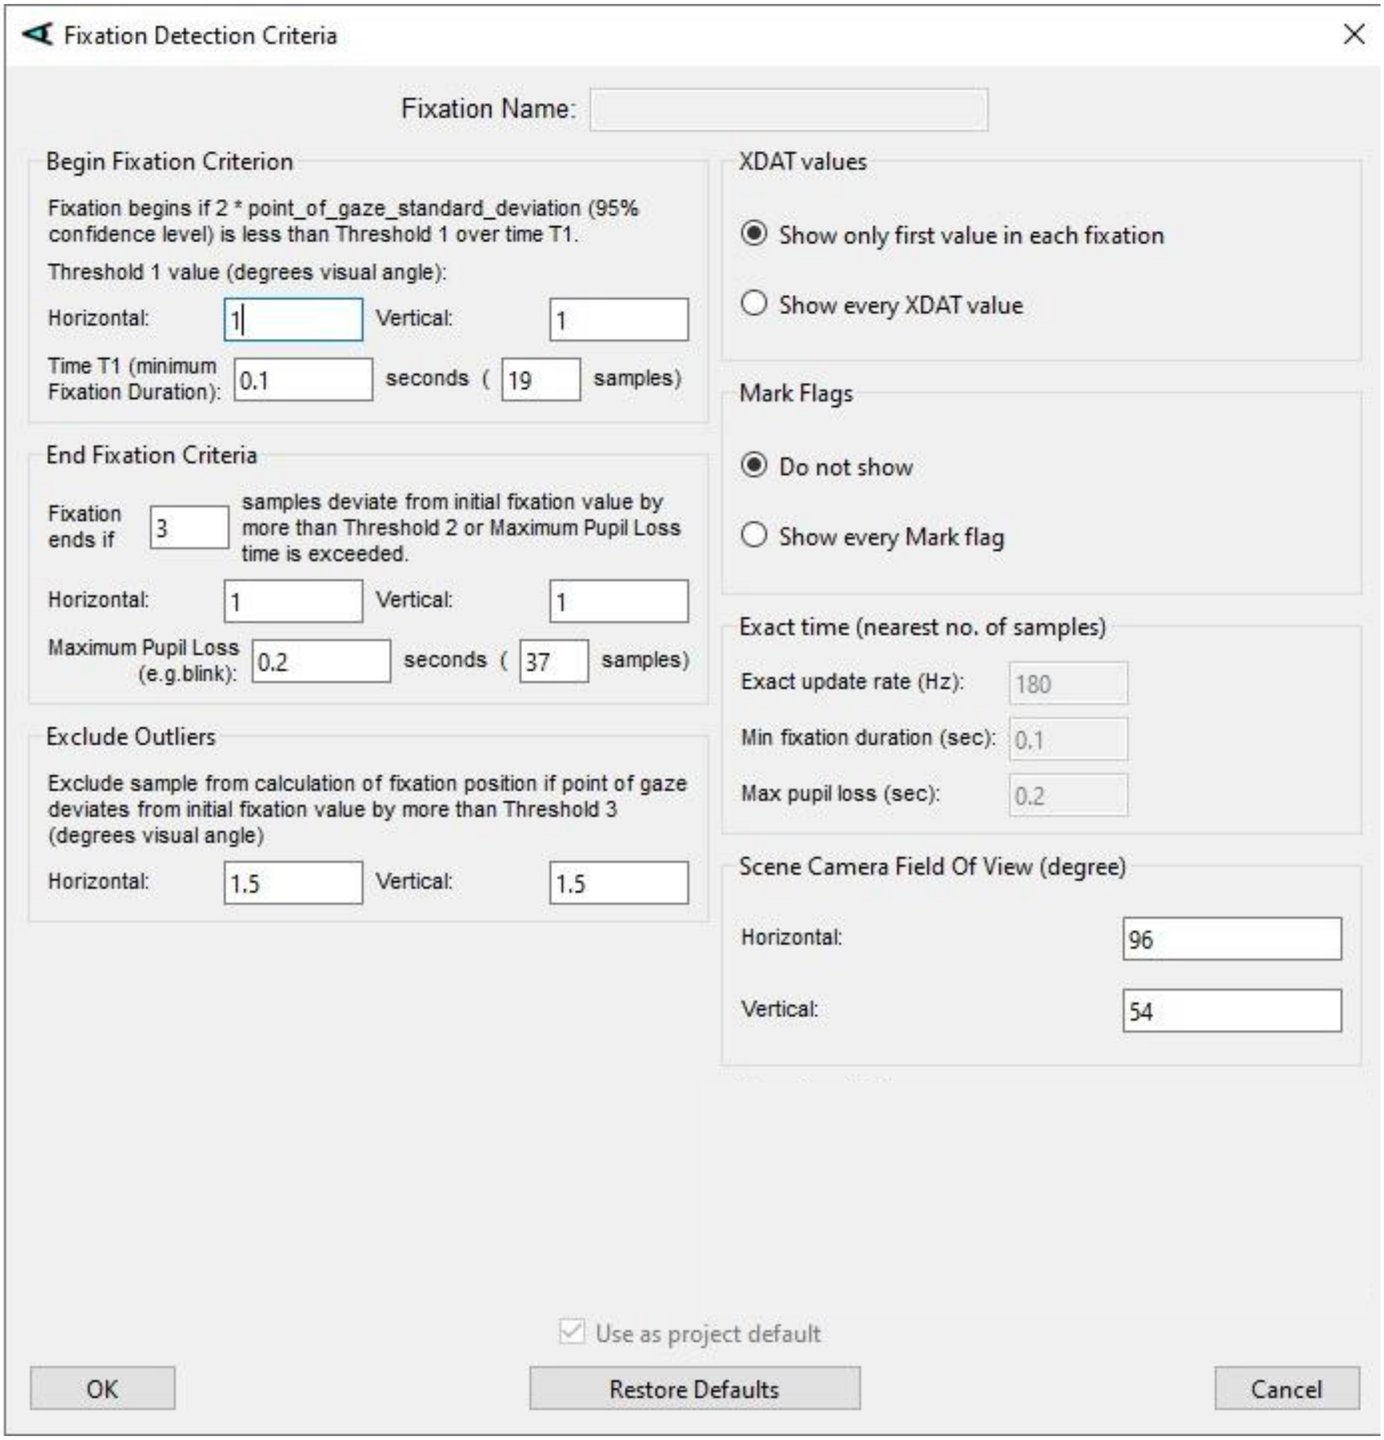

10.1.3.1 Begin Fixation Criteria

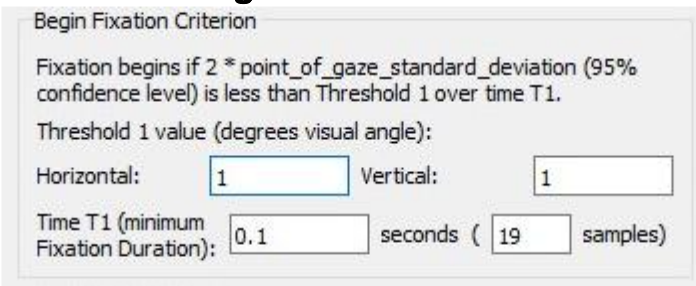

A fixation is considered to "start" when the gaze data is sufficiently stable for a minimum time. More specifically, a fixation "starts" when a minimum number of sequential horizontal and vertical point-of-gaze coordinate samples have a standard deviation below half of Threshold 1. The rational for specifying twice standard deviation as the threshold is that a normally distributed random variable will have about 95% of its values within the range of two standard deviations. The threshold value is expressed in units of degrees visual angle, and Threshold 1 has a default value of 1 degree. The minimum time (T1) is specified in seconds, and the program selects the number of samples that most closely corresponds to the specified time. The default value for T1 is 0.1 sec.

Note that the time interval covered by n samples is the sample period multiplied by n-1. At a 60 Hz update rate, for example, 0.1 sec corresponds to 7 samples; at 180 Hz, 0.1 sec corresponds to 19 samples. The user can change T1 and the program will recalculate the number of samples using the update rate that it reads from the file header.

Once the program finds the minimum number of sequential sample points (corresponding to time period T1) that have a small enough standard deviation, the fixation is considered to start with the first of these data samples, and the average point of gaze value for this set of points is memorized as the fixation start position.

Note that Time T1 is the minimum possible fixation duration.

10.1.3.2 End Fixation Criteria

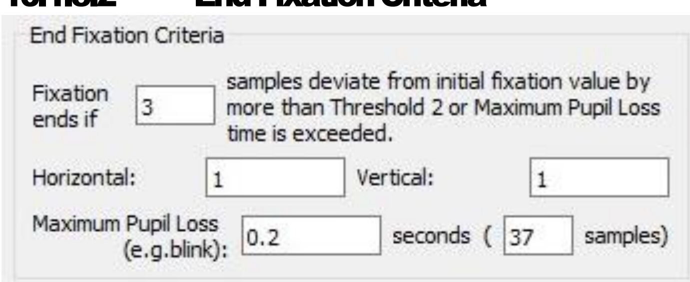

A fixation ends when several (default: three) sequential samples, as well as their average, deviate from the fixation start position by more than the Threshold 2 value (default: 1 degree visual angle). The deviation can be in either horizontal or vertical point of gaze coordinates. The data point preceding these samples is considered to be the last data sample in the fixation.

Another reason to end the fixation is continuous loss of eye recognition for more than a specified time (default: 0.2 sec.). Shorter losses are assumed to be blinks and do not cause the fixation to end. The program uses the update rate to calculate the number of samples most closely corresponding to the specified time.

10.1.3.3 Finding the final fixation position and excluding outliers

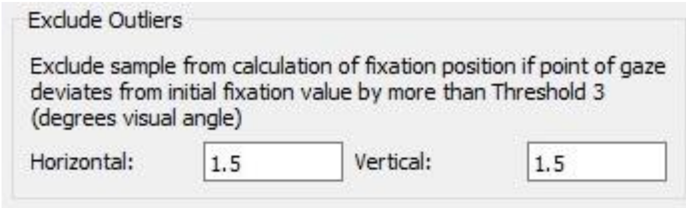

The final fixation position is the average of all the points from the start point to the end point, but excluding some points considered to be outliers. Remember that a single point that is very far from the fixation start point doesn't necessarily end the fixation. There must be 3 (or some other specified number) in a row. This is so that a brief measurement noise spike will not end the fixation. The result is that there may be some "far off" points between the fixation start and end points that should be considered noise. Any points farther than Threshold 3 from the fixation start position will be excluded from computation of the final fixation position. Threshold 3 is specified separately for the horizontal and vertical axis, and the default is 1.5 degrees visual angle.

10.1.3.4 XDAT and Mark Flags

Each fixation record shows the XDAT value that was recorded at the beginning of fixation. However, the XDAT value can change during the fixation period. If the user selects "Show every XDAT value" option, the program will create a line in the fixation list for every XDAT change as shown below:

Data | More Info

Fix#	StartTime	Duration	PupilLoss	StopTime	InterfixDur	InterfixPupilLoss	InterfixDegree	HorPos	VertPos	PupilDiam	Field	StartRecord#	StopRecord#	XDAT
1	0.000	0.200	0.000	0.200	0.000	0.000	0.000	200.0...	50.000	60.000	1	1	13	1
x	0.000													1
2	0.300	0.200	0.000	0.501	0.100	0.000	14.142	100.0...	150.000	40.000	19	19	31	0
x	0.300													1
x	0.400													0
3	0.601	0.200	0.000	0.801	0.100	0.000	14.142	200.0...	50.000	60.000	37	37	49	1
x	0.601													0
x	0.801													1

When the eye tracker is recording a file, the user can create "marks" on the data by pressing numeric keys on the keyboard. These flags can be displayed in the fixation list if user selects the option "Show every Mark Flag".

Data | More Info

Fix#	StartTime	Duration	PupilLoss	StopTime	InterfixDur	InterfixPupilLoss	InterfixDegree	HorPos	VertPos	PupilDiam	Field	StartRecord#	StopRecord#	XDAT	MarkFlag
1	0.000	0.200	0.000	0.200	0.000	0.000	0.000	200.0...	50.000	60.000	1	1	13	1	
2	0.300	0.200	0.000	0.501	0.100	0.000	14.142	100.0...	150.000	40.000	19	19	31	1	
m	0.400														0
3	0.601	0.200	0.000	0.801	0.100	0.000	14.142	200.0...	50.000	60.000	37	37	49	0	
m	0.801														1
4	0.901	0.200	0.000	1.101	0.100	0.000	14.142	100.0...	150.000	40.000	55	55	67	1	
5	1.201	0.200	0.000	1.401	0.100	0.000	14.142	200.0...	50.000	60.000	73	73	85	0	

The default for XDAT is “Show only first value in every fixation”, and the default for Mark Flags is “Do not show”.

10.1.3.5 Exact time

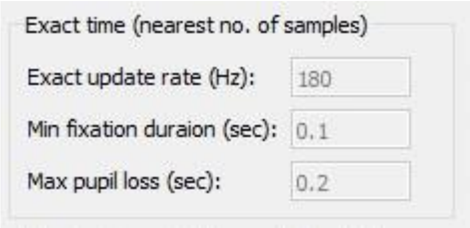

For some fixation criteria, the user specifies a time interval in terms of seconds, and the program calculates the corresponding number of samples that comes closest to this time interval, based on the eye tracker update rate. The actual time interval is based on this number of samples and can be slightly different from the value that user requested since it must round to the nearest whole sample value. The Exact time group of items shows the exact time intervals based on the number of samples and update rate. This is information only.

10.1.3.6 Scene Camera Field of View

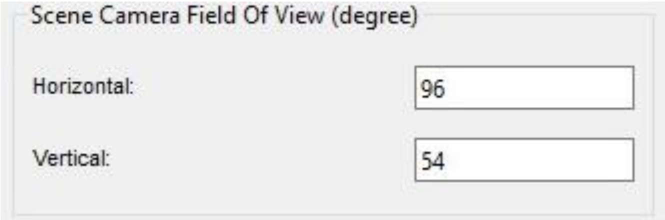

The field of view specification refers to the field of view of the eye tracker head mounted scene camera, and allows the program to accurately compute visual angles from the scene camera pixel positions specified by the gaze data (“horz_gaze_coordinate” and “vert_gaze_coordinate”).

For data collected with *ETVision* eye trackers the default values are “Horizontal: 96” degrees, and “Vertical: 54” degrees. <u>It is strongly recommended that the default values be used for this parameter</u> unless a non standard scene camera lens has been used. If unsure, please consult Argus Science.

Note that if *ET3Space* data is being used to compute fixations, the *ET3Space* data contains all the information needed for the program to calculate visual angles (gaze position on a surface and distance from the head to the surface) and the scene camera field of view specification is not required.

10.1.4 Creating Fixation sets

Fixation sets are computed for data contained in Events and appear as nodes on the *ETAnalysis* tree diagram under the corresponding event node, with a name specified by the user. Each fixation set is created with a user-defined set of parameters that define fixations. An event node can have more than one fixation set since the same data can be processed using different parameter values to define the fixations.

Right clicking an event node, or any node above the event level on the tree diagram produces a context menu that includes a **Find Fixations** item. Selecting **Find Fixations** will bring up a **Fixation Detection Criteria Dialog** used to specify the parameters that will "define" a fixation. Clicking **OK** to close this dialog causes fixation sets to be computed for all events in sub-nodes under the selected node. For example selecting **Find Fixations** by right clicking an event node will create just one fixation set using the data from that event. Right clicking the "Data Files" node and selecting **Find Fixations** will compute fixation sets for every event in the project, etc.

A Basic Fixation Criteria dialog will open.

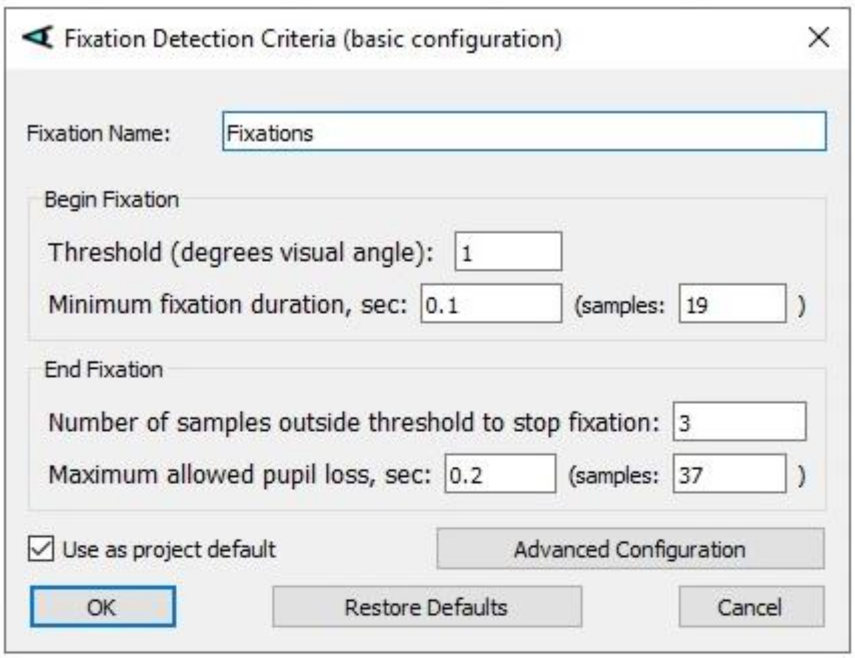

Type a name for the fixation set (or leave the default name shown). The name can also be changed later directly from the node on tree diagram. If default parameters have already been set as described in the previous section, and if these parameters are the ones desired, simply click OK. Basic parameters can be changed, if desired directly on this dialog. Checking the "Use as project default" box, will make these changes the default for any subsequent fixation sets created, but previously created fixation sets will not be changed.

To view or change all of the available parameters, click the "Advanced Configuration" button. This brings up the same dialog discussed in section 10.1.3. Changes made to the defaults values will apply only to the fixation sets currently being created, unless the "Use as project default" box is checked, in which case these changes will become the default for any subsequent fixation sets created.

10.1.5 Fixation data display

Clicking "OK" on the **Fixation Detection Criteria** window will cause a new "Fixation" node to appear on the project tree diagram. When a fixation node is selected on the tree diagram, in the left panel of the main window, the data tab on the right panel displays a list of fixation points.

Fix#	StartTime	Duration	PupilLoss	StopTime	InterfixDur	InterfixPupilLoss	InterfixDegree	ScenePlane	HorPos	VertPos	PupilDiam	GazeLength	StartField#	StopField#	CU_Field#	XDAT
1	0.000	0.534	0.250	0.534	0.000	0.000	0.000	0	0.490	-1.112	32.176	39.533	1	33	-1	0
2	0.584	1.418	0.000	2.002	0.050	0.017	0.508	0	0.449	-0.763	33.291	39.577	36	121	-1	0
3	2.069	0.117	0.000	2.186	0.067	0.000	18.445	0	-10.028	-9.586	31.125	42.558	125	132	-1	0
4	2.202	0.601	0.000	2.803	0.017	0.000	2.911	0	-11.840	-10.808	32.081	43.394	133	169	-1	0
5	2.853	0.184	0.000	3.036	0.050	0.000	17.057	0	1.081	-10.234	36.500	40.915	172	183	-1	0
6	3.053	0.717	0.000	3.770	0.017	0.000	1.147	0	0.266	-10.144	35.523	40.917	184	227	-1	0
7	3.820	0.434	0.000	4.254	0.050	0.000	16.563	0	12.592	-9.752	35.963	42.013	230	256	-1	0
8	4.404	0.701	0.050	5.105	0.150	0.100	31.695	0	-11.697	-0.787	32.100	41.842	265	307	-1	0
9	5.155	0.834	0.000	5.989	0.050	0.000	16.154	0	0.074	-0.653	34.706	39.435	310	360	-1	0
10	6.039	0.751	0.000	6.790	0.050	0.000	17.002	0	12.300	-0.877	34.587	40.548	363	408	-1	0

The table below explains the fields included for each fixation. All time intervals are shown in seconds

Name	Description
Fix#	Fixation number
StartTime	Time stamp of the first record in the fixation
Duration	Difference between stop and start time
PupilLoss	Total time during fixation when point of gaze was not available
StopTime	Time stamp of the last record in the fixation
IntefixDur	Start time minus stop time of previous fixation (zero for first fixation)
InterfixPupilLoss	Total time between fixations when point of gaze was not available (zero for first fixation)
InterfixDegree	Difference between this fixation and previous fixation in degrees visual angle (zero for first fixation). Calculation of InterfixDegree is explained below.
ScenePlane	Fixation scene plane number. ET3Space only
HorzPos	Average point of gaze horizontal coordinate during fixation
VertPos	Average point of gaze vertical coordinate during fixation
PupilDiam	Average pupil diameter during fixation
GazeLength	Average eye to scene distance during fixation. ET3Space only
StartField#	video_field_# of the first record in the fixation
StopField#	video_field_# of the last record in the fixation
CU_Field#	CU_video_field_num of the first record in the fixation (optional, may be missing in the data file)
XDAT	XDAT value of the first record in the fixation. Note: If user selects "Show every XDAT value" option, the fixation list will include a line for each new XDAT value. These lines will be marked "x" in the firsts column to distinguish them from the fixation records.
MarkFlag	Optional column. If user selects "Show every Mark flag" option, the fixation list will include a line for each Mark value. These lines will be marked "m" in the firsts column to distinguish them from the fixation records

The “More Info” Tab lists all of the information that was specified in the **Fixation Detection Criteria** window in order to create the fixation set, and also shows some summary information. The summary information items are:

- Event duration
- Number of fixations
- Average fixation duration
- Average inter-fixation duration
- Average inter-fixation degree
- Frequency (fixations per sec)
- Pupil loss time
    - Before first fixation
    - After last fixation
    - Total within fixations
    - Total between fixations
    - Total loss for event
- Loss due to overtimes

“Event duration” is the length of the data section defined as an “event”. This is the section of data over which the program tried to identify fixations. The first fixation may not have started until some time after the start of the “event” and the last fixation may have ended before the end of the event data.

“Average inter-fixation degree” is the distance between fixations expressed in degrees visual angle.

“Pupil loss time” is the time that a pupil was not recognized by the eye tracker. This includes blinks, lack of recognition due to poor discrimination, and loss of data fields as described below. Of course this cannot include periods during which the system “thought” it was recognizing a pupil, but was in fact mistakenly recognizing some artifact. The first 4 items under “Pupil loss time” should add up to the last item (“Total loss for event”).

The eye tracker may occasionally lose one or more fields of data. This should be a rare occurrence, but the system detects this when it occurs and reports, as “Loss due to overtimes”, the number of fields lost multiplied by the field period.

Other items in the summary list are self-explanatory.

10.1.6 Fixation overlay on data time plots

Data time plots created from a fixation node on the tree diagram will show a blue line to indicate periods identified as fixations. The best way to confirm that the fixation algorithm is classifying fixation periods as the user thinks it should is to examine time plots with the fixation detection overlay. For data gathered using *ET3Space* these should be plots of “ET3S_horiz_gaze_coordinate” and “ET3S_vert_gaze_coordinate”, since these were the data used to compute fixations. For data gathered not using *ET3Space*, the values to plot are "horiz_gaze_coordinate" and "vert_gaze_coordinate".

To generate a time plot with fixation overlay, right click the saccade node on the tree diagram, and select **Display Time Plot**. On the **Configure Time Plot** dialog, select the items to be plotted as discussed in the previous paragraph. It is also suggested that the "Manual Range" check boxes be checked, and that both ranges be manually set to Min = 0 and Max = 1280. This will make the height of time plot vertical axis correspond to the *ETVision* scene camera horizontal field of view. Click **Show Plot** to make the plot appear. Use the "+" button or the mouse wheel to expand the time scale until fixation periods can be clearly distinguished and evaluated. Use the arrow buttons or "drag" with the mouse to move along the time scale. Select "Points and Lines" to see individual data points as well as connecting lines.

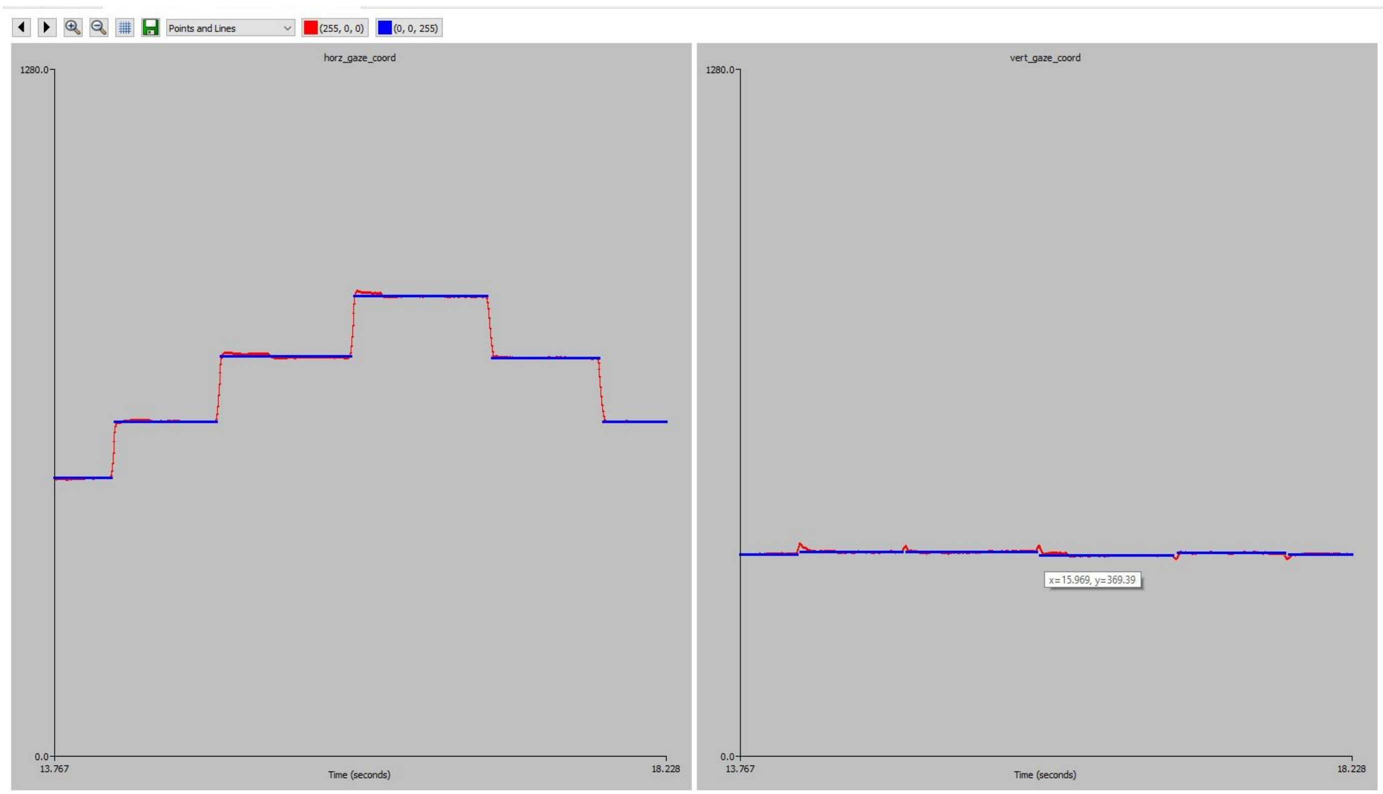

Section 14.1 has a more general description of time plot features and adjustments.

10.2 Saccade Analysis

(Please note that the saccade analysis function is intended primarily for use with data gathered using the Argus Science *ETVision* system, rather than older system types.)

As discussed at the beginning of section 10, saccades are rapid eye movement jumps that typically separate the relatively stable gaze "fixation" positions. These jumps are ballistic in nature and can reach peak velocities on the order of 600 deg/sec. The relationships between amplitude (amount of rotation from beginning to end), peak rotation velocity, and duration usually fall within a typical range referred to in the literature as the "main sequence".

In *ETAnalysis*, saccade analysis is always initiated from a Fixation node, and examines data during periods between the periods identified as fixations. The algorithm does not attempt to identify "micro-saccades" that may occur during fixation.

Since saccades are characterized by rapid "ballistic" rotations of the eyeballs in their sockets the saccade analysis is always based on the data specifying gaze with respect to the head ("horz_gaze_coordinate" and "vert_gaze_coordinate"). Even if fixations have been computed based on gaze relative to an environment surface ("ET3S_horz_gaze_coordinate" and "ET3S_vert_gaze_coordinate") or relative to the position of a Moving Area of Interest (as described in section 16.8), the periods between these fixations are still classified as saccades based on gaze movement with respect to the head ("horz_gaze_coordinate" and "vert_gaze_coordinate").

The saccade analysis function examines data between identified fixation periods to find periods during which gaze movement and its velocity profile are consistent with a saccade, as opposed to measurement noise, or the slower smooth pursuit or compensatory movements used to stabilize an image on the retina.

For a period of data to be classified as a saccade by *ETAnalysis*, gaze movement must remain above a threshold velocity in a consistent direction. Average and peak velocities, as well as amplitude and duration, must fall within specified ranges. Many of these parameters are adjustable in *ETAnalysis*, with suggested values as defaults.

As in the case of fixations, the saccade algorithm provided by *ETAnalysis* is a user adjustable filter. Although default values are provided, it is up to the user to adjust the algorithm to best match periods that the user would like to classify as "saccades". *ETAnalysis* provides a means to examine time plots of gaze data with superimposed graphics to show periods classified as saccades by the algorithm. This is described in section 10.2.4.

10.2.1 Creating Saccade sets

Before selecting saccade analysis it will be necessary to compute fixations with one of the methods described in section 10.1 or section 16.8. Right click the resulting Fixation node and select **Find Saccades** from the drop down context menu . A *Saccade Detection Criteria* dialog will appear. Use the "Restore Defaults" button to use the "factory default" criteria or set individual parameters as desired. The adjustable parameters are described in the following section. Once desired parameters are in place, click OK to proceed with saccade detection. This will result in a new Saccade node under the fixation node in the tree diagram.

Saccade sets are computed for data contained in the selected fixation node and appear as nodes on the *ETAnalysis* tree diagram under the corresponding fixation node, with a name specified by the user. Each saccade set is created with a user-defined set of parameters used by the saccade detection algorithm. A fixation node can have more than one saccade set since the same data can be processed using different parameter values to identify the saccades.

If the “Use as project default” box is checked when the OK button is clicked, the saccade parameters currently in place will become “project default”, and will be the parameters shown next time the *Saccade Detection Criteria* dialog appears. Use the “Restore Defaults” button to return to the “factory default” parameters.

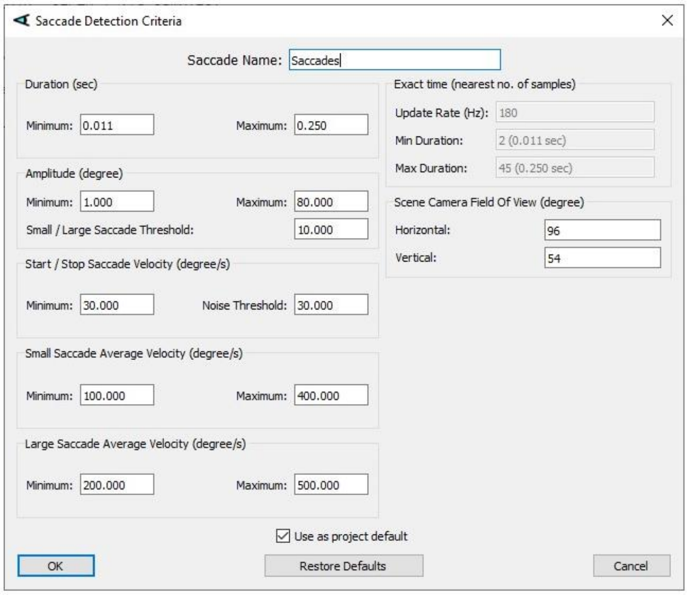

10.2.2 Saccade Detection Criteria

In this section the word “default” is used to refer to the “factory default” values. If these values have been changed for a particular project, “factory” default values can be restored by clicking the “Restore Defaults” button.

10.2.2.1 Saccade Name

This will be the name of the created saccade node. The default name can be replaced with a name of the user’s choice.

10.2.2.2 Duration

The default limits for saccade duration are minimum = 0.011 sec and maximum = 0.25 sec. This minimum represents only 3 data points (2 update intervals) at 180 Hz, and evaluation of shorter intervals is probably not reasonable for 180 Hz data. The type of rapid ballistic eye rotation implied by “saccade” classification cannot realistically last longer than 0.25 sec., and will virtually always be shorter than this. The default value is intended as an “outer bound”.

10.2.2.3 Amplitude

The default minimum and maximum values for saccade amplitude are 1 degree and 80 degrees, respectively. Although saccades can be smaller than 1 degree, fixations are often defined such that movements up to 1 degree are considered part of a fixation. Saccades smaller than 1 degree may also be increasingly hard to distinguish from measurement noise. Saccade amplitudes usually do not exceed 50 degrees visual angle. The default value is set to 80 degrees as an outer bound.

The range of expected saccade average velocities is a function of saccade amplitude. Average velocities tend to be higher for larger saccades. Rather than attempt to define a continuous function the algorithm divides saccade amplitudes into “small” and “large” saccade classifications. The dividing line is somewhat arbitrary, but the “Small/Large Saccade Threshold” default value of 10 degrees, coupled with the default values for minimum and maximum average velocities, described further below, is consistent with data in the literature and seems to produce correct saccade classification results.

10.2.2.4 Start/Stop Saccade Velocity

The “Minimum” value describes a velocity that must be exceeded before a gaze movement will be considered a saccade. If velocity falls below the “Minimum” for more than 3 sample periods (we want to allow occasional “noisy” samples) the saccade will be considered to have ended. (Note that other criteria may also cause the saccade to “end” or may prevent a period of data from being classified as a saccade).

Once the “Minimum” is exceeded along either axis, a velocity in the opposite direction along that axis of greater than “Noise Threshold” will be considered as most likely measurement noise rather than a continuation of a saccade.

Smooth pursuit and compensatory movements used to stabilize images on the retina are usually smaller than 30 deg/sec, and this is the default value for “Minimum”. The same value is the default for “Noise Threshold”.

10.2.2.5 Saccade Average Velocity

These parameters describe the range of average velocities for a period that can be classified as saccade

A period of data will be classified as a saccade only if the conditions previously described are met and the average velocity for the period also falls between the Average Velocity “Minimum” and “Maximum” values. The average velocity range is specified separately for “Small” and “Large” saccades as determined by the “Small/Large Saccade Threshold” previously discussed. The default values specify an average velocity range of 100-400 deg/sec for “Small” saccades (by default, under 10 degrees amplitude); and 200-500 deg/sec for larger saccades.

10.2.2.6 Exact time

For some saccade criteria, the user specifies a time interval in terms of seconds, and the program calculates the corresponding number of sample periods that comes closest to this time interval, based on the eye tracker update rate. The actual time interval is based on this number of sample periods and can be slightly different from the value that user requested since it must round to the nearest whole sample period value. The "Exact Time" group shows the exact time intervals based on the number of samples and update rate. This is information only.

10.2.2.7 Scene Camera Field of View

The field of view specification refers to the field of view of the eye tracker head mounted scene camera, and allows the program to accurately compute visual angles from the scene camera pixel positions specified by the gaze data ("horz_gaze_coordinate" and "vert_gaze_coordinate").

For data collected with *ETVision* eye trackers the default values are "Horizontal: 96" degrees, and "Vertical: 54" degrees. It is strongly recommended that the default values be used for this parameter unless a non standard scene camera lens has been used. If unsure, please consult Argus Science.

10.2.2.8 Use as project default

If the "Use as project default" box is checked when the OK button is clicked, the saccade parameters currently in place will become the project default, and will be the parameters shown next time the *Saccade Detection Criteria* dialog appears. Use the "Restore Defaults" button to return to the "factory default" parameters.

10.2.3 Saccade Data Display

Clicking "OK" on the **Saccade Detection Criteria** dialog will cause a new "Saccade" node to appear on the project tree diagram. When a saccade node is selected on the tree diagram, in the left panel of the main window, the data tab on the right panel displays a list of detected saccades.

Sac#	StartTime	Duration	PeakTime	PupilLoss	StopTime	Amplitude	AverageVelocity	PeakVelocity	DirectionH	DirectionV	StartSample#	StopSample#
1	0.244	0.028	0.023	0.000	0.272	10.391	311.744	456.465	-0.951	-0.309	45	50
2	0.983	0.039	0.023	0.000	1.022	17.583	395.608	577.939	0.898	0.441	178	185
3	1.322	0.028	0.011	0.000	1.350	9.990	299.709	436.378	-0.538	0.843	239	244
4	1.617	0.022	0.011	0.000	1.639	3.549	127.765	207.822	-0.645	-0.764	292	296
5	1.783	0.050	0.023	0.000	1.833	14.958	269.241	491.134	-0.241	-0.970	322	331
6	1.900	0.028	0.028	0.000	1.928	5.945	178.356	314.308	-0.989	0.145	343	348
7	2.233	0.011	0.011	0.000	2.244	1.814	108.854	188.622	-0.218	0.976	403	405
8	2.600	0.050	0.022	0.000	2.650	24.902	448.237	651.798	0.970	-0.242	469	478
9	3.000	0.044	0.033	0.000	3.044	15.560	311.194	569.236	-0.974	-0.228	541	549
10	3.750	0.039	0.022	0.000	3.789	16.037	360.835	530.691	0.977	0.215	676	683
11	4.900	0.033	0.022	0.000	4.933	8.946	230.046	356.765	-0.986	0.168	883	889

The table below explains the fields included for each saccade.

Name	Description
Sac#	Saccade number (saccade 1 is the first saccade identified in the event containing the saccade node)
StartTime	Time stamp of the first data record in the period identified as a saccade
Duration	Difference between saccade Start and Stop times
Peak Time	Time between Start Time and data record with Peak Velocity
Pupil Loss	Time within the detected saccade during which the left pupil was not recognized
StopTime	Time stamp of the last record in the period identified as a saccade
Amplitude	Distance (degrees visual angle) between the gaze positions at Start and Stop times
AverageVelocity	Amplitude divided by Duration (deg/sec)
PeakVelocity	Largest gaze velocity (deg/sec) reached during saccade
Direction H	Horizontal component of a unit vector specifying average saccade direction

	(direction of line connecting gaze position at Start and Stop Times)
Direction V	Vertical component of a unit vector specifying average saccade direction (direction of line connecting gaze position at Start and Stop Times)
StartSample#	Number of the data record corresponding to Start Time (data record 1 is the first data record in the event containing the saccade node)
StopSample#	Number of the data record corresponding to Stop Time (data record 1 is the first data record in the event containing the saccade node)

Gaze velocities for data record *n* are computed for each axis as:

Velocity(*n)* = (gaze_position (*n* +1) – gaze_position (*n*)) / sample_period ;

In the above equation, gaze_position is the visual angle from the center of the scene camera field of view to the position represented by the data sample, along the axis being considered.

Total (2 dimensional) velocity is computed as:

total_velocity = SQRT(horiz_velocity^2 + vert_velocity^2);

Reported Amplitude, AverageVelocity, and Peak Velocity are all total values rather than the values measured separately along horizontal and vertical axes.

The “More Info” Tab lists all of the information that was specified in the **Saccade Detection Criteria** dialog, and also shows the 2 summary information items listed below.

- Event duration
- Number of saccades

“Event duration” is the length of the data section defined as an “event”, and is the section of data over which the program tried to identify saccades between fixation periods.

10.2.4 Saccade overlay on data time plots

Data time plots created from a saccade node on the tree diagram will show a yellow line to indicate periods classified as saccades as well as blue lines to indicate periods identified as fixations. The best way to confirm that the saccade algorithm is classifying saccade periods in the way desired is to examine time plots of “horz_gaze_coordinate” and “vert_gaze_coordinate” data with the saccade detection overlay.

To generate a time plot of “horz_gaze_coordinate” and “vert_gaze_coordinate” data with saccade overlay, right click the saccade node on the tree diagram, and select **Display Time Plot**. On the **Configure Time Plot** dialog, select “horz_gaze_coordinate” and “vert_gaze_coordinate” as the first two data items. It is also suggested that the “Manual Range” check boxes be checked, and that both ranges be manually set to Min = 0 and Max = 1280. This will make the height of time plot vertical axis correspond to the *ETVision* scene camera horizontal field of view (about 96 degrees visual angle). Click **Show Plot** to make the plot appear. Use the + button or the mouse wheel to expand the time scale until saccade periods can be clearly distinguished and evaluated. Use the arrow buttons or “drag” with the mouse to move along the time scale.

Section 14.1 has a more general description of time plot features and adjustments.

On time plots of "horz_gaze_coordinate" and "vert_gaze_coordinate" data created from saccade node, the yellow saccade overlay will be a straight line from the first data value in the saccade period to the last data value in the saccade period. Note that in the case of "horz_gaze_coordinate" data plots, the yellow line slope is the average horizontal velocity of the detected saccade, and for "vert_gaze_coordinate" data the yellow line slope is the average vertical velocity.

On time plots of any other data items, created from the saccade node, the saccade overlay is a yellow horizontal line with a length corresponding to the saccade period, and with a vertical position at the median value of the plotted data item during the saccade period. This is intended only to show the period classified as a saccade.

The fixation overlays are horizontal blue lines with a length corresponding to the fixation period. On time plots of "horz_gaze_coordinate" and "vert_gaze_coordinate", and if the fixation computations were based on "horz_gaze_coordinate" and "vert_gaze_coordinate" data (rather than *ET3Space* data or moving area of interest positions), then the vertical position of the blue fixation overlay shows the "fixation position". On time plots of any other data item the vertical position of the blue line is the median value of the plotted data item during the fixation period, and is intended only to show the period defined as a fixation.

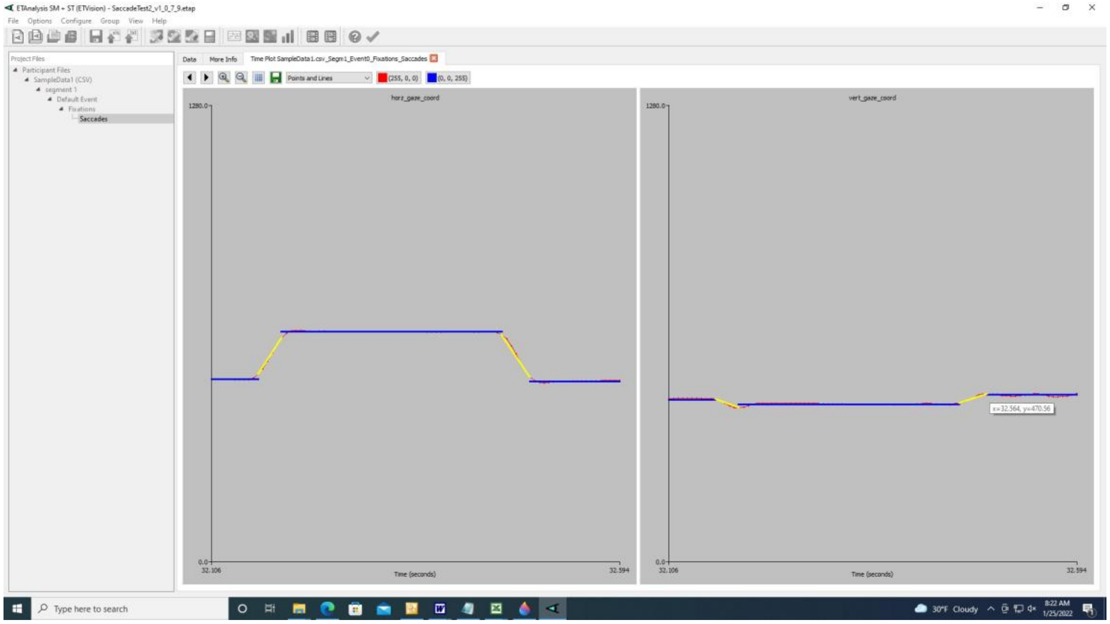

Time plot of *ETVision* "horz_gaze_coordinate" and "vert_gaze_coordinate" data from saccade node, showing fixations (blue overlay) separated by mostly horizontal saccades (yellow overlay). Time scale has been expanded and range has been manually set to 0-1280 pixels, corresponding to about 96 degrees visual angle

11 Fixation Sequence analysis with static AOIs

The Fixation Sequence Analysis compares a Fixations list to a set of Areas of Interest (AOIs). It reports the relation of each fixation to the defined AOIs, and computes related statistics. A similar analysis is done for "Dwells", which are defined as sequential fixations that remain in the same AOI, and are discussed in section 12. AOIs and fixations are discussed in sections 8 and 10, respectively.

This section discusses static Areas of Interest (sets of areas which retain constant shapes and positions for an entire event). It is also possible to define moving areas of interest, on a video showing a dynamic scene presentation. Working with scene video, defining moving areas of interest, and doing Fixation Sequence Analysis with moving areas, is discussed in section 16.

To launch the fixation sequence and dwell computations right click any node in the tree diagram at the "Fixation" level or higher, and select **Find Fixation Sequence (Static AOIs)**. The Fixation Sequence computation requested will apply to all sub-nodes under the selected node.

An AOI Sets and/or Background Correspondences dialog will appear, and its use was previously described in section 9.

When the "OK" button is clicked on this next dialog, the fixation and dwell sequence computations are performed and a fixation sequence node and dwell node are added to the project tree under each Fixation node selected as described above.

The Dwell node is at the same level as the fixation sequence node, and is automatically created whenever fixation sequence computations are done. Dwell analysis is described in section 12. Fixation Sequence nodes and Dwell nodes always include sub-nodes with an "AOI Summary" table, a "Transition table", a "Conditional Probability" table, and a "Joint Probability" table.

Note: when parsing data into events (as described in section 6), be sure that a new event starts at every point in the data where the scene image changed and will be represented by a different AOI set.

11.1 Fixation Sequence Data list and Info tab

Highlighting a Fixation sequence node on the project tree, in the main window left panel, displays a corresponding data list in the right panel.

Data | More Info

Fix#	AOI#	AOIName	StartTime	Duration	PupilLoss	StopTime	InterfixDur	InterfixPupilLoss	InterfixDegree	HorzPos	VertPos	PupilDiam	StartField#	StopField#	CU_Field#	XDAT
1	0	OUTSIDE	4.288	0.601	0.000	4.888	0.000	0.000	0.000	115.851	121.681	39.378	258	294	-1	1
2	4	Duck2_...	4.988	0.434	0.000	5.422	0.100	0.017	5.324	149.048	163.304	31.111	300	326	-1	0
3	4	Duck2_...	5.439	0.567	0.000	6.006	0.017	0.000	1.945	129.934	159.683	33.371	327	361	-1	0
4	2	Duck2_...	6.056	0.484	0.000	6.540	0.050	0.000	6.352	171.080	111.293	34.000	364	393	-1	0
5	2	Duck2_...	6.557	1.235	0.000	7.791	0.017	0.000	0.704	174.375	105.071	34.120	394	468	-1	0
6	2	Duck2_...	7.808	0.284	0.000	8.091	0.017	0.000	0.488	173.828	100.222	36.722	469	486	-1	0

The Fixation Sequence List is the same as the Fixation list except that AOI number and name designations for each fixation are added. In each case the program has determined that the coordinates for that fixation are within the boundaries of the specified AOI. Fixations that are not inside any defined AOI are considered to be within AOI #0, named "OUTSIDE" (all defined AOIs are numbered starting with 1).

The "More Info" tab displays the rules used for selecting the AOI set (as described in the previous section), the AOI set selected, and a list of the areas and boundary coordinates that make up that AOI set. It also contains information about the first three AOIs viewed.

The Fixation Sequence node on the project tree expands to show AOI summary, transition table, conditional probability table, and joint probability table sub-nodes.

11.2 AOI Summary

Highlighting an AOI Summary node on the project tree displays an AOI summary table in the right pane of the main window.

Data | More Info

AOI#	AOIName	FixCount	%FixCount	TotalFixDur	%TotalFixDur	AvgFixDuration	AvgInterFixDur	AvgInterFixDeg	AvgPupilDiam
0	OUTSIDE	9	40.909	4.905	43.235	0.545	0.213	3.671	37.151
1	Duck1_...	0	0.000	0.000	0.000	0.000	0.000	0.000	0.000
2	Duck2_...	6	27.273	3.270	28.824	0.545	0.017	0.916	34.860
3	Duck1_...	2	9.091	1.451	12.794	0.726	0.017	0.991	37.307
4	Duck2_...	5	22.727	1.718	15.147	0.344	0.017	1.694	34.347

The AOI summary includes the following fields that are calculated for each AOI. All time intervals are given in seconds.

Name	Description
AOI#	AOI number. Zero represents the area outside all AOIs
AOIname	AOI name (from AOI properties)
ScenePlane	AOI scene plane; *will be displayed only for EHD data files*
FixCount	Number of fixations that were inside AOI
%FixCount	Above value as a percentage of the total fixation count
TotalFixDuration	Total duration of the fixations inside AOI
%TotalFixDuration	Above value as a percentage of the total duration of all fixations
AvgFixDuration	Average duration of a fixation inside the AOI
AvgInterFixDuration	Average time between the *fixations that stayed inside the same AOI*
AvgInterFixDegree	Average angle between the fixations that stayed inside the same AOI (in degrees)
AvgPupilDiam	Average pupil diameter for the fixations inside the AOI. The average pupil diameter of each fixation is weighted with the fixation duration: AvgPupilDiam = SUM(fixation.avgPupilDiameter * fixation.duration) / SUM(fixation.duration)
FirstFixTime	Start time of the first fixation inside AOI. If AOI contains no fixations, value will be set to –1
XDAT	XDAT value at the start of the first fixation inside AOI

Note that if there are no overlapping AOIs, the %FixCount and %TotalFixDur columns should add up to 100%. If there are overlapping AOIs, the same fixation may sometimes be in multiple areas, and these columns may sum to greater than 100.

It is also possible to display a table showing the summary fixation sequence data corresponding to a single AOI, but for multiple events. From the main menu, under Group, select "Group Aoi Data". From the resulting pop up dialog, select the appropriate AOI set and click OK. From the next pop up dialog, select an AOI and click OK. The next pop up dialog will be a selection chart showing all the fixation sequence nodes in the project. Check the desired event nodes and click OK. A new tab will appear on the right pane labeled "Group AOI Data (*name*)", where *name* is the selected AOI.

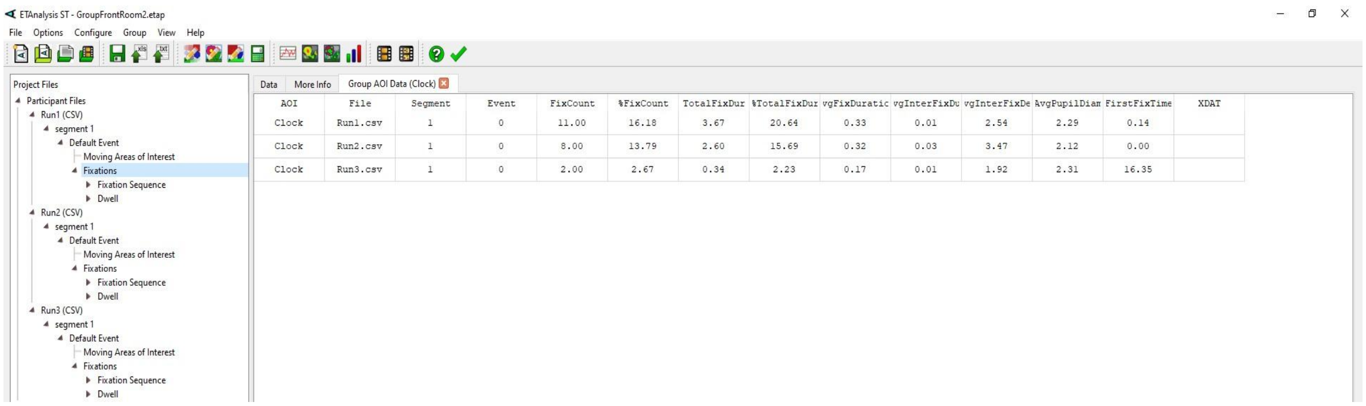

AOI	File	Segment	Event	FixCount	%FixCount	TotalFixDur	%TotalFixDur	vgFixDuratic	vgInterFixDu	vgInterFixDe	AvgPupilDiam	FirstFixTime	XDAT
Clock	Run1.csv	1	0	11.00	16.18	3.67	20.64	0.33	0.01	2.54	2.29	0.14	
Clock	Run2.csv	1	0	8.00	13.79	2.60	15.69	0.32	0.03	3.47	2.12	0.00	
Clock	Run3.csv	1	0	2.00	2.67	0.34	2.23	0.17	0.01	1.92	2.31	16.35	

The example, above shows summary data for an AOI named "Clock" from events on 3 different files (Run1.csv, Run2.csv, and Run3.csv).

11.3 Transition Table

Highlight a "Transition Table" node on the project tree to see the number of transitions between fixations on any two AOIs.

Data | More Info

AOI Name	OUTSIDE	Duck1_Head	Duck2_Head	Duck1_Feet	Duck2_Feet
OUTSIDE	4	0	2	0	2
Duck1_Head	0	0	0	0	0
Duck2_Head	2	0	3	0	1
Duck1_Feet	1	0	0	1	0
Duck2_Feet	1	0	1	1	2

Each cell in the Transition Table shows the number of transitions from a fixation in the AOI represented by the row number to the AOI represented by the column number. For example, the cell on row labeled "Duck2_Feet" and column labeled "Duck1_Feet" shows how many times a fixation in the "Duck2_Feet" area was directly followed by a fixation in the "Duck1_Feet" area (reminder: first row and column represent all parts of the scene not within any defined area).

The table contains non-negative integer values. If no AOIs overlap, the sum of all entrees should be equal to one less than the total number of fixations (total_number_of_fixations – 1). If there are overlapping AOIs than a fixation may sometimes be in more than one area and will appear in the table more than once. In this case the total of all entrees may be greater.

11.4 Conditional Probability Table

Highlight a "Conditional Probability" node to see the probability that a fixation in a given AOI would transition to a fixation in any other AOI.

AOI Name	OUTSIDE	Duck1_Head	Duck2_Head	Duck1_Feet	Duck2_Feet
OUTSIDE	0.500	0.000	0.250	0.000	0.250
Duck1_Head	0.000	0.000	0.000	0.000	0.000
Duck2_Head	0.333	0.000	0.500	0.000	0.167
Duck1_Feet	0.500	0.000	0.000	0.500	0.000
Duck2_Feet	0.200	0.000	0.200	0.200	0.400

The entry on row *n*, column *m*, shows the conditional probability that if a fixation was in *AOI n*, the next fixation was in *AOI m*. The value is calculated as the corresponding entry in the Transition Table divided by the total number of transitions from the AOI defined by the row. The number of transitions from the AOI is equal to the number of fixations inside this AOI with one exception: the last fixation does not transition anywhere. Therefore the divisor is taken from AOI summary table (the number of fixations for the given AOI) except when the AOI contains the last fixation. In that case the divisor is decremented by 1.

If there are no overlapping AOIs, the total of the values in each row must be 0 or 1.

11.5 Joint Probability Table

Highlight a "Joint Probability" node to see the total probability of a transition between two AOIs.

AOI Name	OUTSIDE	Duck1_Head	Duck2_Head	Duck1_Feet	Duck2_Feet
OUTSIDE	0.190	0.000	0.095	0.000	0.095
Duck1_Head	0.000	0.000	0.000	0.000	0.000
Duck2_Head	0.095	0.000	0.143	0.000	0.048
Duck1_Feet	0.048	0.000	0.000	0.048	0.000
Duck2_Feet	0.048	0.000	0.048	0.048	0.095

The entry on row *n*, column *m*, shows total probability that there was a fixation in *AOI n* followed by a fixation in *AOI m*. The value is calculated as the corresponding entry in the Transition Table divided by the total number of transitions. The number of transitions is calculated as one less than the total number of fixations (total_number_of_fixations – 1) because the last fixation does not transition anywhere.

If there are no overlapping AOIs, the total of all the table cell values should be 1.

12 Dwell analysis

The Dwell analysis further constrains the parameters of analysis from the Fixation Sequence. The function takes the results from the Fixation Sequence analysis and applies additional qualifiers.

An individual Dwell is defined as the time period during which a contiguous series of 1 or more fixations remains within an Area of Interest (AOI). That is, a dwell is defined as continuous time spent fixating within an area of interest (without leaving that area) regardless of how many individual fixations this involved.

The Dwell function creates the same set of reports as the Fixation Sequence function, however values are likely to be different owing to the difference in event definition.

This analysis type is generally preferred when the experimenter is interested only in the overall interaction with AOIs, not the individual fixation events within them.

Dwell data appears as a node on the tree diagram, at the same level as fixation sequence data. However, there is no separate command to perform dwell analysis computations. It is done automatically whenever fixation sequence data is computed. The dwell node always contains sub-nodes for an “AOI Summary” table, a “Transition table”, a “Conditional Probability” table, and a “Joint Probability” table.

12.1 Dwell Data list and Info tab

Highlighting a Dwell node on the project tree, in the main window left panel, displays a corresponding data list in the right panel.

Data | More Info

Dwell#	AOI#	AOIName	StartTime	Duration	StopTime	InterDwlDur
1	0	OUTSIDE	4.288	0.601	4.888	0.000
2	4	Duck2_...	4.988	1.018	6.006	0.100
3	2	Duck2_...	6.056	2.035	8.091	0.050
4	4	Duck2_...	8.158	0.100	8.258	0.067
5	3	Duck1_...	8.358	1.468	9.826	0.100
6	0	OUTSIDE	9.843	3.403	13.247	0.017
7	2	Duck2_...	13.280	1.051	14.331	0.033
8	0	OUTSIDE	14.381	0.117	14.498	0.050
9	4	Duck2_...	14.515	0.634	15.148	0.017
10	0	OUTSIDE	15.199	0.300	15.499	0.050
11	2	Duck2_...	15.516	0.234	15.749	0.017
12	0	OUTSIDE	15.766	1.335	17.100	0.017

The table lists sequential dwell periods on successive rows. For each row (dwell period) the columns indicate the number and name of the AOI being viewed, the dwell start time, duration, and stop time, and the time between the end of the previous dwell period and beginning of the current dwell period.

Dwells that are not inside any defined AOI are considered to be within AOI #0, named “OUTSIDE” (all defined AOIs are numbered starting with 1).

Note that, by definition, no two sequential dwells, and thus no two sequential rows on the dwell list, can be in the same area.

The "More Info" tab displays the rules used for selecting the AOI set (as described in section 9), the AOI set selected, and a list of the areas and boundary coordinates that make up that AOI set.

12.2 AOI Summary (for Dwells)

Highlight an AOI summary node (under a Dwell node), on the project tree in main window left panel, to see the AOI summary information in the right panel. The AOI summary includes the following fields, calculated for each AOI. All time intervals are in seconds.

Name	Description
AOI#	AOI number. Zero represents the area outside all AOIs
AOIname	AOI name (from AOI properties)
ScenePlane	AOI scene plane; *will be displayed only for EHD data files*
DwellCount	Number of dwells inside AOI
TotalDwellDur	Total duration of the dwells inside AOI
AvgDwellDur	Average duration of the dwells inside the AOI
MedianDwellDur	Median duration of the dwells inside the AOI. If dwell count is even (2n) than medium dwell is the one with n dwells shorter and n-1 dwells longer
SkewDur	Average dwell duration minus median dwell duration
STDDur	Standard deviation of durations

12.3 Transition Table (for Dwells)

Highlight a project tree "Transition Table" node (under a Dwell node) to see the Dwell transition table in the right window pane. Each cell in the Transition Table shows the number of transitions from a dwell in the AOI represented by the row number to the AOI represented by the column number. For example, the cell on row 3, column 5 shows how many times a dwell in AOI #2 was directly followed by a dwell in AOI #4 (reminder: first row and column represent AOI # 0).

The table contains non-negative integer values. The sum of all entrees should be equal to one less than the total number of dwells (total_number_of_dwells – 1). The diagonal elements (elements with the same row and column number) must always be zero since the definition of a dwell insures that a dwell can never be followed by another dwell on the same area (it would just be part of the previous dwell).

12.4 Conditional Probability Table (for Dwells)

Highlight a project tree "Conditional probability" node (under a Dwell node) to see the Dwell Conditional probability table in the right window pane. The entry on row *n*, column *m*, shows the conditional probability that if a dwell was in *AOI n*, the next dwell was in *AOI m*. The value is calculated as the corresponding entry in the dwell Transition Table divided by the total number of transitions from the AOI defined by the row. The number of transitions from the AOI is equal to the number of dwells on this AOI with one exception: the last dwell does not transition anywhere.

Therefore the divisor is taken from AOI summary table (the number of dwells for the given AOI) except when AOI contains the last dwell. In that case the divisor is decremented by 1.

The total of each row must be 0 or 1, and diagonal elements must always be 0.

12.5 Joint Probability Table (for Dwells)

Highlight a project tree "Joint probability" node (under a Dwell node) to see the Dwell Joint probability table in the right window pane. The entry on row *n*, column *m*, shows total probability that a fixation traveled from *AOI n* to *AOI m*. The value is calculated as the corresponding entry in the Transition Table divided by the total number of transitions. The number of transitions is calculated as one less than the total number of fixations (total_number_of_fixations – 1) because the last fixation does not transition anywhere.

The total of all the table cells should be 1, and diagonal elements must always be 0.

13 Pupil Diameter Analysis

ETAnalysis can process pupil diameter data by scaling and interpolating across blinks. It can also provide statistics for an event period that include average, median, minimum and maximum pupil diameter, and total number, average duration and average frequency of blinks. Of course it can do these things only if pupil diameter has been included on the data file as one of the recorded items (it is included by default; see the Eye Tracker manual for details). For data from the binocular *ETVision* system, a separate pupil diameter analysis is done for each eye.

In many cases, pupil diameter is measured in units that relate to the size of the pupil image on the camera sensor chip. To convert these measurements to the real diameter of the pupil in millimeters, a scaling factor must be determined as described in the next section. The Argus *ETVision* system can be set to report pupil diameter in millimeters (see *ETVision* manual Appendix D – System Configuration tab). In this case *ETVision* estimates the camera to pupil distance and uses the distance estimate to convert the camera pixel measurement to millimeters.

13.1 Determining A Pupil Diameter Scaling Factor

Eye pupil diameter data for eye trackers is usually used to track relative changes in pupil diameter rather than to measure absolute pupil diameter. Pupil diameter is often recorded in eyetracker units that relate to size on the eye camera field of view. Although it may vary a bit from subject to subject, *ETVision* systems will always have approximately the same camera to eye distance. Scale factors for converting eye tracker units to millimeters will be approximately 0.16 millimeters per eye tracker unit for *ETVision*.

As mentioned in the previous section *ETVision* systems can optionally be set to output pupil diameter in millimeters (based on an estimate of camera to pupil distance). In this case, no scale factor is required, but the data is still most appropriate for tracking relative changes in pupil diameter rather than precise measurement of absolute pupil diameter.

13.2 Performing a Pupil Diameter Analysis

Pupil diameter Analysis can be selected from an event node or any parent nodes above the event node level. The Pupil diameter Analysis will be conducted for data in all events under the node where the selection is made. Right click on the appropriate node and select **Pupil Diameter Analysis**.

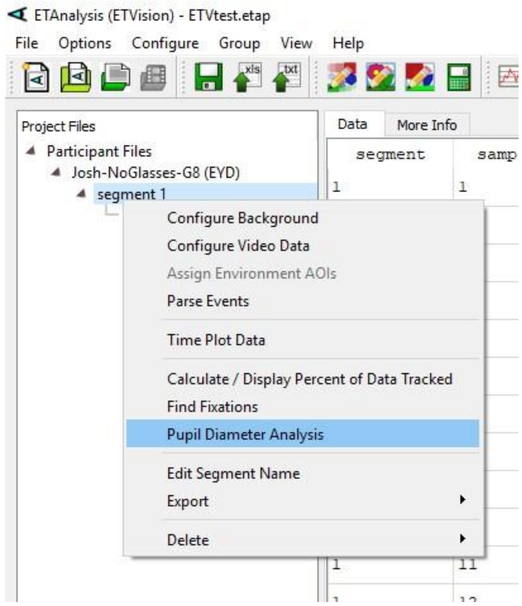

A **Pupil Diameter Analysis Configuration** dialog will appear.

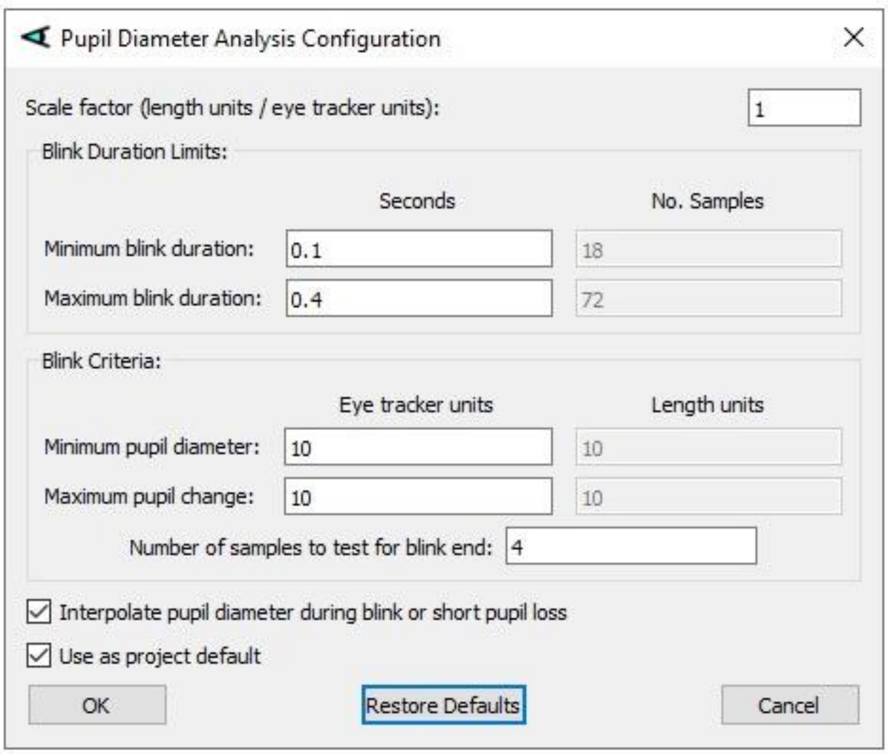

Use the Configuration dialog to define the pupil diameter scale factor (millimeters per eye tracker unit), conditions considered a pupil loss, and pupil loss durations that will be considered "blinks". If *ETVision* data was recorded with "Pupil Size" units set to "Millimeters", set the Scale factor to 1. Note that pupil recognition can be lost for reasons other than blinks, for example, optics being bumped, poor discrimination of pupil edges, etc. Use the "Restore Defaults" button to restore "factory" default values.

A "pupil loss" starts if:

Pupil diameter drops below "Minimum pupil diameter"

OR

Pupil diameter change in one cycle (i.e. abs(pupil – previous_pupil)) exceeds the "Maximum pupil change".

Usually, values corresponding to about 1 mm are appropriate. *Caution: if ETVision data was recorded with "Pupil Size" units set to "Millimeters", the default values for "Minimum pupil diameter" and "Maximum pupil change" will usually not be appropriate. In this case the appropriate values will usually be 1.0.*

A "pupil loss" is considered finished if, for the number of sequential samples specified by "No. samples to test for blink end" (default: 4):

Pupil diameter is greater than "Minimum pupil diameter"

AND

Pupil change in one cycle is smaller than "Maximum pupil change"

If the "Interpolate pupil diameter." check box is checked, interpolation will be performed for pupil losses of duration less than "Maximum blink duration".

When interpolation is done, the program performs a linear interpolation starting 4 fields from the loss beginning (or "No of samples to test for blink end", if this is less than 4) and ending an equivalent distance from the end of the "loss". If a pupil loss exceeds the "Maximum blink duration", the program will not interpolate. The user can also disable interpolation altogether by un-checking the box labeled "Interpolate pupil diameter during blink or short pupil loss".

All interpolated records are marked as "Interpolated".

Not all records with lost or suspicious pupil data are marked as "blinks". A b**link is a group of records during which the pupil was "lost" (not recognized) continuously for a period that is within the specified range** (default: 0.1 sec to 0.4 sec), **and that is preceded by a minimum number of non-loss fields** (default: 4 samples). If a pupil loss period is shorter than the minimum, longer than the maximum, or not preceded by a minimum number of non-loss fields, the program concludes that the pupil loss was probably caused by something other than a blink. Such an interval may be interpolated (provided that it is not too long) but it will **not** be marked as a "blink".

When OK is clicked on the Pupil diameter analysis configuration dialog, a pupil diameter analysis node, labeled "PD Analysis" is created under the appropriate event nodes. If the project data is from *ETVision*, there will be two nodes, one labeled "PD Analysis left eye" and another labeled "PD Analysis right eye".

13.3 Pupil Analysis Display

Data | More Info

sample_#	time_secs	pupil_diam	scaled_pupil_diam	xdat	blink	interpolated
1	0.000	61.090	4.887	0		
2	0.006	61.530	4.922	0		
3	0.011	63.890	5.111	0		
4	0.017	63.030	5.042	0		
5	0.022	61.780	4.942	0		
6	0.028	62.120	4.970	0		
7	0.033	61.080	4.886	0		
8	0.039	61.260	4.901	0		
9	0.044	61.450	4.916	0		
10	0.050	61.710	4.937	0		
11	0.056	61.510	4.921	0		

When a PD Analysis node is highlighted in the tree diagram the data tab in the right window pane will display a list of pupil diameter values in the units originally recorded as well as the scaled values with interpolated data. Other columns show the sample number and time values (starting from zero at the beginning of the event), and XDAT values. There are also columns that flag data records that have been determined to be part of a blink and data records for which the scaled pupil diameter value is an interpolated value. Any of the data columns can be displayed on a time plot by right clicking the PD Analysis node and selecting “Display Time Plot” from the pop up context menu (see section 14.1 for a description of time plots). The data window contents can also be exported to Excel or to a text file by right clicking the PD Analysis node and selecting “Export”.

The “More Info” tab displays all of the parameters that were selected in the “Pupil diameter analysis configuration” dialog, and also shows all of the summary pupil diameter statistics computed for the event. Event duration is available in several places and is repeated here for convenience.

Pupil recognition time is divided into several categories. “Pupil Available” is the total amount of time during which the pupil was not considered to be “lost”, as defined in the previous section. Total pupil loss is the time during which the pupil was considered to be “lost”. “Total pupil loss” plus “Pupil Available” should equal Event duration. “Loss due to blinks” is the total time considered to be part of a blink as defined in the previous section. This is less than or equal to “Total pupil loss” time (some pupil loss may not meet “blink” criteria). “Loss due to overtimes” is the sum of any missing data records in the file.

Pupil diameter statistics include the minimum, maximum, and average, and median pupil diameter, and standard deviation of pupil diameter during the event. All of these quantities are computed from the scaled pupil diameter data considering only data fields for which the pupil was not defined as “lost”, “blink”, or “interpolated”. Blink statistics include total number, average duration, and average frequency, for blinks as defined in the previous section.

14 Graphics Displays

ETAnalysis can display time ("strip chart") plots of raw data and fixation data, X/Y plots of raw gaze data and fixation data superimposed on background images (often called gaze trail and scan path plots), and heat map displays superimposed on background images. Raw gaze data is always shown superimposed on the same time plots with fixations. Raw data can also be superimposed with fixation data on X/Y plots, if desired. Multiple X/Y plots, from different events (often corresponding to different "trials"), can all be superimposed. Heat map displays can either represent data from a single event, or pooled data from multiple events. All displays can be saved as bit map images for inclusion in documents or use in other applications.

In addition, gaze trails, heat maps, and fixation plots, can be played back dynamically (showing their original time course) over static backgrounds.

This section discusses graphics displays on static background images. Dynamic displays on scene video images are discussed in section 16.9.

14.1 Time Plots

Time plots of any item on the data file are available on context menus for data segment nodes, and event nodes. Time plots including an overlay to show detected fixation periods are available only from fixation nodes and saccade nodes, time plots including an overlay to show detected saccade periods are available only from saccade nodes, and time plots of pupil diameter analysis data are available only from PD Analysis nodes.

To plot raw data (data directly from the *.eyd, *.ehd, or *.csv file), select **Time Plot Data** from the context menu at a segment or event node. A segment node will include all of the data in the original recorded data segment (remember that data segments are sections of continuous data). An event node, which may include all or just part of the original data segment, will include only the data that is part of that event.

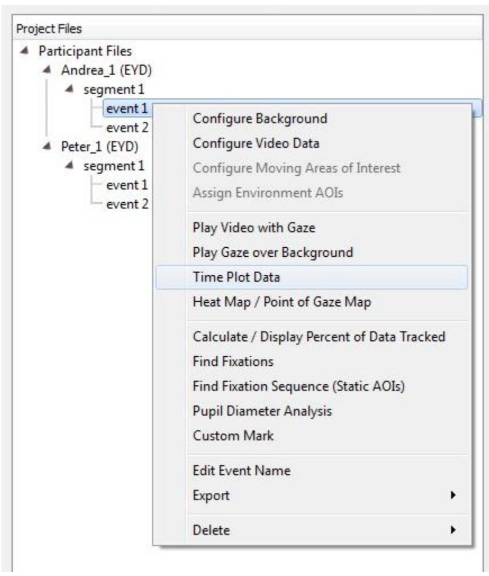

Selecting **Time Plot Data** will bring up a **Configure Time Plot** dialog.

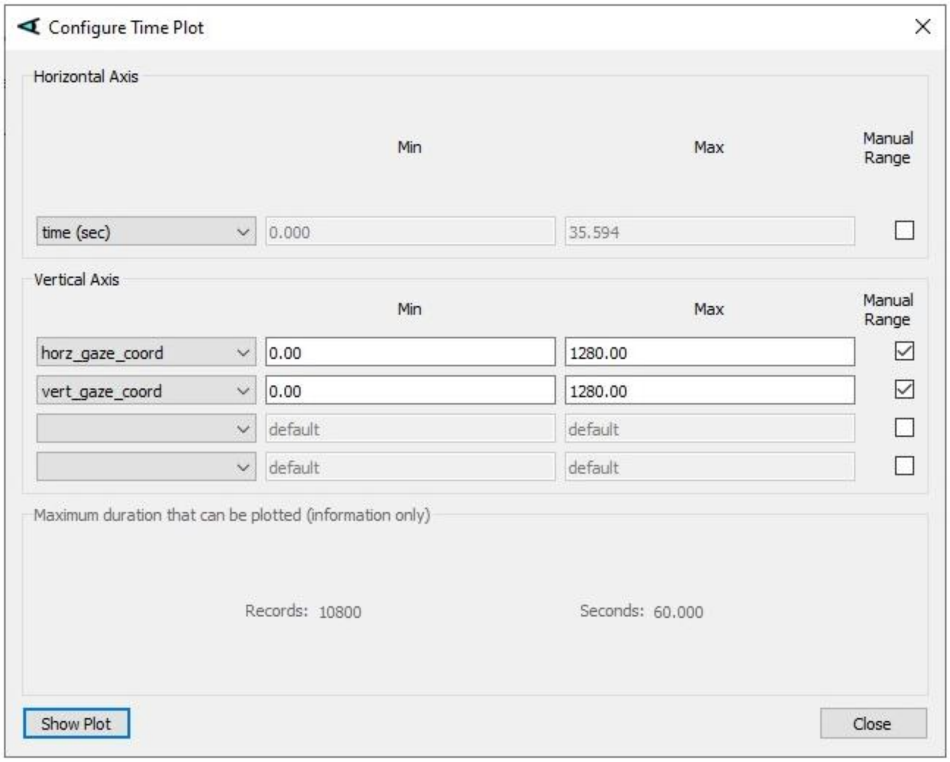

Up to four time plots can be shown on the display area, and the data item to be plotted on the vertical axis of each is selected from the drop down menus under "Vertical axis". Each drop down menu will list all of the data items that available on the data file for plotting. Min and Max values represent the range that will be represented by the vertical axis of the graph. These will default to the minimum and maximum values of the data item during the time period specified at the top of the dialog. To manually specify a range, check the "Manual Range" box and type in the desired values.

The horizontal axis is always time, specified and labeled in terms of seconds. The initial time range will default to either the entire event (or segment), or to 10800 data records, which ever is smaller. To enter a different time range, check the "Manual Range" box and type in the desired value. Once the plot is displayed, a time zoom function is also available on the plot window.

Clicking the "Show Plot" button will create a new tab in the display area showing the time plot display.

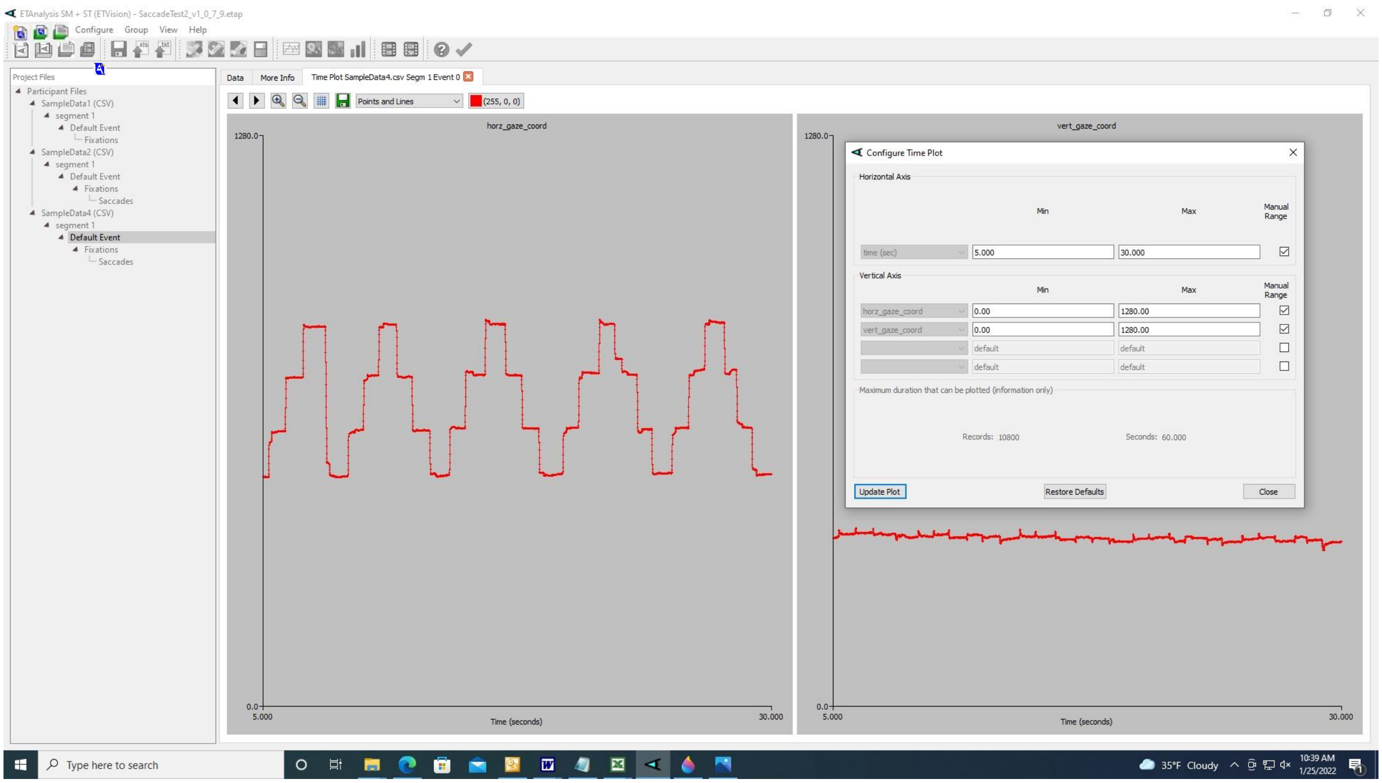

Buttons on the Plot window can be used to zoom in and out, move along the time scale, turn on a grid, or change the color of a plot component. The plot can be saved as an image file by clicking the save-image button

Plot commands:

- Move Back
- Move Forward
- Zoom In (time scale)
- Zoom Out (time scale)
- Show Grid
- Save plot to image file

The “Zoom In” and “Zoom Out" buttons expand or contract the time scale by a set amount for each click, but the mouse wheel (or zoom gesture on touchpad) can also be used for more continuous control. The mouse can also be used to move forward or backward along the time scale by using the right mouse button to “drag” left or right in the plot area.

A drop down menu selects between display of individual data points (useful for high time resolution), lines (default) or points connected with lines.

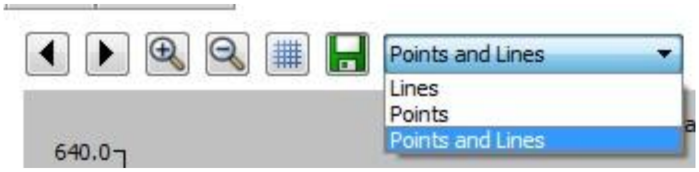

A button to the right of the “Points and Lines” drop down brings up a plot color selection dialog.

The grid lines button will enable or disable a grid display. When enabled, each plot window is overlayed with a set number of vertical and horizontal black grid lines, and light blue vertical lines which indicate 5 second intervals from the start of the event (or segment).

The **Configure Time Plot** dialog remains active, along with the plot, but the “Show Plot” button changes to “Update Plot”. It is possible to modify the time range or any of the data range values by changing these things in the **Configure Time Plot** dialog and then clicking the “Update Plot” button. The set of items to be plotted cannot be modified without first closing the plot.

To see time plots with an overlay showing detected fixations, select **Display Time Plot** from the context menu at a **fixation** node.

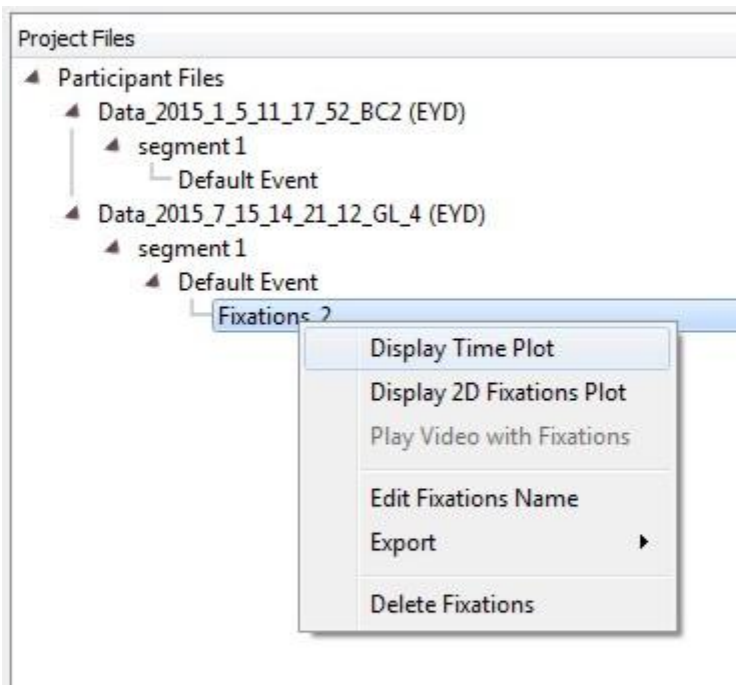

A **Configure Time Plot** dialog will appear just as previously described. Plots of the horizontal and vertical gaze data used to compute fixations will show both raw gaze data (red lines on plot shown below) and the fixation computation results (blue lines on plot shown below). When superimposed on gaze position data, the blue fixation lines show both fixation position and duration. Note that if the *ET3Space* feature was used when collecting data, it is the “ET3S_horz_gaze_coordinate” and “ET3S_vert_gaze_coordinate” items that should be plotted to evaluate fixation computation results. If *ET3Space* was not used, select the “horz_gaze_coordinate” and “vert_gaze_coordinate” data items.

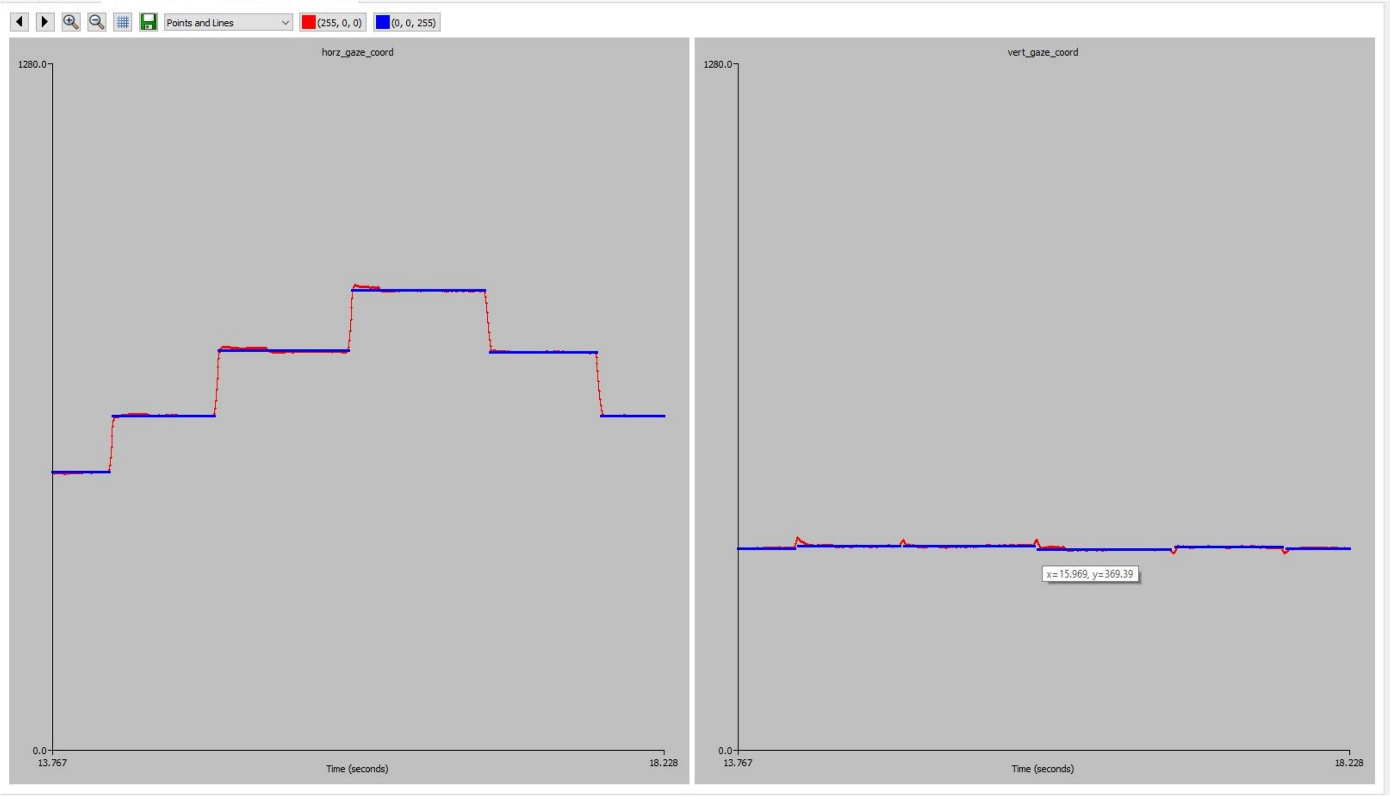

If the plot start and end points are selected by time, any fixation overlapping with the given time interval will be shown. In the example below all three fixations will be shown.

```
< Fixation 1 >          < Fixation 2 >        < Fixation 3 >

       ↑ Start Time                        ↑ Stop Time
```

To see time plots with an overlay showing detected saccades and fixations, select **Display Time Plot** from the context menu at a **saccade** node. To best evaluate the classification of saccade periods, select "horz_gaze_coordinate" and "vert_gaze_coordinate" data items. Periods identified as saccades with be overlayed with a yellow line connecting the first and last data points in the saccade (see section 10.2.4 for additional detail).

14.2 Two Dimensional Plots

This section discusses graphics displays on static background images. Dynamic displays on scene video images are discussed in section 16. ***All of the plots discussed in this section require that a background image has been configured, as described in section 7, and associated with the relevant event as described in section 9.***

Note that two-dimensional plots over a background image will make most sense when the gaze data specifies gaze with respect to a surface in the environment. This is the case for eye trackers that use remote (table mounted) optics, or for head mounted eye trackers like Argus Science *ETVision* when using the *ET3Space* feature or the *StimTrac* feature.

When not using *ET3Space* or *StimTrac*, *ETVision* data specifies gaze position with respect to head mounted scene camera image rather than objects in the environment. Environment surfaces and objects change position within the scene image as the participant's head moves about. In this case, static two-dimensional plots can show gaze patterns with respect to the participant's visual field, and this may sometimes be of interest; but the plots will not show gaze patterns with respect to an environment background unless the head was perfectly stationary.

If an environment object in the head mounted scene camera image is tracked with a Moving Area of Interest, it is possible to plot gaze over a static image of that object. See sfection 16 for an explanation of Moving Areas of Interest, and section 15.1.3 to make 2 D plots over static images of the tracked objects.

14.2.1 Heat map, Peek map, and point-of-gaze scatter plots

Heat maps, peek maps, and scatter plots show the relative density of visual activity. In the case of heat maps, the colors closer to the red end of the spectrum indicate the most visual activity (most time spent gazing towards these areas), while "cooler" colors indicate progressively less visual activity. Peek maps display a dark, semi transparent mask over the background with areas of high visual activity being the most transparent. Another way to visualize the same thing is to simply draw a dot for each gaze data sample. The thickest clusters of dots indicate the highest density of visual activity.

Heat maps, peek maps and point-of-gaze (POG) scatter plots are available only from **event** node context menus.

On the tree diagram in the left panel select any event node and right click to display a context menu. Select "Heat Map / Point of Gaze Map"**.**

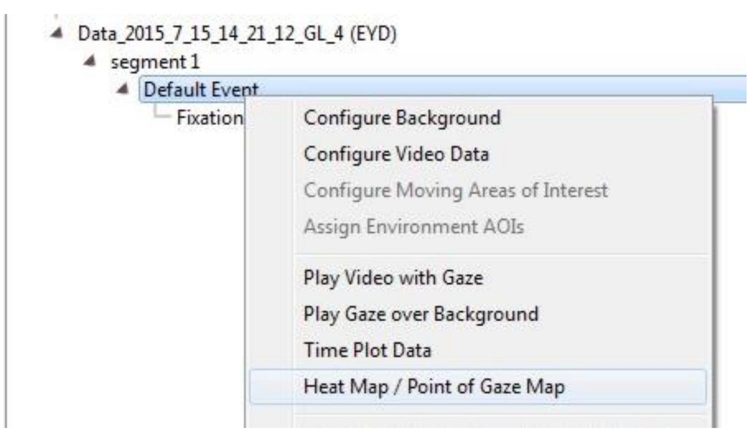

The **Heat Map Configuration** dialog will appear. Select the background image on which the plot will be drawn from the drop down menu labeled "Background". To be on the drop down list an image must have been "configured" and added to the project as described in section 7.

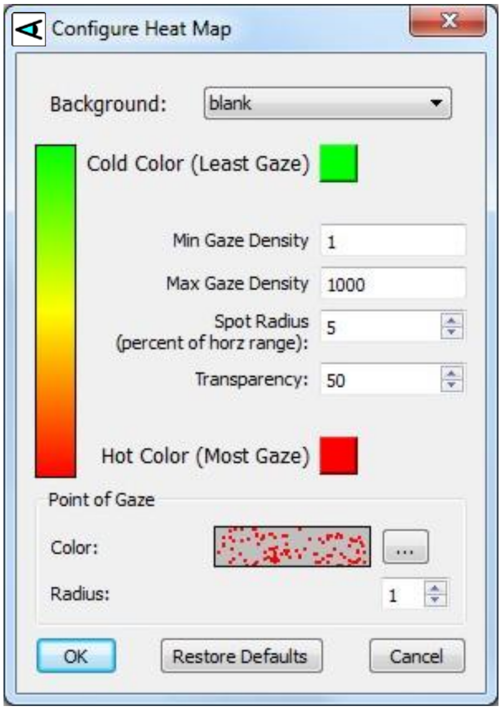

The color bar on the **Configure Heat Map** dialog shows the range of colors, with "hottest" at the top and "coolest at the bottom. The Max and Min "Gaze Density" values control the gaze density that will be mapped to the hottest color and coolest color respectively. "Max Gaze Density" values greater than the number of gaze sample points in the event will have no affect. As the "Max Gaze Density" value becomes progressively smaller than the total number of gaze sample points it will have the affect of progressively increasing the proportion of the map shown as a "hotter" color.

Increasing the "Min Gaze Density" value will have the affect of shrinking the outer boundary of the heat map (increasing the area shown with no heat map color).

"Spot radius" controls the size of the "blob" made by any given density of gaze activity.

"Transparacy" controls the transparency of the heat map overlay with greater values making the overlay increasingly transparent.

"Point of Gaze", "Color:" is simply the color of the dots used for the scatter plot, and "Radius:" is the size of each dot.

The best way to see the affect of the Heat Map Configuration controls is to make a heat map plot by first using the default settings and then experiment by varying the parameters. Increasing "Spot Radius" will make bigger blobs, decreasing "Max Gaze Density" will increase the share of red colored areas. Increasing "Min Gaze Density" will cause more areas of low gaze density to be ignored by the heat map. This can be useful for plots that include a large volume of data. Modifying "Cold Color" and "Hot Color" will affect the colors shown on the heat map, but the defaults will most often be the best choice for these.

Click OK on the **Heat Map Configuration** dialog to make the heat map appear as a new tab in the display area. Note that this may take several seconds. The amount of time it takes to generate the plot will depend on the amount of data in the event.

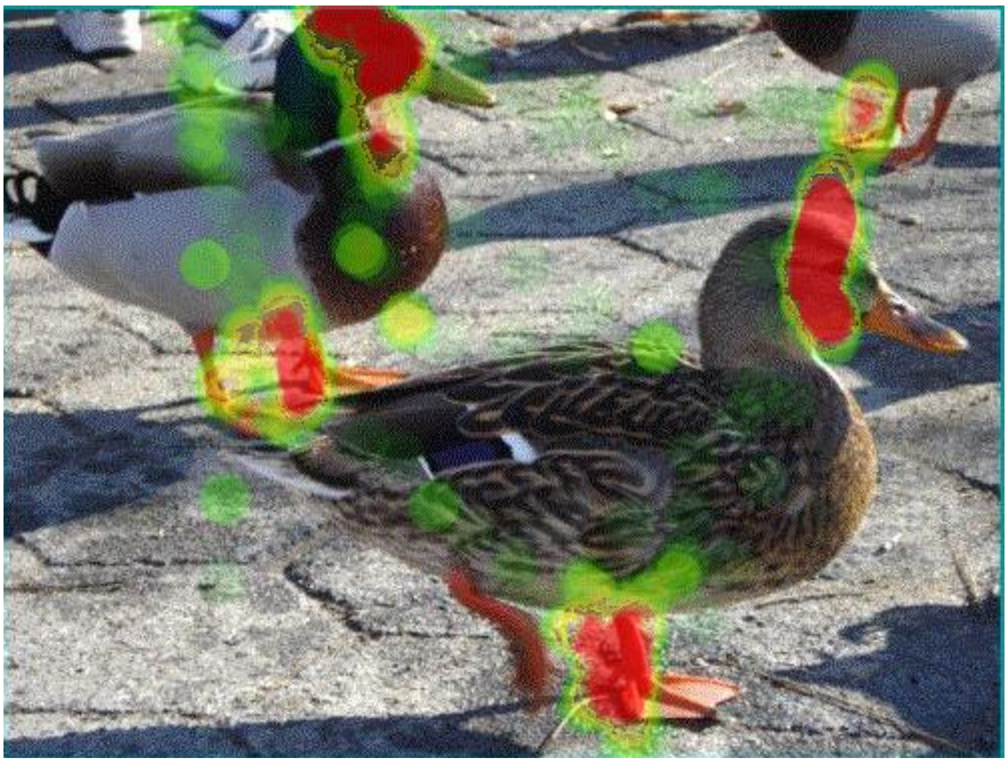

To change one of the items on the **Heat Map Configuration** dialog, click the "Configure Heat Map" button to bring up the dialog, make the desired change and click OK.

A set of radio buttons above the graphics display can be used to select a scatter plot ("Points of Gaze") or a Peek Map instead of the Heat Map.

The "Peek Map" will be similar to the "Heat Map", but with varying transparency values replacing the "Hot" and "Cold" color values. "Hotter" areas will correspond to greater transparency, allowing the background image to be visible, and colder colors will correspond to less transparency (a more opaque overlay) making the background image progressively less visible.

To save the display as an image file, click the "Save as Image" button.

Data from multiple events can be combined in a single heat map, peak map, or scatter plot diagram by selecting **Group→ Group Heat Map** on the *ETAnalysis* main window.

A selection dialog will appear, in the form of a tree diagram, showing all of the event nodes in the project. Each node of the tree is a check box. Checking an event node selects only that node. Checking a higher-level node selects all of the event nodes underneath it. The dialog also has a "Select all" button and a "Unselect all" button.

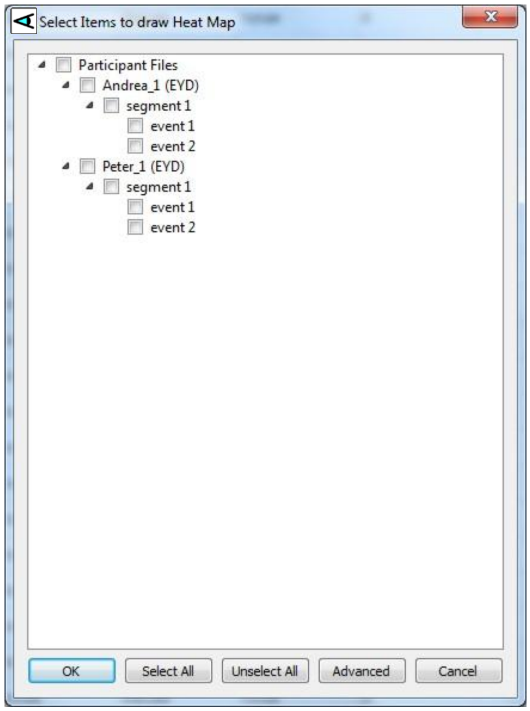

Click the "OK" button to accept the choice of events and bring up the **Heat Map Configuration** dialog previously described. Be sure the desired background is selected, and adjust other parameters as previously discussed. When "OK" is clicked on the **Heat Map Configuration** dialog a heat map plot will be generated based on pooled data from all of the events selected. It desired, use the radio buttons at the top of the display to select "Peak Map" or "Gaze Points" instead of "Heat Map".

14.2.2 Two Dimensional Fixation Scan Plots

Scan plots are 2 dimensional ("X/Y") plots of fixation points, optionally connected by lines between successive fixations, and superimposed over one of the previously configured background images (presumably the image that was viewed by the subject). Scan plots of the data from any single fixation node can be made from the context menu at that node. Superimposed scan plots from multiple fixation nodes (for example, from multiple subjects or multiple trials) can be made from the "Group" menu on the main *ETAnalysis* window.

14.2.2.1 Scan plot from single fixation node

On the tree diagram in the left panel select any fixation node and right click to display the context menu. Select **Display 2D Fixations Plot**.

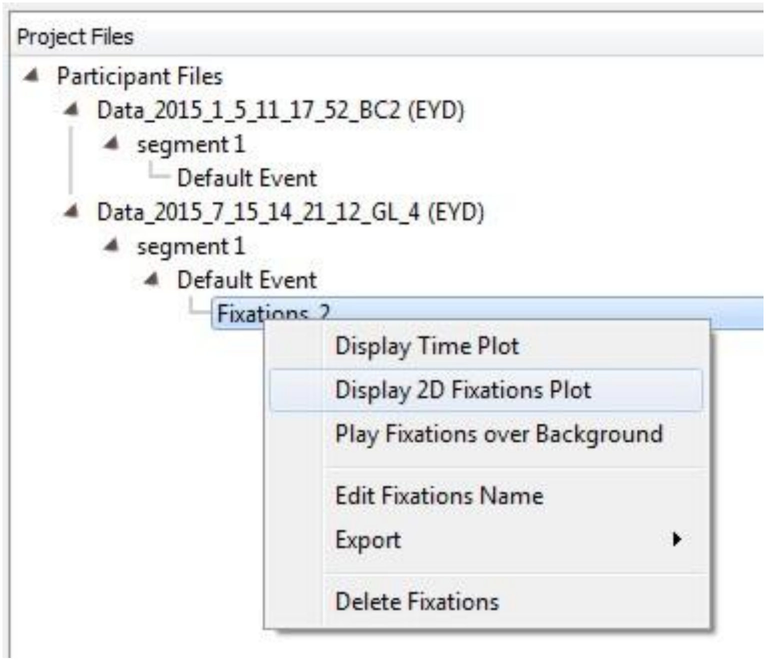

The **Configure Fixation 2D Plot** dialog will appear.

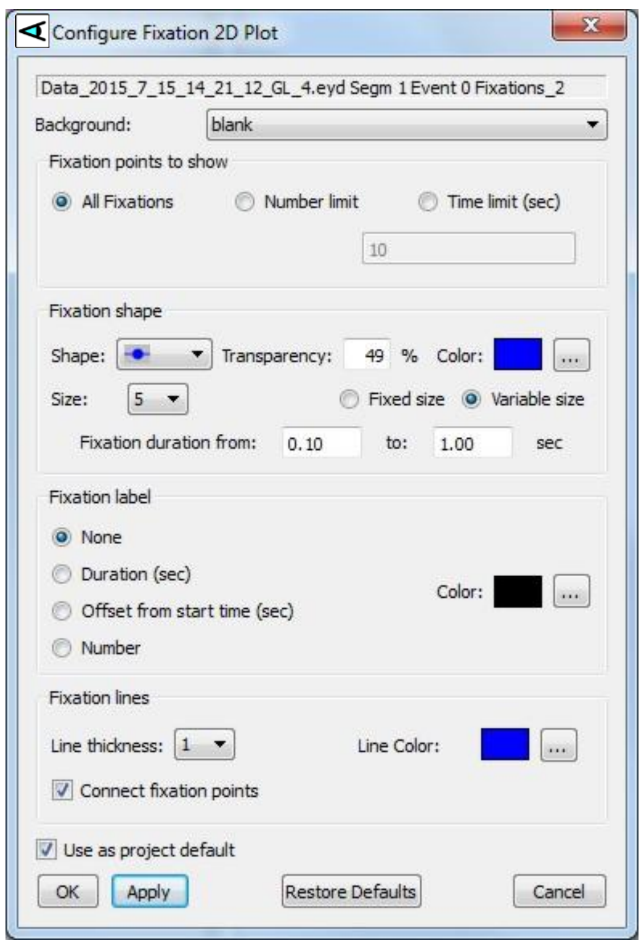

Select the appropriate background from the drop down menu. The other settings on the configure dialog affect the appearance of the plot.

Under "Fixation points to show" select which fixations to plot. All fixations means all fixations in the selected node. Alternately, plot just the fixations between specified time points or fixation numbers. Remember that time always starts at zero at the beginning of the fixation set (beginning of the data represented by the node). The first fixation in the node is fixation 1.

If number or time limit subset is chosen it will be possible to scroll forwards or backwards through the data once the plot is opened.

Under fixation shape, select the shape, color and size of the symbols plotted at each fixation point. If variable size is selected, the diameter of symbol at each point will be proportional to duration of that fixation so long as the duration is between the specified limits ("Fixation duration from"). If longer than the upper limit, the diameter will not grow beyond the diameter associated with the upper time limit. If shorter than the lower limit, the diameter will not decrease. To make all of the symbols proportionately larger or smaller, adjust the "Shape size" value.

If desired each fixation point can be labeled with its duration (in seconds), the start time of the fixation (measured from the beginning of fixation data in the selected node), or the sequential number of the fixation. Select the appropriate radio button under "Fixation label". Select the label color to contrast with the background image so that the labels will be visible. It usually makes sense to specify labels only if a small number of fixations are displayed on the plot. If a lot of fixation points are displayed the labels are usually too crowded to be legible.

"Fixation Lines" refer to the lines that connect sequential fixation points. Line thickness and color can be adjusted, or by not checking the "Connect fixation points" box the connecting lines can be omitted altogether. In this case only the fixation point symbols will appear. If the "use as project default" box is checked the dialog settings will be used as the project default settings the next time **Display 2D Plot** is selected from the fixation node context menu.

Click OK. The program will display a 2D plot of fixations (scan plot).

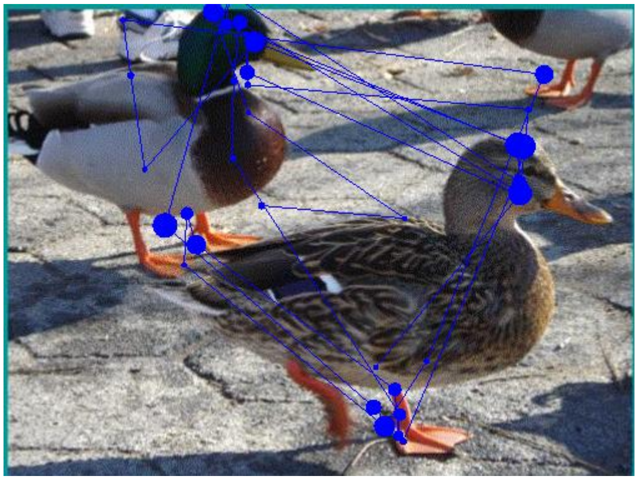

Right clicking on a fixation point will bring up a context menu as shown below.

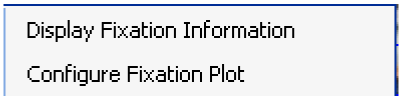

Select "Display Fixation Information" to see the digital data values for this fixation. This is the row corresponding to this fixation from the data table described in section 10.1.5. Select "Configure Fixation Plot to bring up the Configure Fixation 2D plot dialog previously described, and click Apply or OK on this dialog to see the affect of any changes.

A set of shortcut buttons and drop down menus at the top of the plot window allow numerous display options. Hovering the mouse over any button will display a label with its function

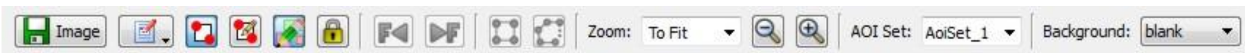

The current display can be saved as an image file, the configuration dialog (previously described) can be brought up for modification; background attachment points (see section 7) can be displayed or hidden; attachment points can be edited (as described in 7); if areas of interest have been created for the background they can be displayed or hidden, and can be unlocked for editing. If areas of interest are unlocked for editing, the same editing controls previously described (see section 8.1) become available.

The Configure button  is actually a drop down menu as shown below.

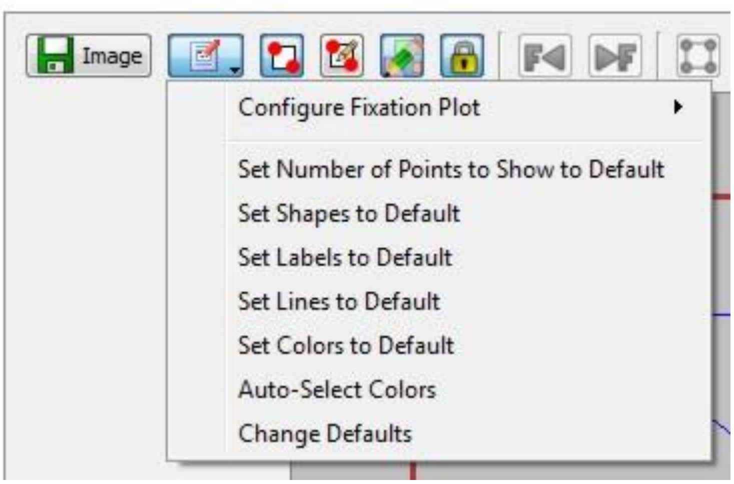

Clicking the first item ("configure Fixation Plot") will display the information that identifies the particular fixation node being plotted; specifically the data file name, segment number, event number, and fixation node name.

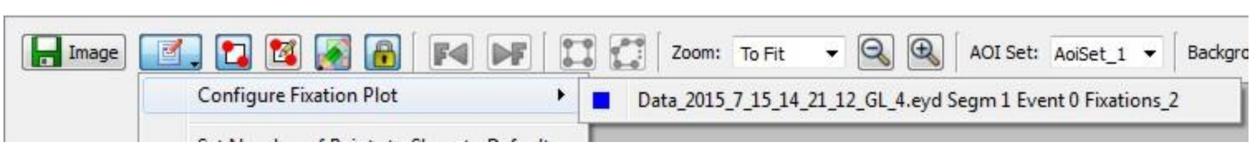

Left clicking this node specification will bring up the **Configure Fixation 2D Plot** dialog for examination or modification. The other menu items can be used to return Fixation Shapes, labels, and Line properties to their default conditions, or to "Auto-Select" colors. **Auto Select Colors** is useful primarily when multiple fixation sets are superimposed, as described in the next section, and need to be distinguished from each other.

If a subset of the fixation set is being displayed, the left and right fixation arrow symbols  can be used to scroll forwards (right arrow) or backwards (left arrow) through the data. For example, if 10 fixation points are being displayed ("Number Limit" selection on **Configure Fixation 2D Plot** dialog), the right arrow will advance to the next 10 fixation points and the left arrow will show the previous 10. If a time interval was selected ("Time Limit" selection on **Configure Fixation 2D Plot** dialog) the arrows will advance or back up by the same size time interval. If the entire fixation set is being displayed, the arrow symbols are grayed out and are not active. Similarly, if there is no additional earlier data or later data than that currently displayed the corresponding arrow is gray.

Other drop down menus are used to select the magnification of the background image, the AOI set that can be displayed over the background image, and the background image itself.

14.2.2.2 Scan Plot showing multiple fixation sets (from Group menu)

To superimpose data from multiple fixation sets on a single background image use the "Group" menu, on the main *ETAnalysis* window. Select **Group → Display Fixation 2D Plot**.

A selection dialog will appear, in the form of a tree diagram, showing all of the fixation nodes in the project. Each node of the tree is a check box. Checking a fixation node selects only that node. Checking a higher-level node selects all of the fixation nodes underneath it. The dialog also has a "Select all" button and a "Unselect all" button.

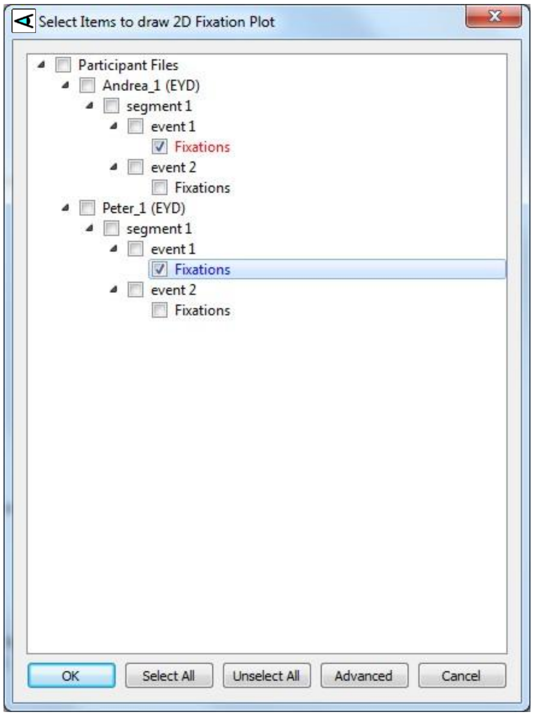

Note that when a fixation node is selected it changes color to show the color that will be used to display fixations from this node. To change the color, right click on the Fixation node, select **Configure Fixation 2D Plot** from the context menu**,** and set the desired color on the dialog.

Click **OK**. The program will display, in a new display area tab, a 2D plot of data from the fixation sets selected.

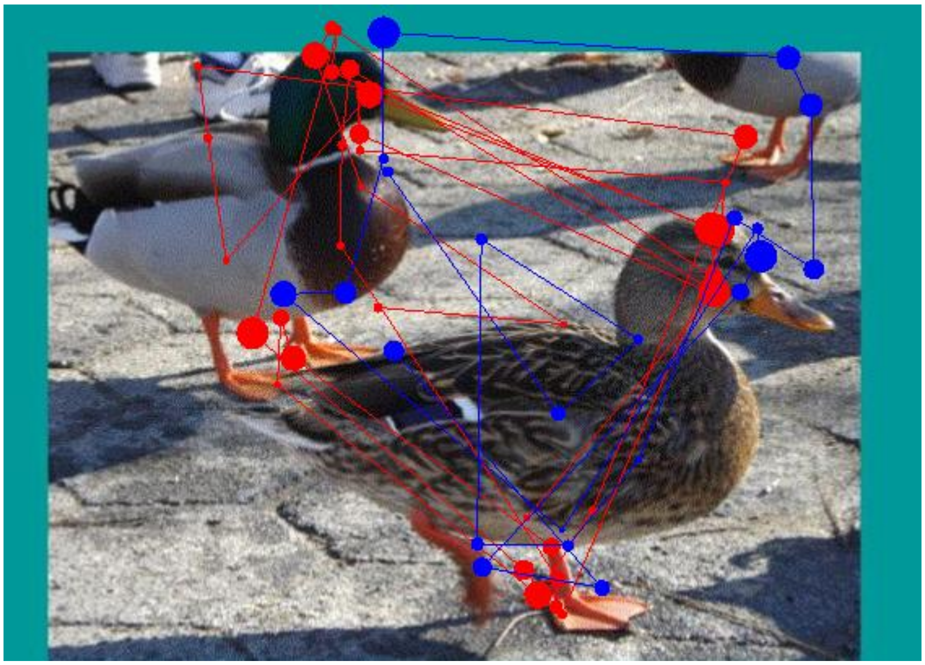

The background image will initially be a default selection. Use the Background drop down menu to select the desired background. Similarly, if an AOI set will be displayed, select the desired AOI set from the AOI drop down menu. A legend will show which display color is associated with which fixation set

The **Fixation 2D Plot** Window includes the same menu bar described in the previous section.

To change the display properties of data from any individual fixation set, right click on one of the fixation points from that set and select “Configure Fixation Plot”, or select “Configure Fixation

Points" from the **configure** button  pull down menu**,** and select the desired fixation set from the list that appears. Note that each list item is displayed with the same color as the corresponding scan path plot.

The display properties of all plots can be simultaneously changed to default values. Note that the **Configure Fixation 2D Plot** dialog has 4 labeled categories: "Fixation Points to Show", "Fixation Shape", "Fixation Label", and "Fixation Lines". The **Configure button** pull down menu also has a separate selection to set each of these categories to its default. For example, to change all fixation sets to use the same line thickness, first use **Change Defaults** to set the desired line thickness, then select **Set Fixation Lines to default.**

14.3 Bar Plots

14.3.1 AOI Bar Plots

After computing fixation sequences as described in section 11, bar plots can be produced showing one bar for each AOI corresponding to a variety of statistics.

View AOI Bar Plots for an individual set of Fixation data by right-clicking the Fixation Sequence or Dwell node and selecting "Display AOI Bar Plots" or by selecting  from the main toolbar while the desired Fixation node or any of its sub nodes is highlighted in the Project Tree. View AOI Bar Plots for data across multiple fixation data nodes by selecting "Display AOI Bar Plots" from the Group menu, and then selecting the nodes to include from the resulting selection dialog. Note that in this case <u>all included fixation sequence nodes must use the same AOI set</u>. The "Display AOI Bar Plots" Group menu selection will be available only if there are Fixation Sequence nodes in the project.

A bar plot will appear in a new "AOI Bar Plots" or "Group AOI Bar Plots" tab on the display window. Use the radio buttons labeled "Plot Options" to select the statistic to be shown by the bar plot. Each of these choices is explained in a following subsection. If it is a group plot, another tab labeled "Fixation Sequence Data" can be used to modify the fixation sets included.

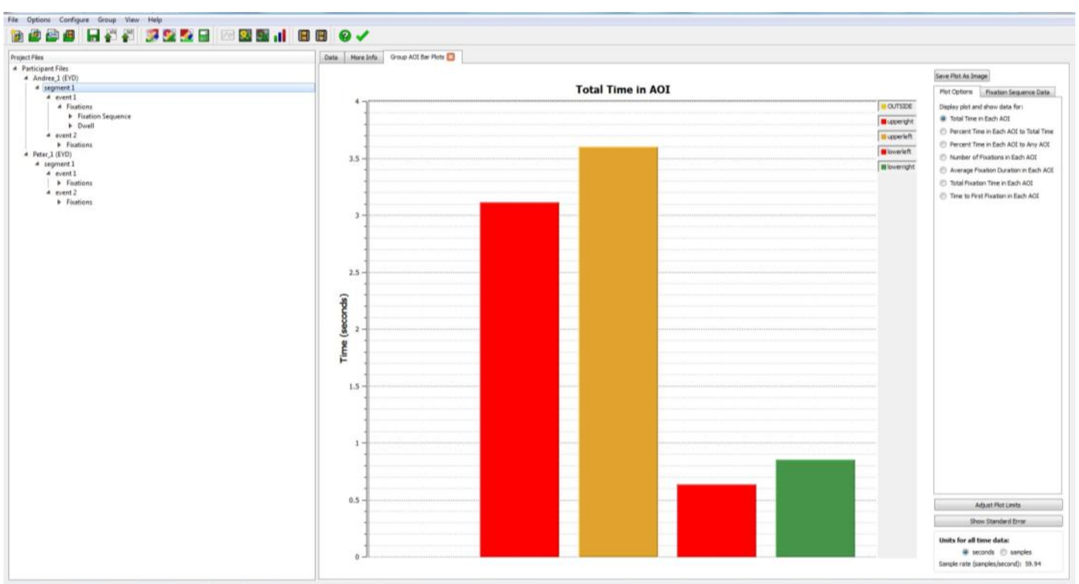

If a plot showing time information has been chosen the desired units (seconds or samples) may be selected at the lower right of the Display Area. The default is to show units of time in seconds. The frame rate used to convert between these values is shown for reference underneath the selection.

Use the "Adjust Limits" button to change the vertical axis maximum or minimum. Click an item in the AOI legend, to the right of the plot, to temporarily hide or unhide the bar associated with that AOI. In the case of Group Bar Plots, click the "Standard Error" button to show or hide a standard deviation display.

If fixation sequence data is computed with respect to "Moving Areas of Interest" as described later on in section 16, bar plots can also be created from those fixation sequence nodes just as described in this section.

The various statistics that can be selected are explained in the following subsections.

14.3.1.1 Total time in each AOI

This value is calculated using the total number of gaze points within fixations and between consecutive fixations in the same AOI; in other words, the total time spent dwelling on an AOI (see Section 9 for description of dwells). These values can be examined from the **Dwell → AOI Summary** node in the "TotalDwellDur" column of the Data tab.

14.3.1.2 Percent time in each AOI to total time

These values are calculated from the Dwell > AOI Summary table by dividing the "TotalDwellDur" by the Event duration and multiplying by 100 to obtain a percentage; the Event duration can be found by looking at the "More Info" tab when the Event is selected.

14.3.1.3 Percent time in each AOI to any AOI

These values differ from the previously mentioned values (Percent time in each AOI to total time) because instead of dividing by the total Event duration, *ETAnalysis* divides by the sum of the total dwell durations for each AOI, excluding the Outside AOI (which represents gaze not in a defined AOI). This sum can be calculated by adding the values in the **Dwell → AOI** Summary "TotalDwellDur" column for all AOIs except the Outside AOI shown in the first row of the table. Since the Outside AOI is not relevant to this bar plot, its bar is not present in the plot. The remaining bars for each AOI will have the same relative sizes as in the previous plot but with a different overall scale factor.

14.3.1.4 Fixations in AOIs bar plots

The fixation bar plots are fairly self-explanatory and include: Number of Fixations, Average Fixation Duration, Total Fixation Time and Time to First Fixation. The time to first fixation bar plot shows the time from the beginning of the event to the first fixation in the AOI.

14.3.1.5 Average Pupil Diameter in each AOI

This plot shows the average pupil diameter corresponding to all gaze points that were within this AOI.

14.3.2 Live AOI Bar Plots

ETVision eye trackers allow scene camera image features to be recognized and tracked, in real-time as "Live Areas of Interest" (LAOIs). Data files recorded by *ETVision* can include a data item, called "Gaze_LAOI", specifying which LAOI or LAOIs contained the point-of-gaze at each data sample. *ETVision* can also compute fixations in real-time ("live fixations"), and recorded files can include a data item ("fix_duration") specifying whether a fixation was detected to be in progress. (See the *ETVision* manual for explanation of LAOIs and "Live Fixations").

If "Gaze_LAOI" was recorded, it can be used by *ETAnalysis* to quickly display "LAOI bar plots" relating gaze time variables to each LAOI. If "live fixation" data was also recorded, "LAOI bar plots" can also show fixation variables. The choice of variables is same as for AOI Bar Plots described in section 14.3.1 except that the areas of interest are the "Live Areas of Interest" tracked by *ETVision* rather than areas of interest specified in *ETAnalysis*; and the fixation variables are "Live Fixations" detected by *ETVision* in real-time, rather than fixations computed off-line by *ETAnalysis*.

To generate a bar plot for a single event, simply right click on an Event node, and select "Display LAOI Bar Plots" from the drop down context menu. A bar plot will appear on an "LAOI Bar Plots" tab.

If multiple events have LAOIs with the same name, it is possible to generate a bar plot that shows the average data across these events. Select "Group Display LAOI Bar Plots" from the "Group" pull down menu on the *ETAnslysis* menu bar, and then select the event nodes to be included from the resulting pop-up dialog. A bar plot will appear on a "Group LAOI Bar Plot" tab.

Selection of the variable to be plotted for each LAOI, and other plot features are the same as for AOI bar plots described in section 14.3.1.

14.3.3 AI Object Bar Plots

ETVision eye trackers can use an AI model for real-time detection objects in the scene camera image. In this case, data files recorded by *ETVision* can contain a data item called "Gaze_AI" specifying which AI detected object contained the point-of-gaze at each data sample. (See *ETVision* manual for explanation of AI object detection).

If "Gaze_AI" was recorded it can be used by *ETAnalysis* to quickly display "AI bar plots" relating gaze time variables to each AI detected object. If "live fixation" data was also recorded, "AI bar plots" can also show fixation variables. The choice of variables is same as for AOI Bar Plots described in section 14.3.1 except that the areas of interest are the AI objects detected by *ETVision* rather than areas of interest specified in *ETAnalysis*; and the fixation variables are "Live Fixations" detected by *ETVision* in real-time, rather than fixations computed off-line by *ETAnalysis*.

To generate a bar plot for a single event, simply right click on an Event node, and select "Display AI Bar Plots" from the drop down context menu. A bar plot will appear on an "AI Bar Plots" tab.

If multiple events have the same set of AI detected objects (objects with the same names and indicies), it is possible to generate a bar plot that shows the average data across these events. Select "Group "Display AI Bar Plots" from the "Group" pull down menu on the *ETAnslysis* menu bar, and then select the event nodes to be included from the resulting pop-up dialog. A bar plot will appear on a "Group AI Bar Plot" tab.

Selection of the variable to be plotted for each AI object, and other plot features are the same as for AOI bar plots described in section 14.3.1.

14.4 Superimpose moving Gaze and Fixation Trail over static backgrounds

A static background image associated with an event can be viewed with a dynamic gaze trail and/or head map display from any event node. The gaze display will progress in time as though the data were being viewed "live". Similarly, fixations can be dynamically displayed on the background from fixation nodes; and areas of interest, along with a time line display showing the areas visited, can be displayed with the background from fixation sequence nodes.

As for static two-dimensional plots (see section 14.2), the moving gaze trail will make most sense when the gaze data specifies gaze with respect to a surface in the environment. This is the case for eye trackers that use remote (table mounted) optics, or for head mounted eye trackers like Argus Science *ETVision* when using the *ET3Space* feature or the *StimTrac* feature (see sections 17 and 20).

From an **Event** node, right click to see the context menu and select "Play Gaze over background".

The video tab will open in the Display Area. The "Configure Display" pull down menu provides check boxes to enable "Draw Options" and "Plot Options" dialogs, and to enable a pupil diameter plot with moving time bar below the main display. The moving time bar (red line that moves from left to right) indicates the current position on the pupil diameter chart. (In the case of *ETVision*, left eye pupil diameter is displayed on the dynamic plot). If enabled, the "Draw Options" tab allows selection of the information to show over background image. Check "Gaze Trial", "Heat Maps", or both. A dialog window for adjusting the length and color of the gaze trail can be brought up by clicking the "Gaze Trail" Configuration button to the right of the "Gaze Trail" check box, and a dialog for heat map properties can be brought up by clicking the "Heat/Peak Map" Configuration… button to the right of the "Heat/Peak Map" check box. The "Plot Options" tab, if enabled, allows adjustments to the pupil diameter display.

The **Configure Heat Map** dialog is similar to the heat map configuration dialog for static displays (see section 14.2.1). In this case, the dialog includes a section labeled "Number of Points to Include". If "All Points" is selected, then the heat map, at any time point, will be based on the gaze data samples from the beginning of the event up until the current time point. If a "Number Limit" is specified, then at any time point the heat map will be based on just the gaze data for the current time point and the n-1 previous data samples (where n is the "Number Limit" selected).

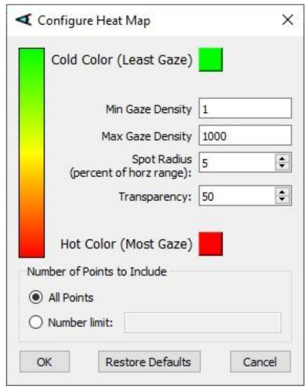

Other specifications in the **Configure Heat Map** dialog operate as described in section 14.2.1. Note that "Max Gaze Density" values will have no affect unless smaller than the number of gaze samples being considered, as described in the previous paragraph. (E.g. if "All points were selected, "Max Gaze Density" values will have affect only when less than the number of data samples from the beginning of the event up to the current time).

The Viewer display includes the usual controls for "play", "pause", step forward or back, and advance to specified frame at the bottom left of the image window. The step forward and back buttons advance or backup by the number of data records specified in the box between the buttons. A playback speed control is at the bottom right of the image window. The number displayed at horizontal center, just under the image window, is the data record number, with total number of records displayed after the slash ("/"). To display a specified data record, enter the desired record number and click the "go to" symbol.

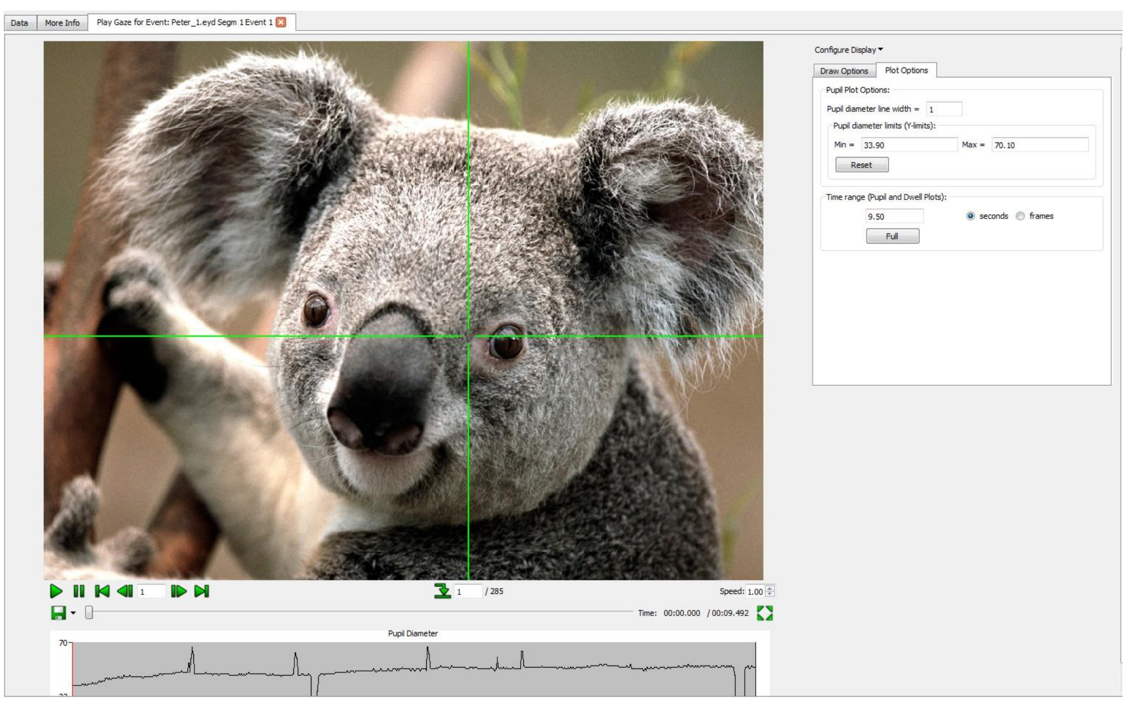

Note that when played at real-time speed (speed = 1.0) the display will update at the eye tracker data update rate. In cases where the gaze data update rate is faster than the PC monitor refresh rate, not all frames will actually display when speed = 1.0; but if played at slower speed or single frame update all frames will display. A slide bar, below the image window, can also be dragged to advance or back up through the video.

There is also a "record" ( ) pull down menu at the lower left of the main display with selections to record the video display as an wmv file, or capture the current frame as bit map image. The wmv file selection brings up a dialog that allows some choice of update rate and resolution. Update rate defaults to the gaze data update rate (E.g., 180 Hz for *ETVision* data), but lower rates can also be specified. Resolution defaults to the resolution of the background image, but also allows selection of the half this resolution.

A "full screen" button at the lower right will toggle to a full screen display that omits the tree diagram, or back to the standard display.

"Play Fixations over background" can be selected from a **Fixation** node. This produces much the same display as the previously described "Play Gaze over Background", accept that fixations can also be displayed. A "Fixations" check box is added to the "Draw Options" list, and a Configure Fixation display button allows adjustment of fixation display parameters.

"Play Statistics over background" can be selected from a **Fixation Sequence** node. This produces display similar to those previously described with a couple of additions. The static Areas of Interest are displayed, and under the "Configure Display" pull down menu, "AOI Dwell Plots" can be selected in place of "Pupil Diameter". In this case the plot below the background image display shows a bar for each AOI with the AOI color, indicating periods during which gaze was in that AOI. A moving time bar (vertical red line that moves from left to right) shows current position on the plot.

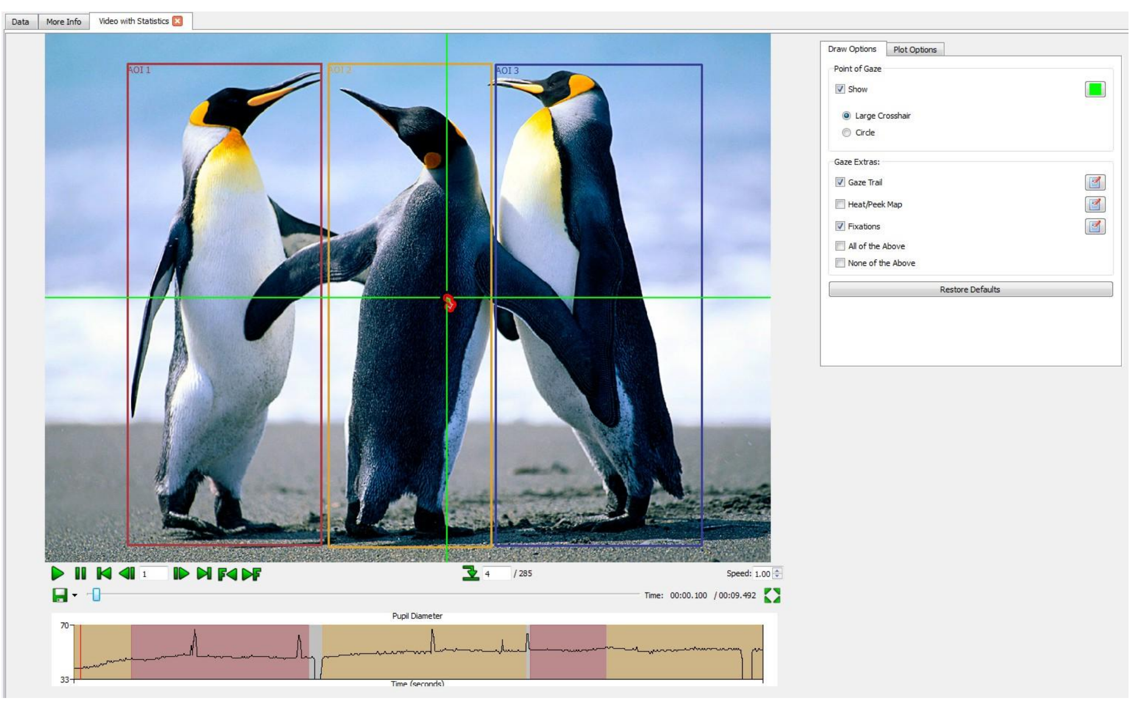

15 Combine data across events

Under "Data Files" the project tree branches out to individual data files, each divided into segments, and subdivided into events. The events often correspond to trials in an experiment. The Group menu on the *ETAnalysis* program allows data across different events to be combined in several ways.

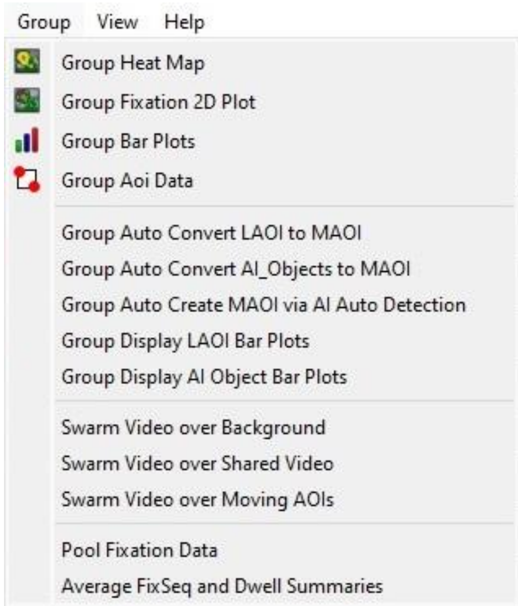

The first 3 selections have been discussed in previous sections (see sections 14.2.1, 14.2.2, and 14.3.1). Note that, as discussed in section 14.2, the 2D fixation plot and heat map usually make the most sense if using *ET3Space* or *StimTrac* so that gaze coordinates for each event represent a position on a surface in the environment (rather than a position on the head mounted scene camera image). See sections 17 and 20 for descriptions of *StimTrac* and *ET3Space*.

"Group Aoi Data" provides a way to display the summary fixation sequence statistics of a single AOI for multiple events. This has been described in section 11.2.

The next 3 selections are automated methods to create Moving Areas of Interest (MAOIs) for multiple events. MAOIs are explained in section 16. Group Auto Convert LAOI to MAOI is discussed in section 16.7.10. Group Auto Convert AI Objects to MAOI is discussed in section 16.7.11. Group Auto Create MAOI via AI Auto Detection is discussed in section 16.7.9.

Group LAOI and AI bar plots are discussed in sections 14.3.2 and 14.3.3.

The Swarm displays are discussed in the next section (15.1).

Statistics can also computed across multiple events in two different ways. One method is to pool all of the fixation data for all selected events, and to compute statistics for the pooled data. The other method is to take the average of the statistical quantities created for the individual events. These two choices are discussed in section 15.2.

15.1 Swarm displays

The Swarm displays are video displays that simultaneously show gaze sequence from multiple events. The gaze point from each event is displayed as a different colored dot, and when data from many events are combined, it looks like a "swarm of bees" flying over the background image, and can provide a visual illustration of whether all subjects followed a similar gaze pattern (dots stay tightly grouped) or a variety of different patterns (dots tend to spread out over display).

There are 3 Swarm display selections under the group menu:, "Swarm Video over Background", "Swarm Video over Shared Video" and "Swarm Video over Moving AOIs".

The first two choices usually make sense only if data has been collected with remote (table mounted) optics, or the head mounted *ETVision* eye tracker when using the *ET3Space* feature or the *StimTrac* feature (see sections 17 and 20). **Swarm Video over Background** can be used to simultaneously display gaze from multiple events over a static background image that corresponds to an *ET3Space* scene plane or to the *StimTrac* display surface, and can create a text file of gaze coordinates with respect to the static image coordinate frame. S**warm Video over Shared Video** can simultaneously display gaze from multiple events over a shared video watched by the participant in each event. Details are in sections 15.1.1 and 15.1.2.

For *ETVision* data not using *ET3Space* or *StimTrac* features, Moving Areas of Interest (MAOIs) can be used to create swarm displays. If multiple events show gaze interaction with the same set of Moving Areas of Interest (MAOIs), **Swarm Video over Moving AOIs** can be used to simultaneously show gaze from all of these events moving over a static image that includes the same set of areas. It can also create a text file of gaze coordinates with respect to the static image coordinate frame. Details are in section 15.1.3.

15.1.1 Swarm Video over Background

The "Swarm Video over Background" display is useful when multiple events show gaze over the same image. The swarm display shows point of gaze for each selected event as a different colored dot which moves about over a static background. When data from many events are combined, it looks like a "swarm of bees" flying over the background image, and can provide a visual illustration of whether all subjects followed a similar gaze pattern (dots stay tightly grouped) or a variety of different patterns (dots tend to spread out over display).

As previously discussed for static 2 dimensional plots, this display will make most sense when the gaze data specifies gaze with respect to a surface in the environment. This is the case for eye trackers that use remote (table mounted) optics, or for head mounted eye trackers like Argus Science *ETVision* when using the *ET3Space* feature or the *StimTrac* feature. For *ETVison* data not using *ET3Space* or *StimTrac*, see section 15.1.3.

Before Selecting "Swarm Video over Background" be sure to have an appropriate static background configured (see section 7). Selecting "Swarm Video over Background" first brings up a selection chart labeled "Select event nodes that share the same background".

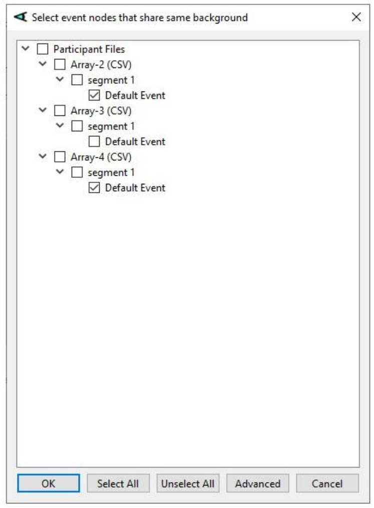

The user can select the events by checking individual event boxes, or higher-level boxes. Checking Segment node, checks all of the events under it. Checking the Participants node selects all events in the project, etc. There are also "Select all" and "Unselect all" buttons. The Advanced button brings up a dialog that allows selection of events based on various criteria.

Clicking OK brings up a viewer similar to that previously described in section 14.4.

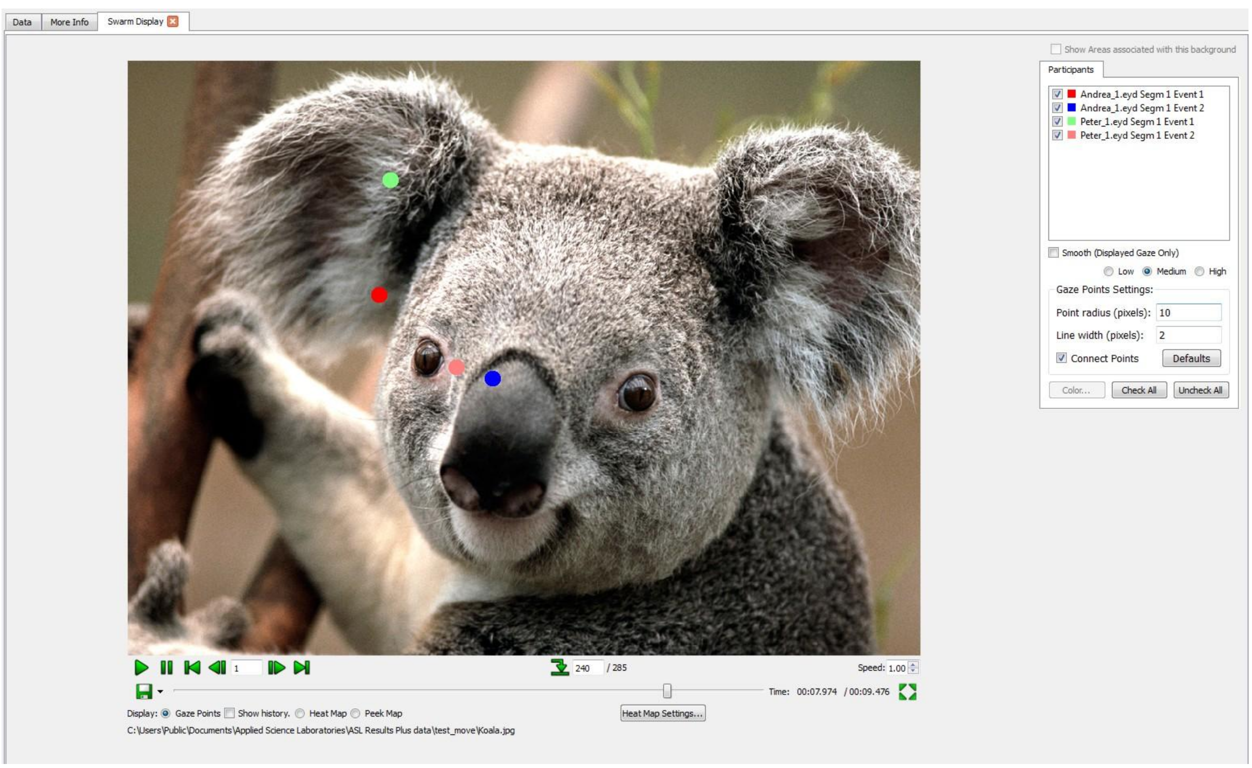

A key on the right, labeled "Participants", shows the color assigned to each event being displayed. Un-checking the box next to one of the event labels will cause the data for that event to <u>not</u> be displayed.

A smoothing filter can be applied to the data by checking the “Smooth” checkbox, just under the “Participants” key. Three levels of smoothing are available, as determined by the low, medium, and high radio buttons. Smoothing will apply only to the display. It will not change the data table or computed statistics.

Use the “Gaze Point Setting” controls, below the “Smooth” selection, to adjust the size of the gaze dots and the width of connecting lines (if “Show History” and “Connect Points” are both selected).

Use the radio buttons and check boxes below the video controls to select the display type. The “Show History” check box applies only to the “Gaze Point” selection. If not checked, each frame will show only the gaze positions for that data frame. If “Show History” is checked, all previous gaze points will also be shown. In other words, once a gaze point is displayed it will remain as subsequent points are added to the display. A Heat map or peek map, rather than gaze points, can be shown for each subject by selecting the corresponding radio button. Note that when Heat Map or Peek Map is selected, it becomes impossible to distinguish between the data from the different events.

The Viewer display includes the usual controls for “play”, “pause”, single step forward or back, advance to specified frame, and playback speed. A slide bar can be dragged to advance or back up through the video. There is also a pull down menu at the lower left of the main display with selections to record the video display as a wmv file, or capture the current frame as bit map image. A “full screen” button at the lower right will toggle to a full screen display that omits the project tree diagram, or back to the standard display.

The “record” ( ) pull down menu includes a selection labeled “Save Gaze”. This selection allows the user to record a text file with a list of all gaze coordinates with respect to the background image pixel coordinate frame.

```
0,Array-2.csv_Segm_1_Event_0,
1,Array-3.csv_Segm_1_Event_0,

sample_#,time_secs,GazeX0,GazeY0,GazeX1,GazeY1,
1,0.000,313.365,225.240,323.365,235.240,
2,0.006,313.210,225.075,323.210,235.075,
3,0.011,312.845,225.360,324.845,237.360,
```

A short sample from such a file is shown above. It begins with a list of events, In this case 2 events numbered 0 and 1. Each event is listed on a separate row with the file name, segment number, and event number. Next is a row of comma separated column labels, followed by a row of comma separated values for each data sample. Time values are from the beginning of each event. The horizontal and vertical gaze positions are labeled “GazeX0” and “GazeY0” for data from the first event, and “GazeX1” and “GazeY1” for the second event. The units for the gaze values are background image pixel values. If, for example, the background used is a 1280 x 720 pixel image, a gaze point at the center would have X,Y coordinates (640, 360).

15.1.2 Swarm Video over Shared Video

The “Swarm Video over Shared Video” is similar to “Swarm Video over Background”, but showing gaze data with respect to video scene images rather than static background images. It is useful when multiple events show gaze while watching the same video presentation. (Use of Video scene images and are discussed in section 16.) As for Swarm Video over Static Background, this display will make sense for *ETVision* data only if *ET3Space* or if *StimTrac* is used (see section 1.2).

Before Selecting "Swarm Video over Shared Video" be sure to have an appropriate shared video configured for all events to be inclded. If data was gathered using *ETVision* with *ET3Space* (section 20), use "Configure Video Data" selection from a file node, or each relevant Segment node or event node to configure the shared video file, as described in section 16.1. If using *ETVision* data with the *Stimulus Tracking* feature (section 17), use the "Configure Stimulus" selection from each relevant "Monitor" node to configure the shared video file, as described in section 17.5.

Selecting "Swarm Video over Shared Video" first brings up a selection chart labeled "Select event nodes that share same video". Select the events to be included, just as described in the previous section for "Swarm Video over Background".

Clicking OK will bring up a viewer just like the described in the previous section except that the background will be the shared video rather than a static image.

"Swarm Video over Shared Video" is also discussed in section 16.10.1.

15.1.3 Swarm Video over Moving AOIs

A Swarm display can be created for *ETVision* data that did not use *ET3Space* or *StimTrac* by creating Moving Areas of Interest (MAOIs) on the scene camera videos (see section 16 for description of Moving Areas of Interest). If multiple events show gaze interaction with the same set of MAOIs, a swarm video can be created which simultaneously shows gaze from all of these events moving over a static image that includes the same set of areas. A fixation sequence computation must first be executed for each event. The swarm display will show only gaze points that were within one of the Moving Areas of Interest.

As an example, assume that data has been recorded from several participants as they walked about in a room containing several objects. Further assume that, on each scene video recording, these objects have been tracked, by *ETAnalysis*, with identically named MAOIs. Although the scene video recorded for each participant is different, if we make a still picture showing the room with all of the objects in view, we can designate the same areas of interest on this image and make a swarm video showing gaze from all the participants simultaneously moving about on this static image. *Note: only gaze points falling within an MAOI will be shown*.

The procedure and resulting viewer for "Swarm Video over Moving AOIs" are described in more detail in section 16.10.2.

15.2 Computing Statistics across multiple events

In most projects the data from each event (or trial) is processed to compute fixations, and these are compared to areas of interest to compute various statistics. To analyze data across trials we may either want to pool (combine) the fixation data from multiple events (or trials) and compute statistics with the pooled data, or we may just want to pool the statistics that were computed for each event and average some of these quantities, or manipulate them in some other way. Both of these can be done from the Group menu on the *ETAnalysis* program. These are discussed in more detail in sections15.2.1 and 15.2.2.

15.2.1 Pool Fixation Data

From the main menu select **Group → Pool Fixation Data.** A selection dialog will appear, in the form of a tree diagram, showing all of the fixation nodes in the project.

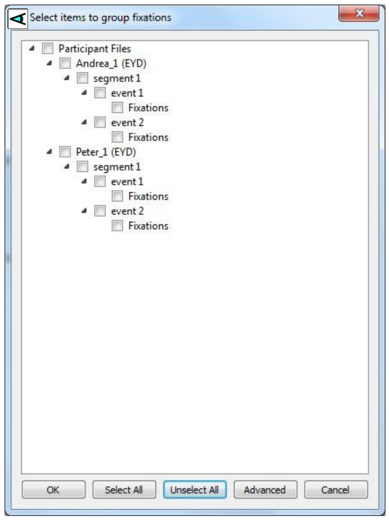

Select the events to be pooled, as described in the previous section (if desired, use the "Advanced" dialog to select events with certain names, certain XDAT flags, or certain associated backgrounds), and Click OK. The program will create a top level node called **Pooled Fixation Data**, and a new branch under this node, labeled **FixationGroup_1,** containing the pooled fixation data specified. Note that in addition to the data columns in the individual event fixation list, there are now also columns for the Filename, Segment ("Segm"), and Event number ("Evt"). These additional columns completely identify the source of the data in each row.

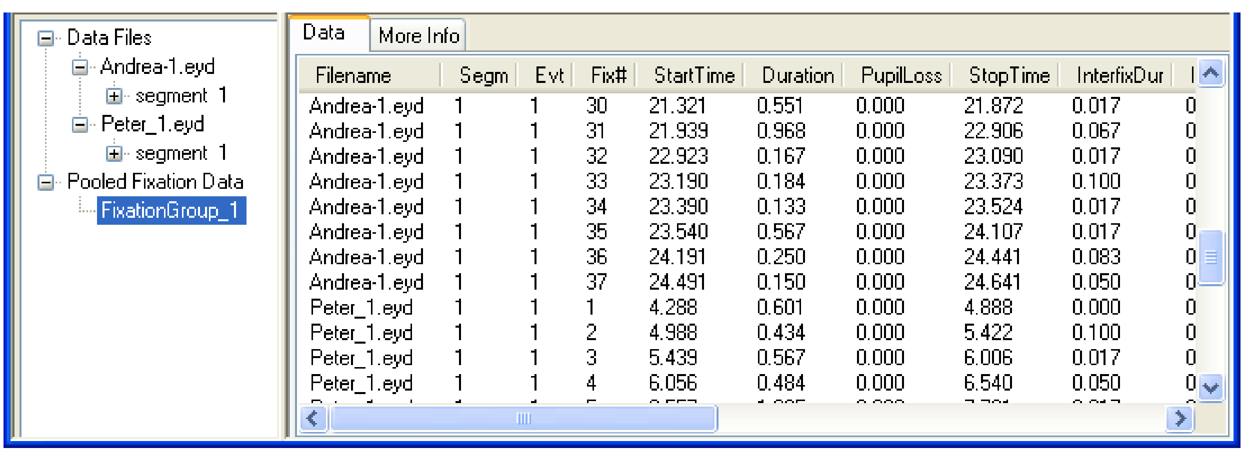

Filename	Segm	Evt	Fix#	StartTime	Duration	PupilLoss	StopTime	InterfixDur	I
Andrea-1.eyd	1	1	30	21.321	0.551	0.000	21.872	0.017	0
Andrea-1.eyd	1	1	31	21.939	0.968	0.000	22.906	0.067	0
Andrea-1.eyd	1	1	32	22.923	0.167	0.000	23.090	0.017	0
Andrea-1.eyd	1	1	33	23.190	0.184	0.000	23.373	0.100	0
Andrea-1.eyd	1	1	34	23.390	0.133	0.000	23.524	0.017	0
Andrea-1.eyd	1	1	35	23.540	0.567	0.000	24.107	0.017	0
Andrea-1.eyd	1	1	36	24.191	0.250	0.000	24.441	0.083	0
Andrea-1.eyd	1	1	37	24.491	0.150	0.000	24.641	0.050	0
Peter_1.eyd	1	1	1	4.288	0.601	0.000	4.888	0.000	0
Peter_1.eyd	1	1	2	4.988	0.434	0.000	5.422	0.100	0
Peter_1.eyd	1	1	3	5.439	0.567	0.000	6.006	0.017	0
Peter_1.eyd	1	1	4	6.056	0.484	0.000	6.540	0.050	0

15.2.2 Average Fixation Sequence and Dwell Summaries

As an alternative to (or in addition to) pooling the fixation data events as described in the previous section, it is also possible to take the AOI summary data from the original Fixation Sequence and Dwell segments, and to average each item in this summary data.

To look at why this might be sometimes be useful consider the following simple example. Suppose we have recorded basketball free throw results for 3 players, and have observed 10 free throws for players 1 and 2, and 50 free throws for player 3. If we pool the data and calculate the percentage of successful attempts out of the 70 total free throws we have observed, the result will be weighted heavily to reflect the skill of player 3. If we want to estimate what the team average will be when they each shoot an equal number of free throws we might be better off computing the percentage separately for each player, and then averaging those results. (Determining the most appropriate way to combine data from different population groups is actually a significant statistical problem too complex to address here).

WARNING: it will usually make sense to do this only across fixation sequence and dwell sets that use the same AOI set, or at least AOI sets with the same number of areas and areas that have the same name and meaning. For example, if subjects looked at pictures of faces, there may be a different AOI set for each face. It may still make sense to average across data from different faces if, in all AOI sets there are a set of areas corresponding the same facial features (E.g. left eye, right eye, nose, mouth).

To collect and average fixation sequence and dwell statistics across events, proceed as follows.

From the main menu select **Group → Average FixSeq and Dwell Summaries.** Select the Fixation Sequence nodes (if desired use the advanced selection criteria as described for Swarm displays in section 15.1) and click OK.

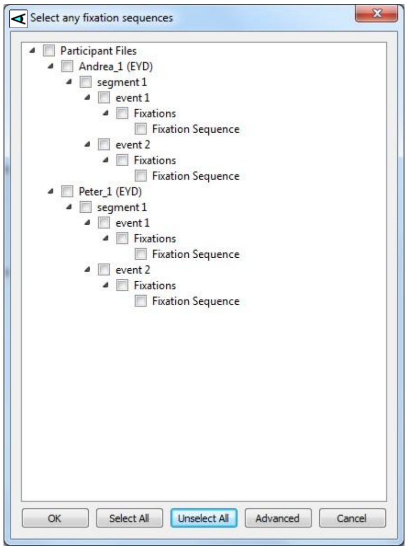

The program will create a new top-level node called **Summary Averages**, and a new branch under this node called **FixSeqDwellGroup_1**.

The sub branch labeled **Fixation Sequence Group Average** will contain combined AOI summary data, with each row showing the file, segment and event from which the summary statistics are taken. Under these nodes will be **AOI Summary Average** (shown below), **Transition Table Average**, and **Joint** and **Conditional Probability Average** tables. The **AOI Summary Average** provides the mean and, in some cases, standard deviations, across the different fixation sequence and dwell sets selected, for each item in the **AOI Summary** tables. For example, the first item in the AOI Summary Average, after the AOI name, is "FixCount". For the first row this will be the average of the "FixCount" items in all the AOI 0 rows from the table shown above. The next item, labeled "STD_FC" will be the corresponding standard deviation.

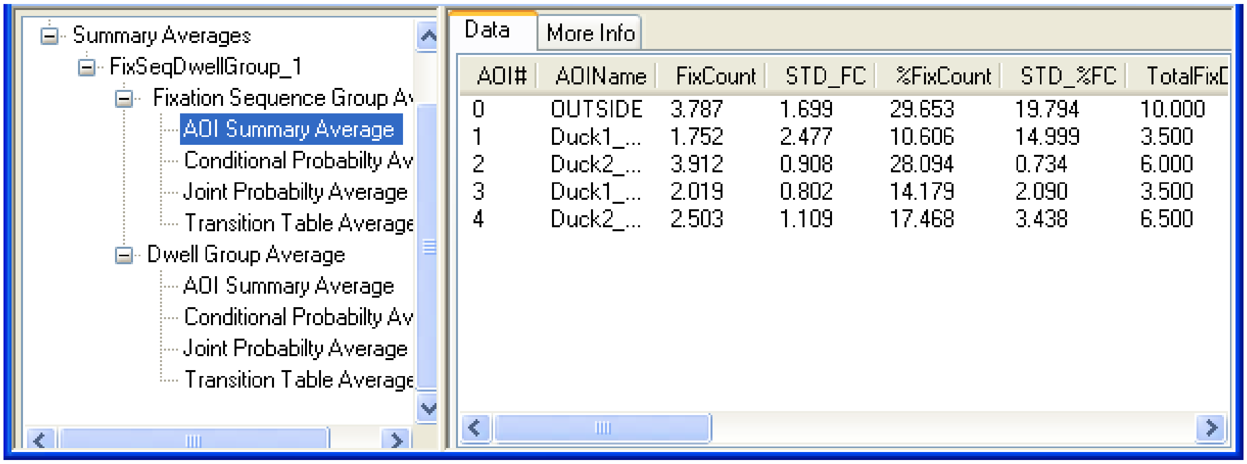

Each cell in the **Transition Table Average** is the average of the corresponding cells from the transition table in each included fixation sequence event. Similarly each cell in the Joint and Conditional probability tables are averages of corresponding cells from all included events.

16 Working with Scene video files and Moving Areas of Interest

If the project type has “Stimulus Type” set to “Videos” or “Both” (see section 4.1), *ETAnalysis* allows the user to create moving areas of interest over scene video files and to analyze gaze data with respect to these moving areas. All of the fixation sequence and dwell analysis statistics available with stationary areas of interest are available with respect to moving areas defined on scene video recordings. In addition, the scene video can be played with fixations, gaze trail, and heat map displays superimposed, and recorded, in this form, as a new wmv type video file. Gaze data may be any data type recorded by *ETVision* eye tracking systems.

To work with scene video recordings, *ETAnalysis* must have the means to properly match digital gaze data with corresponding video frames and to scale gaze data with respect to the video images. If scene video is recorded by the eye tracker computer, usually from the head mounted scene camera or from an external camera using the *ETVision* **Stationary Scene Camera** feature, temporal synchronization is handled automatically.

If gaze was recorded as subjects watched the playback of some video file, then it may be possible to use this same file in *ETAnalysis*, but only if arrangements were made to start gaze recording, or to set a specific XDAT value, on the eye tracker at same time that the video presentation to the subject began; and to stop gaze data recording, or set a specific XDAT value, on the eye tracker at the same time the video file presentation ended. In other words either the beginning and end of the gaze data file (eyd, or ehd file), or XDAT marks on the gaze data file, must correspond to the beginning and end of the video file. If XDAT values on the data mark the beginning and end of periods that correspond to the beginning and end of video files, use XDAT to divide (parse) the data into “events” that correspond to these periods. See section 6 for instructions on “Parsing” data in *ETAnalysis*.

The next two sections describe the general procedure for “Configuring Video Data”. Once video data is configured, gaze data can be superimposed, moving areas of interest (MAOIs) can be assigned, and statistics computed. Bar plots, as described in section 14.3, can be created for moving areas of interest just as for static AOIs.

If a single video file with MAOIs is to be shared by multiple Segments or Events, then an “Environment” Node containing the video file can be created. In this case scaling for the video file and the MAOIs are specified under the Environment node. This “Environment” video, with MAOIs can then be associated with multiple Segments and/or events. This procedure is described in section 16.3. An environment video can also be used to implement a “batch” process for auto-detecting MAOIs on individual event nodes (sections 16.5 and 16.7.7).

If a particular Video file is associated with only one Segment or Event, it can be completely configured at the Segment or Event node level and Moving Areas of Interest can be specified for configured video files at individual event nodes. In this case the MAOIs apply only to that event (section 16.4).

Example 1: *ETVision* data will be analyzed with respect to the recorded head mounted scene camera video. In this case, the scene camera video data will be associated with the data by default and it will not be necessary to use the **Configure Video Data** dialog.

Example 2: *ET3Space* data has been gathered with a video screen as a scene plane, and the video display has also been recorded as an "SSC" video file using the *ETVision* **Stationary Scene Camera** feature. Data can be analyzed with respect to this video file rather than the scene camera video, but the **Configure Video Data** dialog must be used to associate the video file with the desired *ET3space* data and to specify the data scaling needed to match gaze data with the video file. If it is later decided to re-analyze the same data event with respect to the *ETVision* scene camera video data, the **Configure Video Data** dialog will need to be used to re-associate the data with the scene camera video file and to select "Standard" scaling for the video.

Example 3: *ET3Space* data has been gathered in an environment with multiple scene plane surfaces. All or some of these surfaces are in view of an external, stationary video camera, and the video from this external camera has been recorded as an "SSC" video file using the *ETVision* **Stationary Scene Camera** (SSC) feature. Data can be analyzed with respect to this video file rather than the scene camera video, but as in the previous example, the **Configure Video Data** dialog must be used to associate the video file with the desired *ET3space* data and to specify the location and scaling for each scene plane within the video field of view. Note that this works only if the external video camera was stationary with respect to the *ET3Space* global coordinate system.

Example 4: As in Example 2, *ET3Space* data has been gathered with a video screen as a scene plane, and the video display has also been recorded as a video file. Unlike Example 2, the external video has not been recorded using the *ETVision* SSC feature, but the recording started and stopped at the same time as *ET3Space* data recording, or *ET3Space* data was marked by an external data (XDAT) value change when the video recording started and ended. As in Example 2, the **Configure Video Data** dialog must be used to associate this video file with the desired *ET3space* data and to specify the data scaling needed to match gaze data with the video file.

Example 5: As in Example 3, *ET3Space* data has been gathered in an environment with multiple scene plane surfaces, and all or some of these surfaces are in view of an external, stationary video camera. Unlike example 3, the external video has not been recorded using the *ETVision* SSC feature, but the recording started and stopped at the same time as *ET3Space* data recording, or *ET3Space* data was marked by an external data (XDAT) value change when the video recording started and ended. As in Example3, the **Configure Video Data** dialog must be used to associate the video file with the desired *ET3space* data and to specify the location and scaling for each scene plane within the video field of view.

16.1 Using the Configure Video Data dialog

If data was recorded with *ETVision*, the head mounted scene camera video is associated with the data by default and properly scaled by default. In this case it is not necessary to use the **Configure Video Data** dialog unless a different video source (not the head mounted scene camera video) is to be associated with the data.

If data will be analyzed with respect to a video source other than the *ETVision* scene camera video (see examples 2 through 5 in previous section) then the video does need to be "configured".

The **Configure Video Data** dialog is available from the context menu under all nodes down to the event level. It is used to associate a specified video file with events that are below the node chosen; to specify whether the video file beginning and end correspond to the beginning and end of the segment or the beginning and end of each event; and to specify the data scaling needed to match gaze data with the video file.

Unlike static backgrounds, there is not a correspondence table to associate configured videos with events. If different video files correspond to different events, the Configure Video Data dialog must be used individually on each.

Bring up the **Configure Video Data** dialog by right clicking the appropriate node and selecting "Configure Video Data" from the context menu.

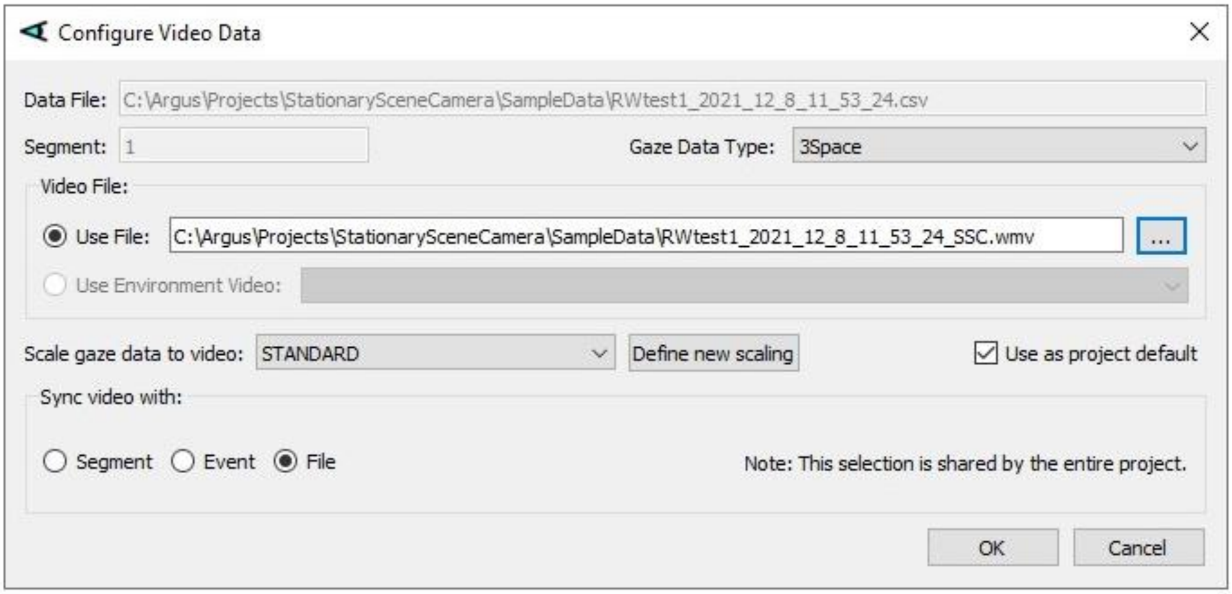

16.1.1 Gaze Data Type

If the data file being analyzed does <u>not</u> contain *ET3Space* data, the field labeled "Gaze data type", at the upper right on the dialog, will be gray, and can be ignored. If the data file <u>does</u> include *ET3Space* data (section 20), the *ET3Space* data (gaze position on stationary scene plane surfaces) can be superimposed on video from a room-fixed source, as described by examples 2 though 5. In these cases, "Gaze data type" should be set to "ET3Space".

ET3Space data files do still include data specifying gaze position with respect to the head-mounted scene camera image (computed without the help of the *ET3Space* feature), and to view this non-*ET3Space* data superimposed on the head-mounted scene camera video, "Gaze data type" can be set to "Scene Image".

16.1.2 Video File

Set the radio button to "Use File" and browse to the video file. Alternately select a video file that has already been loaded as an **Environment** node, by setting the radio button to "Use Environment Video" and using the pull down menu to select the file. (The procedure for creating an "Environment Video" is explained in section 16.3).

If *ET3Space* data is to be displayed with respect to a video recorded using the *ETVision* **Stationary Scene Camera** feature (Examples 2 and 3), select the SSC video. The SSC video file will have the same file name as the gaze data file, but with "SSC" appended to the end of the name (as shown in the sample screen shot on the previous page).

16.1.3 Sync video with Event, Segment, or File

If re-associating *ETVision* scene camera video with the data, set "Sync video with:" to "File".

If using a **Stationary Scene Camera** (SSC) video file, set "Sync video with:" to "File".

If using an external video not created with the SSC feature, it will be necessary to know whether the video started and stopped at the same times as the data file start and stop, the data segment start and stop, or the data event start and stop point. *Note that "csv" type files can have only one "Segment", and in this case "File" and "Segment" are the same. Binary "eyd" or "ehd" type files can have multiple "Segments" (multiple recording starts and stops) on the same file. Any data segment can be parsed, by ETAnalysis, into multiple "events".*

In the case of "Example 4" or "Example 5", with an external video <u>not</u> created with the SSC feature, if the recorded video was started and stopped at the same time as *ET3Space* data recording was started and stopped, then "Configure Video" data should be selected at the Segment node, and "Sync video with:" should be set to "Segment". If the recorded video was started after *ET3Space* data recording, and XDAT values were used to mark the *ET3Space* data with the video start and stop times, then these XDAT value changes should be used to define an event in *ETAnalysis* (see instructions for "parsing events" in section 6). In this case, "Configure Video" data should be selected at the Event node, and "Sync video with:" should be set to "Event".

16.1.4 Scale gaze data to video

If re-associating *ETVision* scene camera video with the data, set "Scale gaze data to video:" to "STANDARD". In the case of *ETVision* scene camera video, "STANDARD" data scaling is from 0,0 at the upper left to H=1280, V=720 at the lower right of the video image.

In other cases, such as Examples 2 – 5 data scaling for the video is determined by capturing a video frame, as a still image, and specifying the gaze data coordinates that correspond to four visible landmarks in the image.

In the case of Examples 2 and 4, it is usually easiest to use the 4 corners of the video image, and the *ET3Space* coordinates corresponding to the positions of these image corners on the *ET3Space* scene plane. In the case of Examples 3 and 5, four attachment points need to be specified <u>for each *ET3Space* scene plane</u> visible in the video.

To compute new scaling data for the current video, click "Define new scaling", and proceed as described in section 16.2.

Once the scaling procedure is complete, the scaling data is saved with a name composed of the video file name and frame number used for the still image. All such scaling data files that are currently part of the project will be available from the "Scale gaze data to video" pull down menu. To use the same scaling previously computed for the same video file (or for a video file of the same type same resolution, and displayed by the same application as the currently selected video) select the scaling file from the pull down menu.

16.2 Scaling Data to Video Files

To define scaling for an Environment video (see section 16.3), select "Configure Attachment Points for Gaze" from the Environment Video node context menu. To define scaling for a video selected in the **Configure Video Data** dialog (described in the previous section), click the "Define new scaling" button.

A window will appear showing the first frame of the video.

The slider can be used to advance to any desired video frame. It will be necessary to identify 4 landmarks (recognizable image features) as near as possible to the corners of the display. The corners of the image can be used if these were visible to the subject (and if the entire image was on a single scene plane, if using *ET3Space* data. In this case any video frame is OK. Click the "Capture Frame" button to produce a bitmap image of the frame. The file name with an appended frame number will be listed as "Current Background:". An **Add/Edit Attachment Points** dialog will also appear.

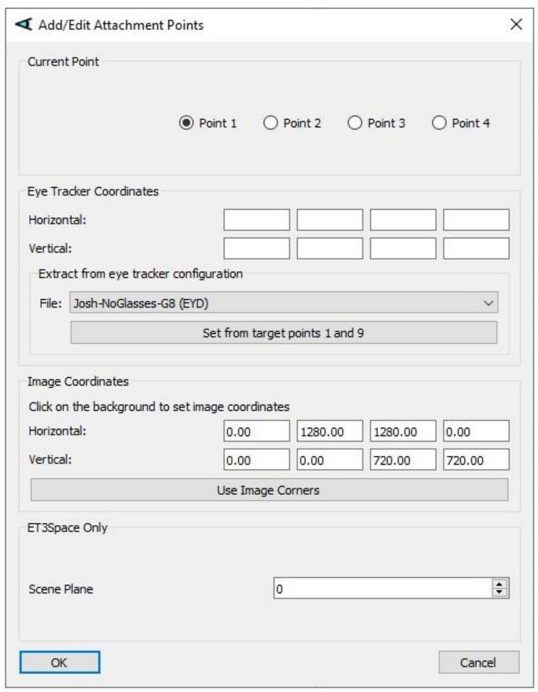

From this point, the procedure for specifying attachment points is the same as that for specifying attachment points on static background images (section 7).

Find the eye tracker coordinates associated with the landmarks chosen. If *ET3Space* was used, use the "pointer test" function, or measure to find the *ET3Space* coordinates associated with any point in the scene image. See *ET3Space* manual for details. Note that this can be done in advance, with coordinates recorded for use at this point in the analysis process. If using landmarks that move as the video advances, be sure to measure the eyetracker coordinates using the same video frame as that selected here. If using *ET3Space* data, remember that a set of 4 attachment points must be specified for each scene plane visible in the video image.

On the **Add/Edit Attachment Points** dialog, type in the "Eyetracker Coordinates" for each of the landmark points chosen. Ignore the "Scene Plane" number field unless using *ET3Space* data. If using *ET3space* data, remember to specify the scene plane number for the set of attachment points being entered.

If the corners of the video image are the landmark points, click the "Use Image Corners" button. Otherwise, use the mouse to click on the corresponding points in the image. A red dot with the point number label ("P1", "P2", etc.) should appear at each point, and the pixel coordinates of the point will appear in the "Image Coordinates" section.

If using *ET3space* data, repeat the procedure (enter Eyetracker Coordinates and click on image points) for each scene plane visible in the video image. As soon as attachment points for more than one plane are specified, the labels will also include the plane number ("P0,1", "P0,2", etc.).

Click OK to close the **Add/Edit Attachment Points** dialog.

Click "Save and Close" to close the video frame image.

16.3 Environment Video

16.3.1 Description and purpose

A special node in the project "tree" diagram can be created to hold video files that will be used by multiple events. These are called "environment videos".

If multiple subjects watch the same video display presentation, and data is recorded with a remote eye tracker or a head mounted system using the *ET3Space* feature, then moving areas of interest (MAOIs) can be created on a single "environment video" and applied to multiple events.

If multiple subjects, wearing a head mounted system, move about in the same environment, then even though each "event" will have a unique scene video, creation of a special "environment video" file can be useful. A specially created environment video file can allow the "Auto Detect MAOI" feature to be used in a batch processing mode to automatically create MAOIs for each individual video.

16.3.2 Creating an environment video node

To create an environment video click the Open Environment Video button on the short cut bar , or select **File → Open Environment Video**, and browse to the video file. An "Environments" node will appear with the video file as a sub-node.

Section 16.6 describes using an environment video to define MAOIs that will be shared by multiple events. Section 16.5 describes a procedure for using an environment video to implement the "Auto Detect MAOI" feature in a "batch" mode. This batch "Auto Detect MAOI" feature will automatically create moving AOIs for multiple events.

16.4 MAOIs for individual events

Once configured video data is associated with an event (see section 16.1) moving AOIs can be created for that event.

Follow instructions in section 16.7 to create an MAOI set. The MAOIs can be created by manually or by one of several automated techniques as described in section 16.7. When done be sure to click the "Save and Close" button near the top right of the Display Area.

A Moving Areas of Interest node will appear below the event node. Fixations can be analyzed with respect to the moving areas of interest, and dynamic gaze and fixation data can be superimposed on the video with moving areas of interest also shown (sections 16.8 and 16.9).

16.5 Using an Environment video to Automate MAOI creation

If multiple subjects move about in the same environment while wearing a head mounted *ETVision* system, then even though each "event" will have a unique scene video, a "batch" processing feature can be used to Auto Detect MAOIs in all the separate event videos. Record scene camera video as the *ETVision* optics are moved around the environment, and make this an "environment video" (section 16.3.2). Follow instructions in section 16.7.7 to designate MAOIs in the environment video, and to "auto detect" these areas in scene videos from multiple events.

16.6 Sharing MAOIs, with multiple Segments or events

The most efficient way to share a single video and attached MAOIs with multiple data event nodes is to first make the video file an Environment Video, and then create the MAOI set on the Environment Video. It is possible to share an Environment Video, including MAOIs, with multiple events. *Note that this usually makes sense only if all events comprise data recorded as subjects watched the* <u>*same*</u> *video presentation.* ***Furthermore, if using a head mounted eye tracker it will usually make sense only if the ET3Space feature or StimTrac feature was used***.

Follow instructions in section 16.3 to make an Environment Video. To add moving Areas of Interest (MAOIs) to the video, right click the video file node, under the "Environments" node, and select a method for creating MAOIs. See section 16.7 for instructions on creating or editing moving areas of interest. When done be sure to click the "Save and Close" button on the "Configure MAOIs…" tab.

If the beginning and end of the video file corresponds to the beginning and end of data Segments, select "Configure Video Data" from the context menu at a Segment node or whatever level above the segment node includes all events that will share the video and MAOIs. For example, to attach the video file with MAOIs to all files and all data Segments in the project, use the Participant File Node context menu. If Segments are divided into multiple events, each event will automatically use the appropriate section of the video and MAOI file.

If the video beginning and end correspond to the beginning and end of one or more events (rather than an entire segment), the event node must be used, and the following procedure will need to be repeated for each such event.

A "Configure Video Data" dialog will appear.

If the video file corresponds to the beginning and end of the entire data Segment, set the "Sync with" button to "Segment". If at an event node (the video file corresponds to the beginning and end of the event) set the "Sync with" button to "Event".

Set the "Video File" radio button to "Use Environment Video" and make sure the proper video file name is specified. Next to "Scale Gaze Data to Video:", the same file name and frame number label should appear as that shown as "Current Background:" when data scaling was specified for the environment video (see section 16.3). Click OK.

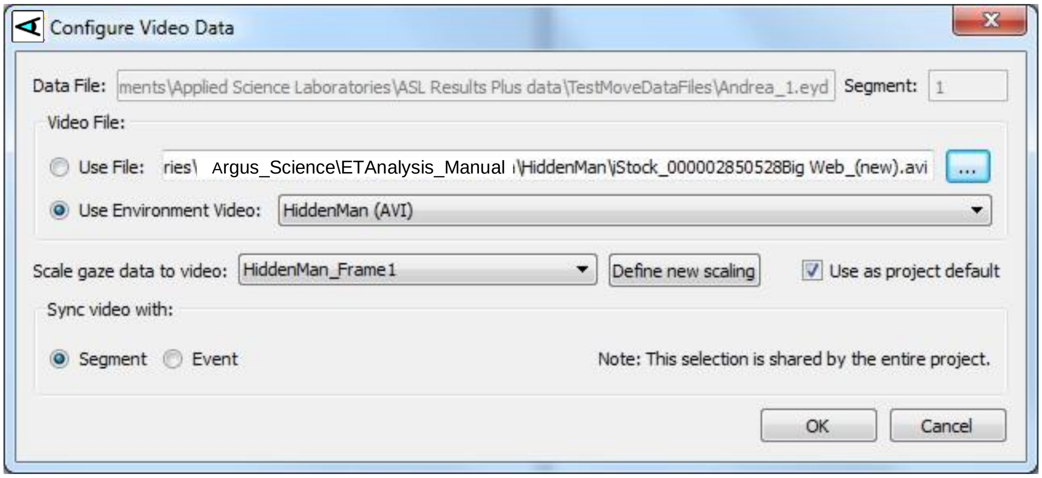

A "Moving Areas of Interest" node should now appear under each applicable event node. Fixations can be analyzed with respect to the moving areas of interest, and dynamic gaze and fixation data can be superimposed on the video with moving areas of interest also shown.

16.7 Methods for Creating Moving Areas of Interest (MAOIs)

Moving Areas of Interest are rectangles or polygons that form boundaries around image objects, and which move and change size and shape to appropriately to track those image objects through the video.

There are several possible methods for creating Moving Areas of Interest.

1. Draw areas on a single video frame and manually adjust the area boundaries to track image objects through the video. (Sections 16.7.1 - 16.7.4).
2. Draw areas on a single video frame and then use an "Auto Detect MAOI" function that uses feature detection to automatically track objects through the video, with manual correction as needed. (Section 16.7.5).
3. Automatically associate the size and position of a new MAOI with an existing MAOI. (Section 16.7.8).
4. Use an AI model in *ETAnalysis* to detect image objects and create MAOIs. (Section 16.7.9).
5. Convert "Live Areas of Interest" created by *ETVision* to MAOIs. (Section 16.7.10).
6. Convert Artificial Intelligence Objects ("AI Objects") detected in *ETVision* to MAOIs. (Section 16.7.11).

The first 3 methods allow creation of MAOIs that exactly conform the shape of the image object. Manually moving area outlines to track objects through the video can be very precise, but may be very labor intensive if there is a lot of motion and if the videos are long. Automating this process with the "Auto Detect MAOI" function (method 2) is very effective if the image objects have well defined edges and if the image content within area of interest boundaries remains constant. However, if feature detection does not work robustly with the image objects involved, frequent manual correction requirement may lead to a labor-intensive process.

Method 3 is very effective if an image object maintains a very consistent position and size relative to an existing, accurately defined MAOI.

Some significant effort may be involved in "training" a custom AI model to recognize particular set of objects, but once this is done AI object detection (method 4) is extremely robust. The AI model can be used with multiple videos and usually does not requiring manual intervention to accurately track the set of objects it has been "trained" for. The resulting MAOI outline for an image object is always a rectangular bounding box. It is not possible, with this method, to create polygons that exactly conform to the object outline.

If *ETVision* created "Live Areas of Interest" in real-time the data file contains the feature detection parameters used for each area. These can be used to implement automatic feature detection (method 5). This is essentially the same as method 2, but without the need to draw the areas on a video frame first.

Methods 2, through 5 all require the program to "play through" the scene video. Time required will depend on the capabilities of the computer being used, but will usually be similar to the real-time length of the scene video.

If *ETVision* used an AI model to detect "AI Objects" in real-time, the position and dimensions of each object bounding box is recorded on the data file and this information can be converted to MAOIs by *ETAnalysis* (method 6). The program does <u>not</u> need to "play" the scene video to do this conversion and the process usually takes only several seconds even for long scene videos.

Once MAOIs are created they behave in the same way no matter which of the above methods were used to create them, and they can be manually manipulated as described in section 16.7.4.

Note that if LAOI or AI Object bounding boxes were recorded on the scene image by ETVision these will interfere with automatic methods for detecting MAOIs in ETAnalysis (methods 2, 4 and 5). If these methods will be used in ETAnalysis, it is recommended that ETVision be set to not "Record LAOI / AI" (See section 9.4.2.2 of ETVision manual).

Detailed instructions for each of the methods discussed above are in the following subsections.

16.7.1 "Configure MAOI" dialog

Moving areas may be manually defined in a scene video corresponding to the eye tracker data file. This can be either a video from a head mounted scene camera, or a stimulus video that has been presented to multiple subjects. Therefore, the option to "Configure Moving Areas of Interest" can be found on the context menus for both event nodes and environment video nodes. The Configure Moving Areas of Interest function can also be selected by left clicking the moving AOIs button  when the relevant node is selected in the Project Tree.

The "Configuring MAOIs….." tab will appear in the Display Area showing the first frame of the video file corresponding to the selected node.

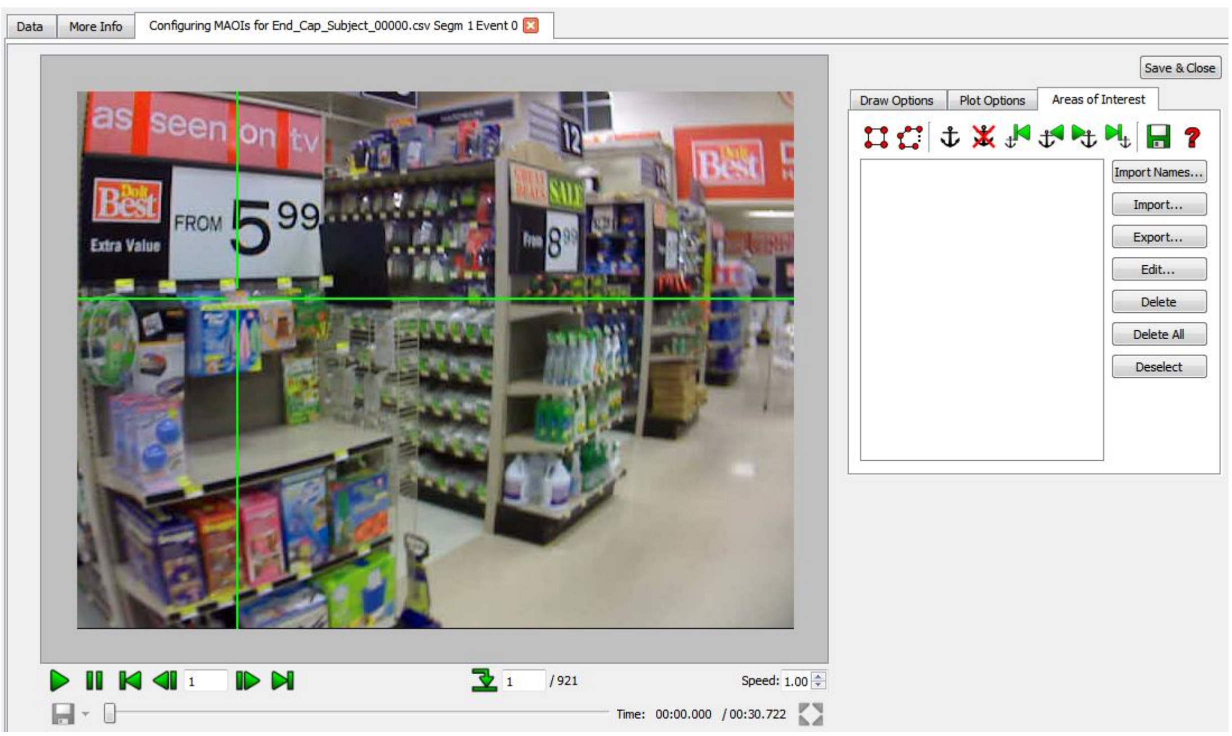

The "Configure MAOIs" tab has controls that allow the user to play the video, jump to any specified frame in the video, or to move forward or backwards by steps of any specified size. Areas of Interest can be drawn as either rectangles or polygons with any number of sides. If an area is created or moved on any video frame, that frame becomes an "anchor" point for that area. Movements of the entire polygon or of any individual vertices are interpolated between anchor points.

Initialize (draw) an AOI on the current frame by selecting the "Draw rectangular AOI" or "Draw polygon AOI" button.

Below is a chart that describes the buttons associated with moving AOIs found in the "Areas of Interest" tab to the right of the video display.

	Draw a rectangular AOI to manually manipulate.
	Draw a polygonal AOI to manually manipulate.
	Anchor all AOIs in the current frame.
	Remove anchor from all AOIs in the current frame (calculate AOI positions from surrounding anchors). Useful for undoing a manual change.
	Go to the first frame in which any AOI has been drawn.
	Go to previous "anchor" frame in which AOIs were manually moved or computed from head motion information (positions in intermediate frames are estimated from these anchor frames).
	Go to next frame in which one or more AOIs are anchored. If any AOIs are selected (during manual editing) then it will go to the next anchor for a selected AOI.
	Go to the last frame in which any AOI is anchored. Note, this option is helpful if, after starting to define manual AOIs, the tab is closed and reopened at a later time. This will take the program to the last frame that was edited ("where you left off").
	Save (backup) AOI data (AOI data will automatically be saved when tab is closed, this button is purely for backup purposes).
	Display help information regarding manually manipulating AOIs.

16.7.2 Draw Moving Areas of Interest on a Video frame

Open the "Configure MAOIs" dialog as described in the previous section and advance to the first frame showing image objects to be tracked with moving areas of interest (MAOIs).

It is often wise to make areas a bit larger than the object they are designating to allow for some measurement error; however it is also usually best not to let areas overlap.

To define a rectangular AOI on the currently displayed video frame, select and then left-click anywhere on the frame. Holding down the left-mouse button, drag a rectangle of any desired size and release the button when done.

A pop-up "AOI Properties" window will appear displaying the properties of the AOI just created. A default name will be provided using the number of the new AOI.

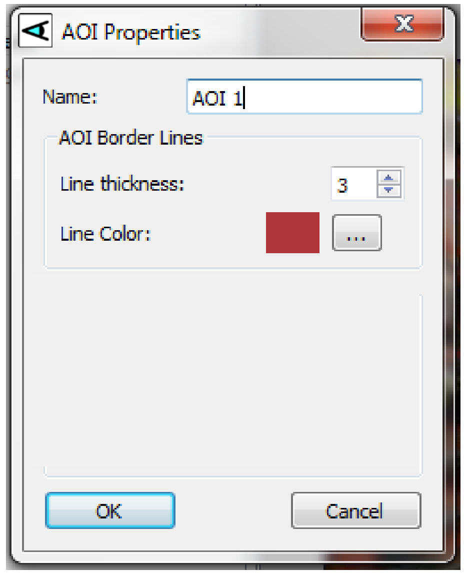

This window allows the AOI name, color, and border size properties to be specified.

Note that the border of the AOI will always be drawn using the color assigned to the AOI, but the AOI will be filled by its color only when its "anchored" and will be filled with gray when it is not anchored; this makes it possible to see when the position of the AOI has been edited and therefore anchored. A good rule of thumb is if the AOI appears in its correct location, anchor it (click it or click ⚓ to anchor all AOIs at once). (Remember that on any frame that is not an "anchor" for that AOI, its position is "calculated" by interpolating between the previous and next anchors).

Click OK to close the "AOI Properties" dialog. Properties can be edited at any time by right clicking on the AOI and choosing "Properties" from the context menu or selecting the AOI from the list in the "Areas of Interest" tab and choosing "Edit".

In order to define a polygon AOI, select (left click on the symbol to make it appear "depressed") and left click in the video image to start a polygonal AOI at that location. An example is shown in the following sequence of images. Upon left clicking within the video image, the AOI Properties dialog will appear. Enter the AOI name and properties as described previously. After selecting "OK", a triangle will appear in the image at the spot originally clicked.

Now, click anywhere on one of the lines of the triangle to add another vertex. To enlarge the polygon or move an existing vertex, left-click on the white square and drag the vertex to the desired location. Continue this process until the entire AOI has been outlined. If at any point a vertex is incorrectly added to the AOI, right-click the vertex and choose “Delete” from the context menu.

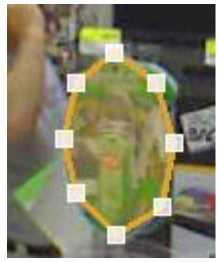

If the mouse has a scroll wheel, the AOI can be enlarged (for example, to allow some room for error around the edges of the image object) or made smaller by hovering the mouse inside the AOI and scrolling the mouse wheel. Left clicking the symbol to make it appear “un-depressed” will still allow modification of the current AOI; but a click on the image outside the current AOI will no longer start a new AOI. This feature may help to avoid unintentionally starting a new AOI before completing the current polygon. When the polygon is complete, press <Enter>. After pressing <Enter> it will no longer be possible to add additional vertices to the polygon. However, it will still be possible to move vertices, or to delete the entire AOI and start again. Attempting to Navigate away from the current frame without first pressing <Enter> will pop up a dialog asking if the AOI is complete. If the symbol was “undepressed, as described above, be sure to left click to make it appear “depressed” again before clicking on the image to start a new polygon AOI.

Vertex positions and overall AOI size can be edited in any frame at any time within the Configure Moving AOIs tab.

To permanently delete an AOI, left click the AOI to select it, and choose “Delete selected AOI” under Menu (at the upper left corner of the “Moving AOIs” window. Be sure that AOI to be deleted is the AOI name being displayed in the “Warning” pop-up window. To delete all existing AOIs, select “Delete all AOIs” from the Menu.

16.7.3 “Clone” Moving Areas of Interest if needed

There may be occasions where two or more identical objects appear in the scene video simultaneously and it may be desirable to treat gaze on any of these copies as the same for statistical purposes. For example there may be multiple sheep, in the scene image. The user may want to consider time during which gaze is on a “sheep” without regard to which one is being viewed.

If an Area of Interest is drawn around one such object it can be “Cloned” to create other areas with the same name designating other copies of the same generic object. To “Clone” an AOI, right click on the area and select “Clone this *name*”, where *name* is the name of the AOI. An identical boundary, with the same number of vertices, the same color, and the same name will appear slightly offset from the original.

Drag the new “cloned” boundary and adjust its size or shape as needed to form a boundary about one of the other images of the generic object (for example, one of the other sheep in the scene image).

Using the "sheep" example, *ETAnalysis* statistics will recognize only one object called "sheep". Time viewing "sheep" will include time spent on the original "sheep" and all of its clones, etc.

The cloned area can be deleted by right clicking it and selecting [Delete "*name* (Clone1)"], where name is the AOI name. If only the original ("parent") AOI is deleted, by right clicking on it and selecting [Delete "*name*"], then the clone will take the place of the parent AOI.

16.7.4 Manually Adjust MAOIs Throughout Video

As the video progresses, Areas of Interest drawn on a single video frame do not automatically remain fixed to their objects. In some cases an "Auto Detect" or "Auto Associate" feature can be used as described in sections 16.7.5 and 16.7.8. Otherwise, it is necessary to manually move or modify the areas as the video advances. Remember that area positions are interpolated between anchor points. The more anchor points used, the more accurately the area will follow the intended object, but the amount of manual work is also increased.

Sometimes a good strategy is as follows. Advance to the frame at which the designated object (object being tracked by the AOI) first begins to move. With the AOI selected, click the anchor button to make this an anchor frame for the AOI (i.e., tell the program the AOI is in the correct position in this frame, before it starts to change direction or move). Advance the video to the next frame at which the object either stops moving or obviously changes direction or rate of motion. Drag the AOI to the proper position on this frame. Click  to return the frame where motion started, and play the video (or drag the slider) to see if the AOI follows the object closely enough. If not, stop about half way between the two anchors and make an adjustment. Repeat the process to add anchor points at as many intermediate frames as necessary.

To move an individual AOI, simply place the mouse cursor within it, so that the mouse cursor changes to the 4-way arrow symbol, and hold down the left mouse button while dragging it to the desired position. To stretch, shrink or resize a rectangle AOI, place the mouse cursor over one of the handles (located at the corners and at the center of each side), so that the mouse cursor changes to the 2 way arrow symbol, and hold down the left button to drag the handle.

Polygon shapes can be adjusted by using the left mouse button to drag individual vertices or hovering inside the AOI and using the mouse scroll wheel to scale it up or down. Click  in the row of Manual AOI buttons in the Areas of Interest tab at any point to view the various commands for resizing/editing AOIs.

It is also possible to move or stretch several AOIs simultaneously by constructing a "Multiselect" rectangle. To define a Multiselect rectangle, hold down the CTRL key, and use the left mouse button to drag a rectangle over the set of areas to be included. Release the mouse button. The Multiselect box will be visible as a light gray rectangular area. To adjust the Multiselect region with respect to the areas, drag the entire area, or one of its handles with the left mouse button (<u>no</u> CTRL key). To drag all of the included areas along with the multiselect box, hold down <CTRL> and drag with the left mouse button. To resize the entire multiselect group, hold down <CTRL> and drag one of the multiselect box handles. Navigating to a different frame will automatically make the Multiselect box go away.

Tip: a Multiselect rectangle can also be used to resize a single polygon AOI or multiple polygon AOIs without having to move each polygon vertex separately. This is extremely helpful for adjusting all vertices at once using a rectangle that surrounds the polygons, just as with a rectangular AOI.

Whenever an AOI is moved it becomes anchored in the current frame. Notice that after moving an AOI for the first time in a given frame, the area color of the AOI changes from gray to the color of the AOI. An AOI is also anchored when resized, first created, hidden or unhidden (more on hiding later), or when it is simply selected by clicking on it.

An individual AOI can be un-anchored by right clicking on it, and selecting “Remove Anchor”. Anchors are specific to individual AOIs; so some AOIs may be anchored on a particular frame while others are not.

All the AOIs on a frame can be simultaneously anchored by clicking the “Anchor all AOIs” button. Alternatively all the anchors can be removed from all AOIs on the current frame by clicking the “Remove anchor” button. Note that all AOIs must have at least one anchor, and it will not be possible to remove the only anchor for an AOI.

If all AOIs in a current frame are in their correct positions, it may be a good idea to click on the “Anchor all AOIs” icon to ensure that they will remain in that position and will not be recalculated based on later changes that might be made to their positions in surrounding frames.

Use the next anchor button to advance to the next frame that is an anchor point for any AOI. The previous anchor button will back up to the closest previous frame that is an anchor for at least one AOI. Note that the video will advance (or backup) to the next frame with any AOI anchor if no AOI is selected. If an AOI is selected, it will advance to the previous or next anchor point for that particular AOI.

It is also possible to hide and unhide AOIs in different frames. Frames on which an AOI changes visibility (from visible to hidden or visa versa) must be anchor points. If an AOI is "hidden" on a particular frame, it will become invisible from that frame to its next anchor point, or to the end of the video file if there are no subsequent anchor points. This feature can be useful when objects in the video move in and out of view. ***It is highly recommended to "Hide" an AOI (via the method described here) when it is not visible as opposed to just moving the AOI outside the image bounds. If an AOI is simply moved outside the image bounds, it will be anchored there and may be hard to retrieve later.***

To hide an AOI, right click on it and select "Hide" from the pop up menu. An anchor point will be created for that AOI on the previous frame (last visible frame), and it will become invisible from the current frame to its next anchor point. To make it invisible past the next anchor point, advance to the anchor point, right click on the AOI, and select Hide, or un-anchor the AOI to make it invisible up to the next anchor.

To unhide an AOI, right click anywhere on the video not within a displayed AOI, hover the mouse over "Unhide" to see a drop down list of hidden AOIs, and select the one to be unhidden. This frame will become an anchor point for that AOI, and it will become visible on subsequent frames, at least up to the next anchor. If the AOI was not visible on previous frames it will remain hidden on those.

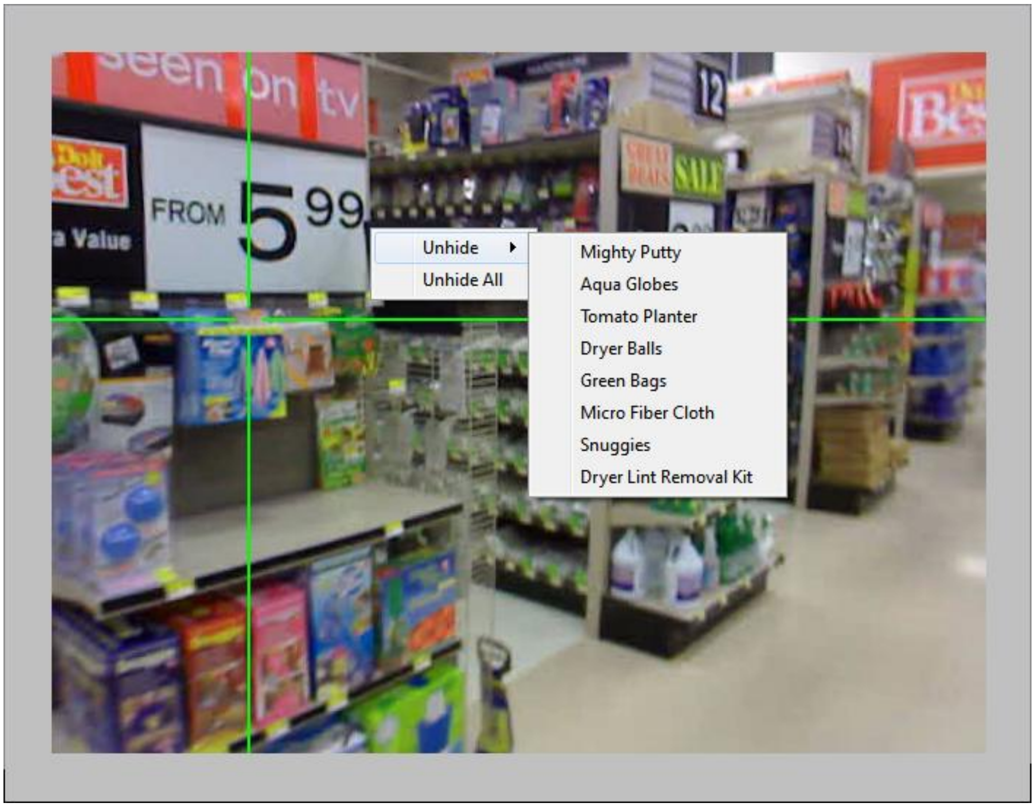

When cloned items are present (see section 16.7.3) it is possible to hide just the original, just one of the clones, or some subset of these. If cloned objects are hidden, the "Unhide" list will differentiate between clones by appending the name of hidden objects with "(Clone *n*)" where n is clone number. "(Clone 1)" will indicate the first clone of the object "(Clone 2)" the second clone, etc.

Be sure to "Save & Close" (button at upper right of "Configure MAOIs" dialog) when finished defining MAOIs.

16.7.5 "Auto Detect" MAOIs

If an MAOI has been defined around the boundaries of an image feature with well defined edges, and with the same image content throughout the video, the program can attempt to "Auto Detect" the MAOI. In this case, the program will attempt to keep the MAOI boundaries attached to image feature boundary as the video advances from beginning to end and the image feature position and perspective changes.

Note that if LAOI or AI Object bounding boxes were recorded on the scene image by ETVision these are likely to interfere with MAOI "Auto Detect" in ETAnalysis. If "Auto Detect MAOIs" will be used in ETAnalysis, it is recommended that ETVision be set to not "Record LAOI / AI" (See section 9.4.2.2 of ETVision manual).

First draw an MAOI (or multiple MAOIs) on a video frame as described in section 16.7.2.

There are two available auto-detection functions that differ slightly. One, called "Auto Detect" attempts to use the entire contents of the enclosed area as the "template" to recognize the same bounded area on other video frames. In most cases this is the technique that will produce the best results. The alternate function, called "Auto Outline Detect" uses just the area close to both sides of the area boundary line as the template. This alternate technique may work better when the contents of the bounded area is constantly changing (E.g., a cell phone or tablet with a changing screen image, etc.), but has a physical boundary such as a monitor bezel, window frame, etc.

To set an MAOI to "Auto Detect", first create an MAOI on any video frame, as described in section 16.7.2. After creating the MAOI, right click within the AOI and select either "Auto Detect *name*" or "Auto Outline Detect *name*" from the context menu, where *name* is the MAOI name. An "Auto Detect MAOI" dialog will appear as shown below.

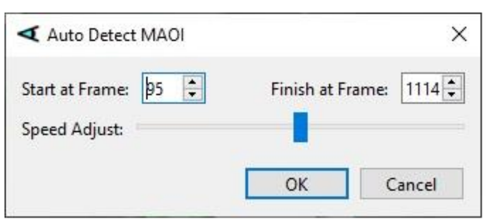

"Start at Frame" will default to the current frame and "Finish at Frame will default to the last frame of the video. Adjust these items if desired either by typing in values or using the up/down arrows. The "Speed Adjust" slider defaults to the center position. The slider can be moved to the right to increase auto detect speed by using fewer detection features at the possible expense of less detection accuracy.

The slider can be moved to the left to use more detection features resulting in more accurate, but slower auto detection. Click “OK” to start the process.

The program will use the entire image within the boundary (“Auto Detect”) or just the image area near the boundary outline (“Auto Outline Detect”) on the current frame as a “template” to try to find the same area on frames between the specified “Start” and “Finish” frames.

A “Detecting MAOI…” window will pop up to show the MAOI being detected as the video advances from the “Start” to the “Finish” frame.

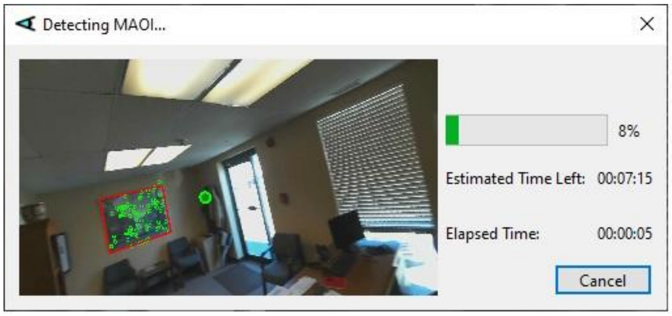

The pop up window also shows a progress bar and the estimated time remaining to advance through the specified frames. When the process reaches the “Finish” frame a pop up “Finished” dialog appears, with an “OK” button. Click OK to complete the process.

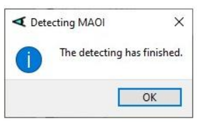

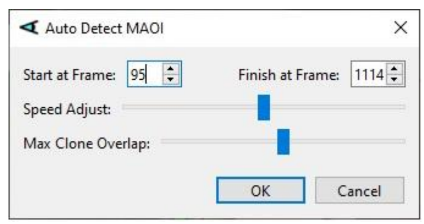

If the selected MAOI has a clone (section 16.7.3) the “Auto Detect MAOI” dialog will also have a “Max Clone Overlap” slider, set to 50% by default. If a clone is detected to overlap the parent by more than this percent, it is assumed that the system is probably mistakenly detecting the same image object as both the clone and parent and only the parent will be considered as “detected”.

The program does not attempt to match the area image content template on every video frame. If the video viewer speed is set to 1.0, it attempts to find the area of interest on every 4^{th} frame, setting these frames as “anchor frames” and interpolating area of interest position between the anchor frames. The matching interval is always 4 times the viewer speed setting; so if the viewer speed is set to 2.0, every 8^{th} frame will be an anchor, and if the speed is set to 0.5 every 2^{nd} frame will be an anchor, etc. Note that the more frames used as anchors, the longer the auto-detect process will take.

Once the auto detect process is finished the MAOI vertex point positions are defined to the system in exactly the same way as manually defined MAOIs, with an anchor frame every *n* frames, where *n* is 4 times the viewer speed setting.

To set all MAOIs to “Auto Detect” select “Auto Detect All MAOIs”, or “Auto Outline Detect All MAOIs”. In this case the program will use the frame on which each area was first defined as the "template" or "definition frame" for that image object. Note that this might be a different frame for each area of interest. As with Auto Detect for a single area, a "Detecting MAOI…" window will pop up to show the MAOIs being detected as the video advances from the beginning from the beginning to the end of the video. As with a single MAOI, the anchor frame interval will be 4 times the viewer speed setting.

Note that Auto Detect will be most successful when the image object or feature is well defined (good contrast between image features within the MAOI) and has an appearance that changes only in position and perspective, and remains in full view. Image features that appear two-dimensional will usually work best. For example, a picture hanging on a wall or a computer monitor that is showing a static image will usually work well. A computer monitor showing changing content, or a window with a changing outside view may not work well for "Auto Detect". Such objects may track more reliably with "Auto Outline Detect" assuming there is enough static content adjacent to the area outline (monitor bezel, window frame, etc.).

A 3 dimensional object that is always viewed from a similar angle in the video will often work well, but if the viewing angle changes such that a different face of the object is shown, that will probably not be well recognized by "Auto Detect".

During periods for which the feature is not in view (or is not recognized by Auto Detect) the AOI will be treated as "Hidden". At other times Auto Detect may simply place MAOI vertex points improperly. If Auto Detect fails to recognize the feature, or improperly recognized the feature for a short time during the video, it can be manually corrected using the manual AOI adjustment techniques described the previous section.

If a feature is improperly detected for a large section of the video, it may be most efficient to delete the MAOI and replace it with a manually adjusted MAOI as described the previous sections. If a feature is improperly detected for only a short period during which it is stationary or moves smoothly, it can often be corrected by simply deleting the anchor points for that MAOI during the period in question. The MAOI position will then be determined by linear interpolation between the remaining anchor points. This technique is often useful when a feature is partially occluded for a brief period.

TIP: When making manual corrections, shortcut keys can be often be used to speed the process. When an MAOI is selected, the "D", "F", and "B" keys are shortcuts for "un-anchor current frame", "advance to next anchor frame", and "go back to previous anchor frame", respectively. Note that it is possible to keep 3 fingers of one hand on the "D", "F", and "B" keys, and the other hand on the mouse while gaze remains on the screen.

Be sure to "Save & Close" (button at upper right of "Configure MAOIs" dialog) when finished defining MAOIs.

16.7.6 Advanced Strategy for Auto MAOI Tracking

The Auto Tracking feature may sometimes struggle when the appearance of the tracked object changes. This may be due to change in object orientation with respect to the camera, change in the background, or change in surface content if the object is a monitor or a window, etc.

Following are two strategies for handling this. If the change in object appearance is subtle proceed as follows:

> Start Auto Detect (or Auto Outline Detect), and at the point where the object appearance changes and Auto tracking begins to struggle click "Cancel". Adjust the AOI to properly indicate the object. Right click in the AOI and select Auto Detect (or Auto Outline Detect). Auto detect will proceed from this point using both the new object appearance as well as its original appearance as recognition templates. Repeat as necessary.

If the object appearance changes radically, a slightly different strategy may be best. For example, a 3 dimensional object may have a completely different shape when viewed from the side than when viewed from the front. When the orientation changes from front to side view it may not be an advantage to use both the original and new view as feature templates, since the original front view template may interfere rather than help. In this case it may be better to use following strategy:

> Start Auto Detect (or Auto Outline Detect), and at the point where the object appearance changes and Auto tracking begins to struggle click "Cancel". Create an AOI clone (see section 16.7.3), hide the original AOI, and adjust the clone to properly indicate the object. Right click on the clone AOI and select Auto Detect (or Auto Outline Detect). Auto detect will proceed from this point using only the new object appearance as a recognition template for detecting the cloned AOI.

16.7.7 Use Environment video to Auto Detect MAOIs for multiple events or for other environment videos

If MAOIs are defined on an environment video, the MAOI definition frames from the environment video can be used to find the same image objects, and to Auto Detect MAOIs, in multiple event videos, or in other environment videos.

A video containing suitable images of all the image objects to be tracked with MAOIs must be designated as an environment video. This may be a video from one of the project events, or a special video made by recording a scene video as the eye tracker optics are moved about the environment. If a special video is made for this purpose, care can be taken to include an optimal image of each visual object of interest. A subject video, including gaze measurement, can be used if desired, but it is not necessary to actually measure eye movements for the environment video since only the scene camera video will be used. Note also that all objects of interest do not need to be visible at the same time, so long as each is clearly visible somewhere in the video.

Make this scene video an "environment video" as described in section 16.3.2. Right click the video file node, under the "Environments" node, and select "Configure Moving Areas of Interest". A "Configure MAOIs…" tab will appear in the Display Area showing the first frame of the video file corresponding to the selected node.

For each object of interest, choose a video frame on which to draw a "Moving Area of Interest" around the object image. This can be a different video frame for each object. There is no need to actually move any of the areas to track image objects on the environment video; the "Moving Area of Interest" for each just needs to be properly drawn on one frame. See instructions for drawing an individual MAOI in section 16.7.2.

After drawing an MAOI for each visual object of interest, “Save & Close" the MAOI file (button at upper right of “Configure MAOIs…” dialog). A “Moving Areas of Interest" node should now appear under the environment video node. The next section describes how to use this node to Auto Detect the same set of MAOIs on multiple event scene videos. The subsequent section describes using the node to Auto Detect the same set of MAOIs on multiple environment videos.

16.7.7.1 Auto Detect MAOIs for multiple events

If multiple subjects move about in the same environment while wearing a head mounted *ETvision* system, then even though each “event” will have a unique scene video, a “batch” processing feature can be used to Auto Detect MAOIs in all the separate event videos, using information from the environment video “Moving Areas of Interest" node.

Right click the “Moving Areas of Interest” node and Select “Auto Detect AOIs for Multiple Participants” from the resulting context menu. ***Note****: this multiple event feature is not available for StimTrac projects (section 17), but in this case the multiple environment feature can be used, as described in the next section (16.7.7.2).*

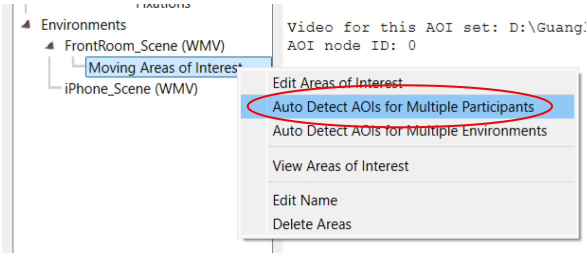

A dialog will appear showing a list of events.

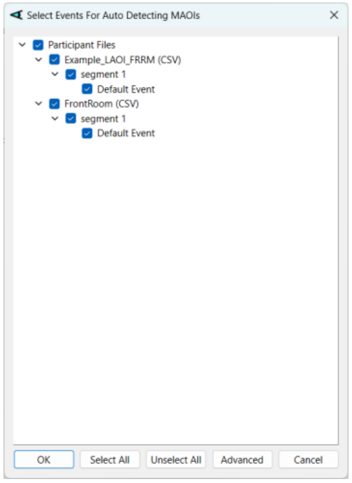

Select the events to be processed, and click OK. An "Auto Detect MAOI" dialog will appear as shown below.

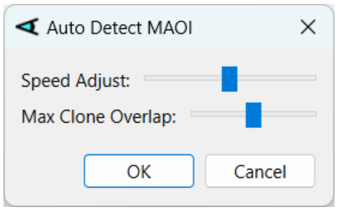

The "Speed Adjust" slider defaults to the center position. The slider can be moved to the right to increase auto detect speed by using fewer detection features at the possible expense of less detection accuracy. The slider can be moved to the left to use more detection features resulting in more accurate, but slower auto detection.

The "Max Clone Overlap" slider applies only to MAOIs that have clones (described in section 16.7.3). The slider defaults to 50%. If a clone is detected to overlap the parent by more than the amount specified by the slider, it is assumed that the system is probably mistakenly detecting the same image object as both the clone and parent and only the parent will be considered as "detected".

There is no start and finish frame choice. Each event video will be processed from start to finish. Adjust the parameters if desired and click "OK". The program will cycle through each designated event and attempt to identify the same set of MAOIs defined on the environment video.

It will be important to review each event scene video after Auto Detect is complete. If the image object is partially obscured or not clearly discernable in the image during some video frames, auto detect may sometimes place MAOI vertex points improperly and these sections will need to be corrected manually as discussed in section 16.7.4.

16.7.7.2 Auto Detect MAOIs for multiple Environment Videos

If multiple environment videos have been created depicting the same environment with the same objects of interest, then MAOIs identified on one of them, as previously described, can be used to Auto Detect the same set of MAOIs on all (or some subset) of the environment videos.

Under the environment node for which MAOIs were designated, right click the "Moving Areas of Interest" node, and Select "Auto Detect AOIs for Multiple Environments" from the resulting context menu. A dialog will appear showing a list of Environment nodes. Select the ones to be processed, and click OK.

An "Auto Detect MAOI" dialog will appear with the same "Speed Adjust" and "Max Clone Overlap" sliders described in the previous section. Adjust the sliders if desired and click "OK". The program will cycle through each designated environment video and attempt to identify the same set of MAOIs.

It will be important to review each video after Auto Detect is complete. If the image object is partially obscured or not clearly discernable in the image during some video frames, Auto Detect may sometimes place MAOI vertex points improperly and these sections will need to be corrected manually as discussed in section 16.7.4.

16.7.8 "Auto Associate" MAOIs

An MAOI can be "Associated" with (or "slaved to") the motion of another "parent" MAOI. The MAOI will change perspective and move in tandem with the associated MAOI. This is most often done with an image feature that is coplanar and adjacent to or within an MAOI that is being "Auto Detected", as described in the previous section. For example if the entire outline of a framed wall picture is "Auto Detected", an object within the picture might be defined by an MAOI, using a video frame where the object is in good view, and then "Associated" with the picture MAOI. The MAOI outlining the object will now move with and change perspective along with the parent MAOI.

To "Auto Associate" an MAOI with another, right click within the MAOI; select "Auto Associate To >"; and then, from the pop up list, select the MAOI that will be the "parent". An "Associating MAOI" window will pop up to show the MAOI being "associated" (or "slaved") as the video advances. No matter what the current frame is, the "Associating" process will advance from the beginning of the video to the end.

The pop up window also shows a progress bar and the estimated time remaining to advance through the video. When the process reaches the end of the video a pop up "Finished" dialog appears, with an "OK" button. Click OK to complete the process.

The "Associate" function works best when the "child" MAOI is coplanar with the "parent" MAOI, and when the "child" MAOI is of comparable size or smaller than the "parent" MAOI.

If an associated parent and child object both have clones (see section 16.7.3), the association is automatically applied to all of the clone pairs. For example, if AOI_1 and AOI_2 both have 2 clones, and if AOI_2 is "Associated" to AOI_1, then the first AOI_2 clone will be automatically associated with the first AOI_1 clone, and the 2^{nd} AOI_2 clone will be associated with the 2^{nd} AOI_1 clone, etc.

If finished defining MAOIs be sure to click the "Save & Close" button at the upper right of the "Configure MAOIs…" dialog.

16.7.9 Use AI Object Detection to create MAOIs

An artificial intelligence model ("ArgusAIModel") can be used in *ETAnalysis* to detect image objects and create MAOIs.

Note that if LAOI or AI Object bounding boxes were recorded on the scene image by ETVision, these image overlays will interfere with AI Object detection by ETAnalysis. ***If AI object detection will be used in ETAnalysis, it is important that, when recording data, ETVision be set to <u>not</u> "Record LAOI / AI"*** *(See section 9.4.2.2 of ETVision manual).*

Use of AI Object Detection by *ETAnalysis* requires that the PC running *ETAnalysis* meet the minimum specs listed in section 1.3.

ArgusAIModel

Argus Science provides a default "ArgusAIModel" trained to recognize a list of 80 objects such as "person", "chair", "keyboard", "clock", "car", etc. (see full list in *ETVision* manual). An optional add on to the Argus Science *ETAnalysis* application enables users to create custom "ArgusAImodels" that will recognize objects specific to their requirements and environment. This is described in section 18 and in a separate "ETVision AI Training" manual (consult Argus Science for details). ArgusAIMmodel files have names of the form *nnn*Model*mmm*.onnx and are always occompanied in the same folder by a file named *nnn*Names*mmm*.txt. *nnn* and *mmm* can be any number of characters as long as they form a legal Windows file name, and are exactly the same on both files.

The standard ArgusAIModel consists of files named "ArgusAIModel.onnx" and "ArgusAINames.txt". The *ETVision* install program places these files in C:\Program Files (x86)\ArgusScience\ETVision. If running *ETAnalysis* on a computer other than the *ETVision* PC, these two files can simply be copied to a directory on the PC running *ETAnalysis*. To use a custom AI model, follow the directions in section 18 and in the separate "ETVision AI Trainin" manual to create the AI model files.

Create MAOIs for a single event or environment video using an ArgusAIModel

To create MAOIs using AI Object detection select "Auto Create MAOI via AI Auto Detection" from an event node or from an environment video node on the *ETAnalysis* project tree. Note that this selection will be inactive ("grayed out") if an appropriate GPU driver is not detected (see section 1.3).

When "Auto Create MAOI via AI Auto Detection" is selected, a browser dialog titled "Choose AI model" will appear. Use the browser to select the ".onnx" file for the AI model to be used, and click OK.

A "Moving Areas of Interest" node will appear under the event or environment video node. A small progress window will also appear labeled "Creating MAOIs for *name*", where *name* is the name of the participant file or of the Environment file node. The progress window contains a progress bar and digital values labeled "Estimated time left" and "elapsed time". The time required to complete the process depends on the specific capabilities of the computer and GPU, but will usually be slightly less than the real-time length of the video file. Another pop up window will indicate that MAOI creation is finished. Click "OK" to close the progress window.

To view the newly created MAOIs, right click the new "Moving Areas of Interest" node and select "Edit Areas of Interest" or "View Areas of Interest ….".

Created MAOIs will have the same names as the AI model Object names. Every frame will be an "anchor" frame, and interpolation between anchors will therefore not be needed (unless some "anchors" are later manually removed). All MAOI bounding boxes will be rectangles. The "More

Info" tab under the "Moving Areas of Interest" node will show the names and ID numbers for all object types in the AI model whether or not an object of that type was ever recognized in the video file. If multiple AI Objects of the same type are simultaneously detected (for expample, multiple "golf balls") these will be treated as MAOI clones (see section 16.7.3).

Once created the MAOIs are stored and behave in the same way as MAOIs created by any other method, and can be manually manipulated in the same ways.

Create MAOIs for multiple events or environment videos using the same ArgusAIModel
To use a single ArgusAIModel to create MAOIs for multiple events, either select "Auto Create MAOI via AI Auto Detection" from the "Participant Files" node, or pull down the Group menu and select "Group Auto Create MAOI via AI Auto Detection". To create MAOIs for multiple environment videos, select "Auto Create MAOI via AI Auto Detection" from the "Environments" node. In all cases, an event selection menu will appear. Check the desired events and click OK. A browser dialog titled "Choose AI model" will appear. Use the browser to select the ".onnx" file for the AI model to be used, and click OK.

A pop up window with a progress bar will appear in sequence for each selected event or environment video. When the progress bar completes for the last selected event or environment video, a dialog with an OK button will appear to indicate that the process has finished. Click "OK". A "Moving Area of Interest" node will appear under each selected event node or environment video node.

16.7.10 Auto Convert LAOIs to MAOIs

If *ETVision* data includes "Live Areas of Interest" (LAOIs), the LAOI definitions can be used to auto detect MAOIs. Instructions for creating LAOIs can be found in the *ETVision* manual. Note that for csv files a separate xml configuration file (same name as the csv file, but with an xml extension) contains the LAOI definition information. This file will be required by *ETAnalysis* and is expected to be in the same directory folder as the csv file.

Auto Convert single event
To use the LAOI definition information to auto-detect MAOIs in a single event for analysis by *ETAnalysis*, click on an event node and select "Auto Convert LAOI to MAOI". A pop up window titled "Converting *file-name segment-name event-name*" will appear, and will show progress as the program advances through the video file and attempts to identify the defined areas as MAOIs.

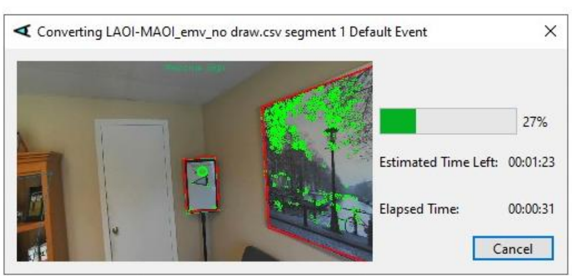

When the process is complete, the "Converting …." window will disappear, and a "Moving Areas of Interest" node will be added to the tree diagram under the event node.

Once created, the MAOIs are the same as those created by previously described methods, and can be modified in the same way as other MAOIs. Additional MAOIs can be added to the "Moving Areas of Interest" node in the ways previously described. The "Auto Convert LAOI to MAOI" process, however, can only be implemented from an event node or from the Group menu, and will always create a new "Moving Areas of Interest" node (or nodes).

When recording *ETVision* data with LAOIs defined, the user has the option to record or not record the LAOI boundary outlines (see *ETVision* manual). If the LAOI outlines have been recorded, they are a permanent part of the video image, but do not designate the MAOIs recognized by *ETAnalysis*. The MAOI outlines produced by *ETAnalysis* will appear in addition to any outlines that are already part of the video file. Note that if LAOI outlines were recorded on the video by *ETVision* they may sometimes affect the ability of *ETAnalysis* to detect the desired image object. Therefore, if it is anticipated that ETAnslysis will be used to make MAOIs it is best not to record LAOI outlines when recording data with *ETVision*.

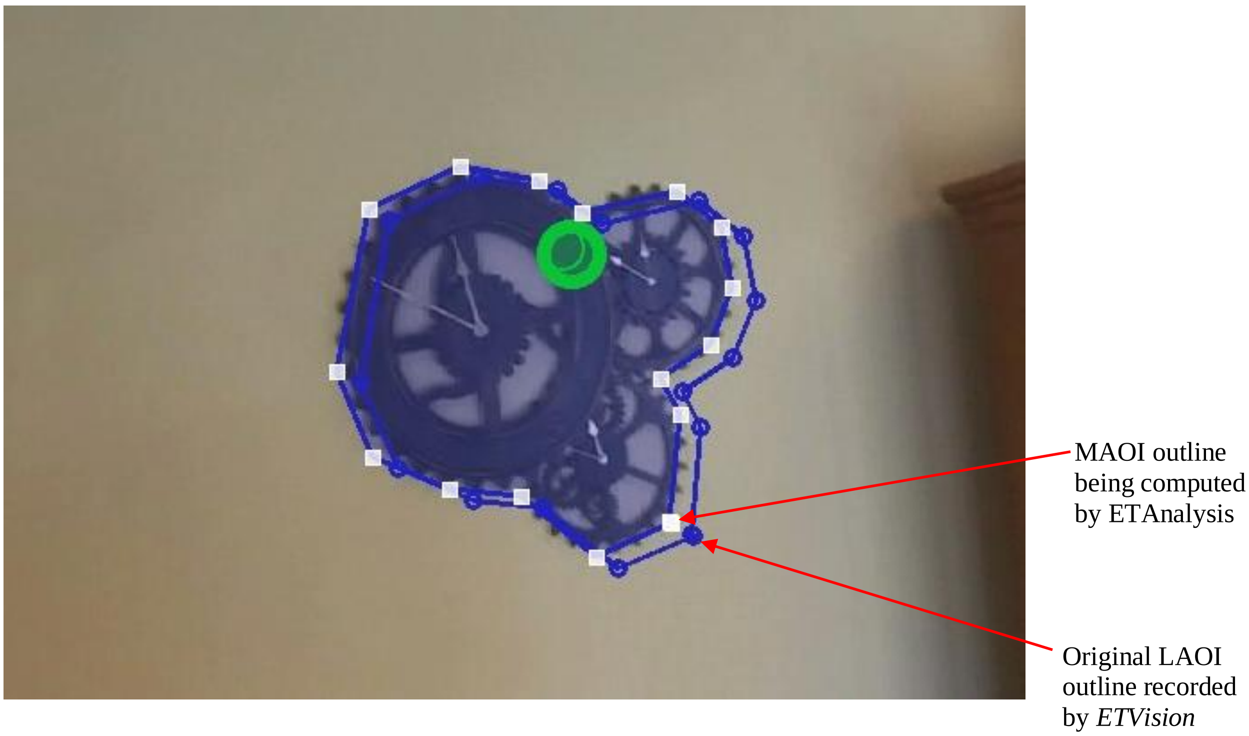

Auto Convert multiple events

To Auto Convert LAOIs for selected multiple events, either select " Auto Convert LAOI to MAOI" from the "Participant Files" node, or pull down the Group menu and select "Group Auto Convert LAOI to MAOI". In either case, an event selection menu will appear.

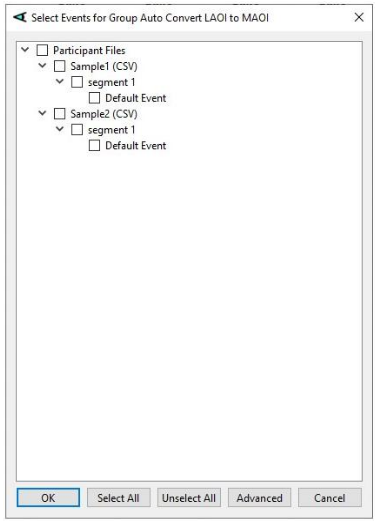

Check the desired events and click OK. A pop up window with a progress bar will appear in sequence for each selected event. When the progress bar completes for the last selected event, a dialog with an OK button will appear to indicate that the process has finished. Click "OK". A "Moving Area of Interest" node will appear under each selected event node.

Once created the MAOIs are stored and behave in the same way as MAOIs created by any other method, and can be manually manipulated in the same ways.

16.7.11 Auto Convert "AI Objects" to MAOIs

If *ETVision* data includes "AI Objects", the *ETAnalysis* can convert these to MAOIs. Instructions for detecting and recording AI Objects in *ETVision* can be found in the *ETVision* manual.

Note that for csv data files a separate xml configuration file (same name as the csv file, but with an xml extension) contains some required information about the AI Objects. This file will be required by *ETAnalysis* and is expected to be in the same directory folder as the csv file.

Auto Convert single event

Right click the event node, and select "Auto Convert AI_Objects to MAOI". A pop up window with a progress bar will appear. The process should take only several seconds. A pop up dialog with an OK button will appear to indicate that the process has finished. Click "OK". A "Moving Area of Interest" node will appear under the event node.

To do the conversion, *ETAnalysis* needs only the information recorded on the digital data file and does not need to process the scene video; so the process should take only several seconds, even for a long data file.

The AI objects recorded by *ETVision* have rectangular bounding boxes, and the MAOIs created from this data will all be rectangular. Every frame will be an "anchor frame", and interpolation between anchors will not be needed (unless some "anchors" are manually removed). MAOI names and bounding box colors will be the same as the names and colors of the AI Objects recorded by *ETVision*. If multiple AI Objects of the same type were simultaneously detected and recorded (i.e., multiple "golf balls") these will be treated as MAOI clones (see section 16.7.3). Once created the

MAOIs are stored and behave in the same way as MAOIs created by any other method, and can be manually modified in the same ways.

Auto Convert multiple events

To Auto Convert AI Objects for selected multiple events, either select "Auto Convert AI_Objects to MAOI" from the "Participant Files" node, or pull down the Group menu and select "Group Auto Convert AI_Objects to MAOI". In either case, an event selection menu will appear. Check the desired events and click OK. A pop up window with a progress bar will appear in sequence for each selected event. When the progress bar completes for the last selected event, a dialog with an OK button will appear to indicate that the process has finished. Click "OK". A "Moving Area of Interest" node will appear under each selected event node.

16.7.12 Creating Multiple Moving Area of Interest sets

Once a **Moving Areas of Interest** node has been created under an event node, using any of the methods discussed in the preceding sections, another **Moving Areas of Interest** node can be created by repeating the procedure for any MAOI creation method. The new node will be labeled **Moving Area of Interest 2.** This can be repeated for additional MAOI nodes, which will be numbered sequentially.

16.8 Fixation Sequence and Dwell Analysis with MAOIs

Fixations can be related to moving areas of interest, with associated statistics, in just the same way as they are related to stationary areas. There are two options available when computing Fixations with Moving AOIs. The user can choose to compute Fixations with respect to the scene image frame, or/and with respect to the moving areas of interests; these two options are described in the following sections.

Once computations are complete, the fixation node can be expanded to show the fixation sequence and related nodes. The Fixation Sequence and Dwell information and statistics produced are the same as that described for stationary areas of interest in Sections 11 and 12, and bar plots relating fixations to moving AOIs can be created as described in section 14.3.

16.8.1 Calculating Fixations with Respect to the Scene Image Coordinate Frame and applying them to MAOIs

When moving areas of interest are available, fixations can be determined either with respect to these moving areas; or with respect to scene image coordinate frame, as with stationary areas. This section addresses finding Fixations with respect to the scene image frame.

In the case of head mounted optics, this means fixations will be considered to be periods during which gaze was stable with respect to the head mounted scene camera image, even though areas of interest may move within this image. In the case of table mounted optics this means fixations will be considered periods during which gaze was stable with respect to the stationary scene display. In the case of *ET3Space* data, it means fixations are periods of stability with respect to a stationary surface defined to be part of the *ET3Space* environment.

Note that the position of fixations, will not depend on Area of Interest positions. However, the fixation sequence analysis will determine when these fixation positions fall within one of the moving Areas of Interest. If desired, fixations can be computed first just as described in section 10.1.4, followed by selecting “Find Fixation Sequence (Moving AOIs)” from an appropriate node. Alternately, just select “Find Fixation Sequence (Moving AOIs)”.

If “Find Fixation Sequence (Moving AOIs)” is selected, a “Configure Fixation Sequence for Moving AOIs” dialog will appear to ask which type of fixations should be computed. For fixations with respect to the scene image frame (the subject of this manual section) set the radio button to “Fixations with respect to Head”.

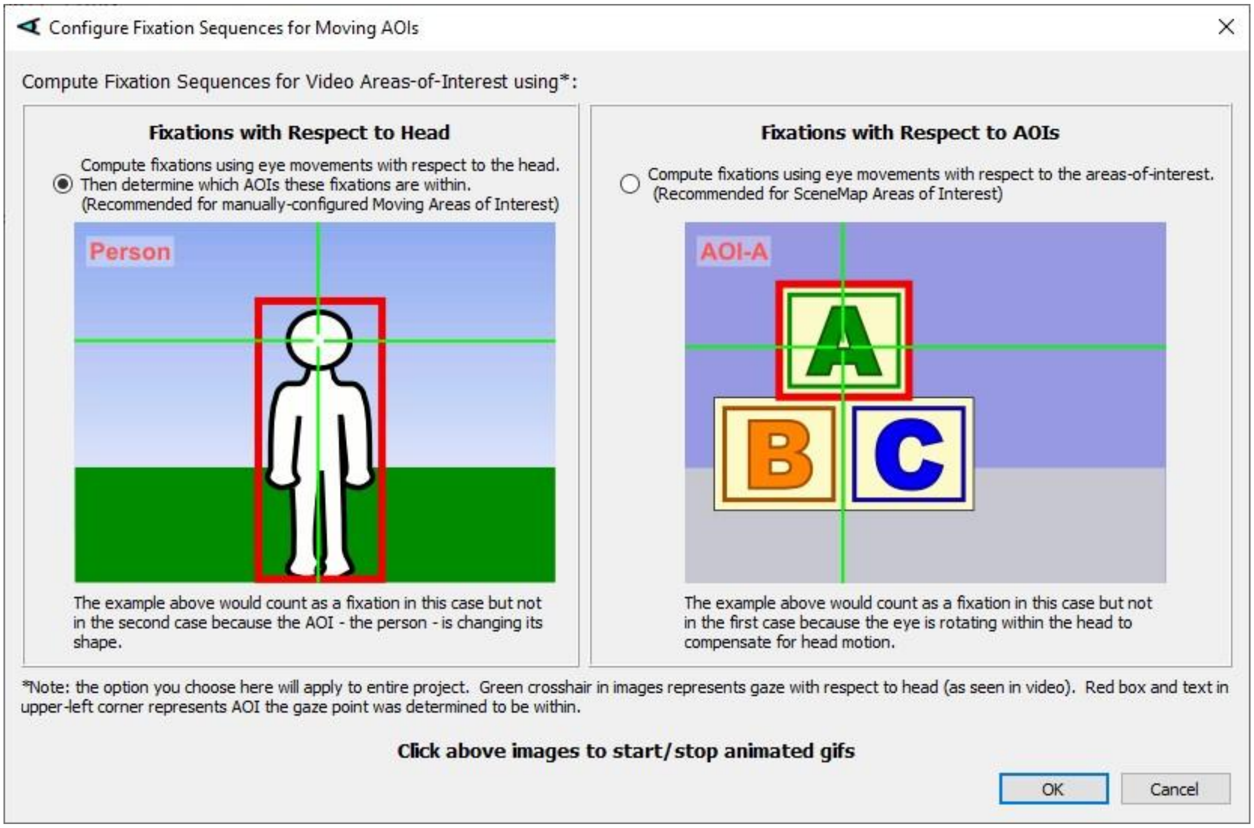

Although the label reads “Fixation with Respect to Head”, this really means with respect to the scene image coordinate frame as previously described.

This type of fixation sequence analysis can be selected from an Event, Segment, Participant (individual data file), or Participant Files (multiple data files) level. The computations will be done for all events below the selected node.

If fixations have not already been computed for events under the selected node, a “Fixation Detection Criteria” dialog will appear. The fixation algorithm and parameters have been discussed in detail in section 10. Use default parameters or modify parameters appropriately, using the Advanced Configuration button if necessary, as described in section 10. Click OK when done.

If there are multiple **Moving Areas of Interest** nodes available, then a “Select Moving Areas” dialog will appear to allow selection of the appropriate MAOI node. Fixations will now be assigned to the areas-of-interest that they fell within. Two new nodes will appear under the “Fixations” node, “Fixation Sequence” and “Dwell” nodes.

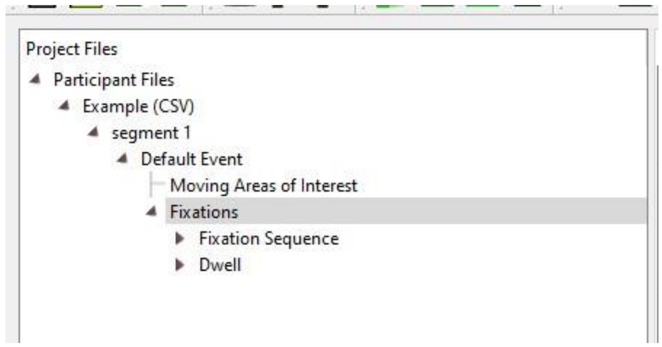

Note, if the MAOI data is edited after generating the fixation sequence and dwell nodes, it will be necessary to rerun the fixation sequence and dwell calculations to reflect the MAOI changes.

The Fixation Sequence and Dwell nodes both have an arrow sign to their left as well. Clicking on this arrow sign reveals additional data information (AOI Summary, Transition Table, Conditional Probability, and Joint Probability) with respect to fixations and dwells.

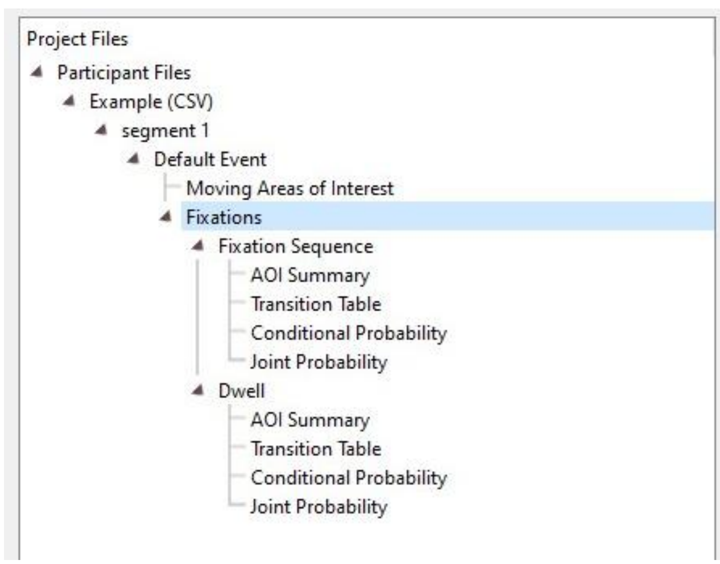

For a detailed explanation of the statistics reported see the explanation for stationary Areas of Interest in sections 11 and 12. Be aware, however, that since, in this case, areas of interest are moving with respect to fixation positions, it is possible for a given fixation to be within an area for only part of its duration.

16.8.2 Calculating Fixations with Respect to MAOIs

When moving areas of interest are available it is also possible to define fixations as periods during which gaze remains stable with respect to one of the moving areas. If, for example, ocular smooth pursuit keeps gaze in a stable position on a moving target (and assuming the moving target is bounded by a moving AOI), that will be regarded as a fixation.

In this case, note that fixations can only be defined when gaze is within a moving AOI. This type of fixation set appears as a node under a **Moving Areas of Interest** node.

Caution: ***if computing fixations in this way it is important that moving area of interest precisely follow the image target, both in terms of position, and apparent size.*** Remember that the program will look for periods of gaze stability with respect to the moving AOI boundaries, not the actual target image. Gaze position is treated as a horizontal and vertical position relative to the MAOI bounding box. For example, gaze at any particular sample, may be one third (0.33) the distance from the left edge of the bounding box to the right edge, and half (0.5) the distance from the top edge to the bottom edge.

To compute these Fixations with respect to the moving areas, right-click on a "Moving Areas of Interest" node and select "Find Fixation Sequence", or any node above that and choose "Find Fixation Sequence (Moving AOIs)". The "Configure Fixation Sequence for Moving AOIs" dialog will appear, and the radio button should be set to "Fixations with respect to AOIs".

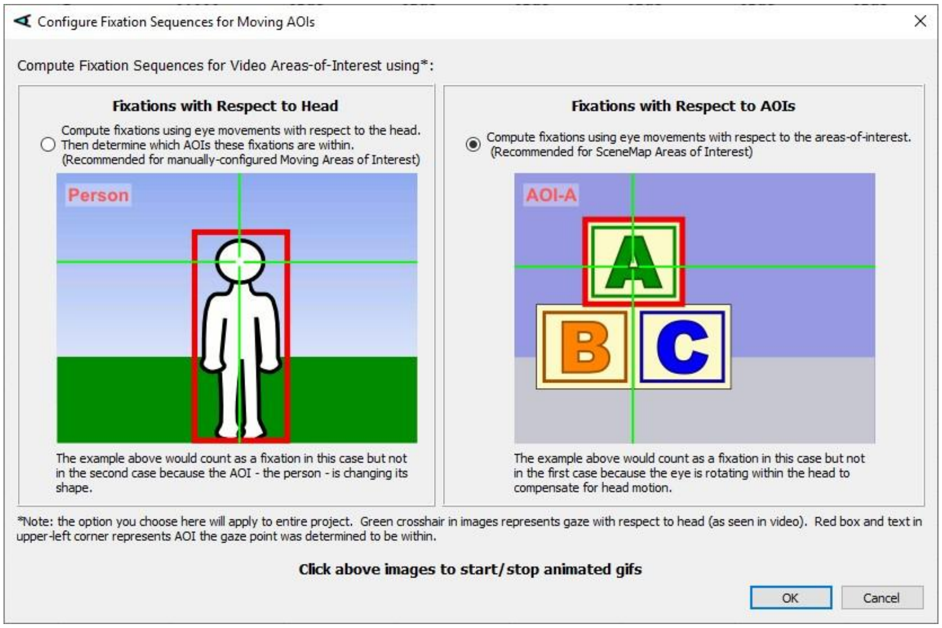

The "Fixation Criteria Detection" will appear once again, allowing the user to set their desired parameters for defining fixations (see section 10).

Two new nodes, labeled "Fixation Sequence" and "Dwell", will appear under the "Moving Areas of Interest" node. Note that there is no Fixation node, since this type of fixation is only defined by association with a Moving Area of Interest.

Expanding the Fixation Sequence and Dwell nodes will reveal additional data information (AOI Summary, Transition Table, Conditional Probability, and Joint Probability) with respect to fixations and dwells. The Fixation Sequence and Dwell information and statistics produced are the same as that described for stationary areas of interest in Sections 11 and 12.

For periods during which gaze is not within any defined MAOI, fixations are computed with respect to the scene image frame (as discussed in the previous section), and fixation positions are reported with respect to the eye tracker coordinates. Fixations within MAOIs, on the other hand, have positions that are reported as a fraction of the horizontal distance from the left to right edge of the MAOI bounding box, and a fraction of the vertical distance from top to the bottom of the bounding box.

16.9 Playing the scene video with superimposed Gaze Trail and other information

A scene video associated with an event can be viewed with a dynamic gaze trail and/or heat map display from any event node. Fixations can be dynamically displayed on the video from fixation nodes; and areas of interest, along with a time line display showing the areas visited, can be displayed with the video from fixation sequence nodes.

From an event node, right click to see the context menu and select “Play Video with Gaze”. If a video file has not already been associated with this event, a “Configure Video Data” Window will appear.

The video tab will open in the Display Area. The “Configure Display” pull down menu provides check boxes to enable “Draw Options” and “Plot Options” dialogs, and to enable a pupil diameter plot with moving time bar below the main display. The moving time bar (red line that moves from left to right) indicates the current position on the pupil diameter chart. (In the case of *ETVision* system data, left eye pupil diameter is displayed on the dynamic plot). If enabled, the “Draw Options” tab allows selection of the information to show over background image. Check “Gaze Trial”, “Heat Maps”, or both. A dialog window for adjusting the length and color of the gaze trail can be brought up by clicking the “Gaze Trail Configuration” button, and a dialog for heat map properties can be brought up by clicking the “Heat Map Configuration…” button. The “Plot Options” tab, if enabled, allows adjustments to the pupil diameter display.

The Viewer window includes the usual controls for “play”, “pause”, single step forward or back, advance to specified frame, and playback speed. A slide bar can be dragged to advance or back up through the video. Note that in some cases gaze data updates at a faster rate than the scene video. For example, *ETVision* systems update gaze data 180 times per second, while the head mounted scene camera video updates 30 times per second. “Frame” advance will always advance one scene video frame; however the “Gaze Trail” feature can show data from all of the gaze data points.

If the scene video includes audio, the audio will be heard only when the video is played at normal real-time speed (speed = 1.0); at other speeds audio will be muted.

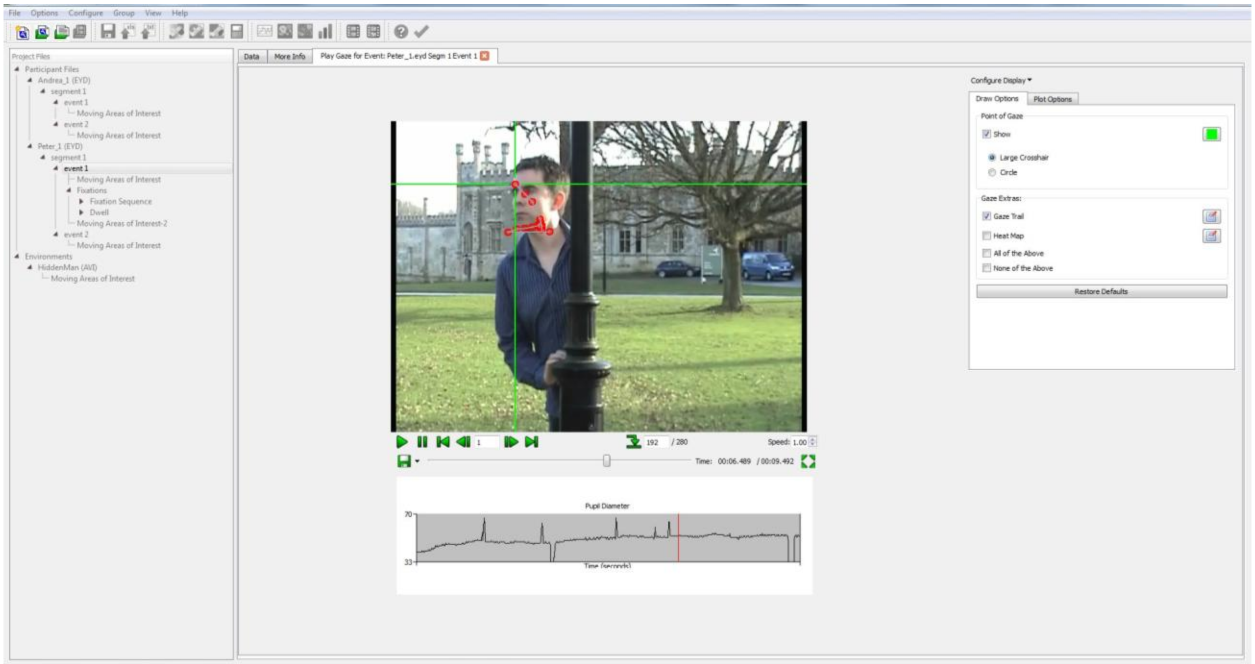

There is also a pull down menu at the lower left of the main display  with selections to record the video display as a wmv file, or capture the current frame as bit map image. A "full screen" button at the lower right  will toggle to a full screen display that omits the tree diagram, or back to the standard display.

To see the same display with MAOIs also superimposed, select "View Moving AOIs with Gaze Data" from a **Moving Areas of Interest** node.

Select "Play Video with Statistics" from a **Fixation Sequence** node to also see MAOIs, and fixation trails. Using the "Configure Display" pull down menu, "AOI Dwell Plots" can be selected in place of "Pupil Diameter". In this case the plot below the background image displays a time line plot showing when the various moving areas were visited by the subject. The display shows a horizontal bar for each AOI with the AOI color indicating periods during which gaze was in that AOI. A moving time bar (vertical red line that moves from left to right) shows current time position on the plot.

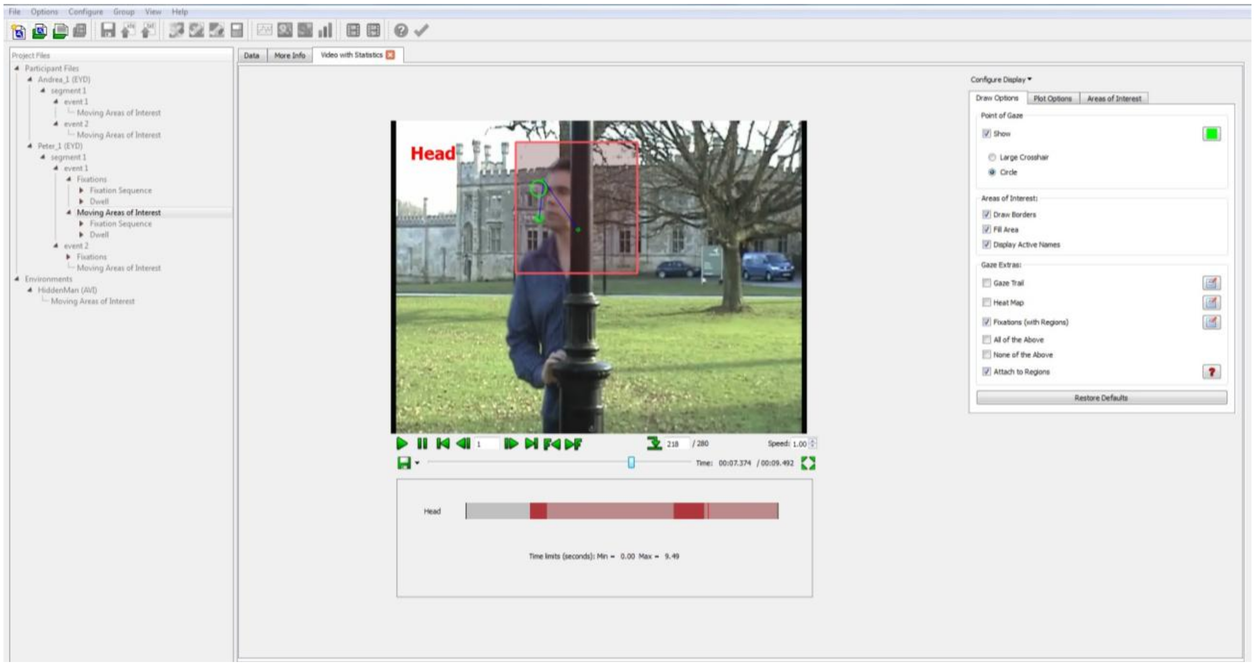

As with other video displays, Heat/Peak map display is available as a selection on the draw options tab. In the case of "Play Video with Statistics" the Configure Heat Map dialog (available by clicking the "Heat Map Configuration…" button) includes an additional set of radio buttons labeled "AOI color scaling".

Selection of "with respect to a single AOI" will mean that gaze densities for the display in each AOI will be computed using only gaze data samples that were within that AOI during the applicable time period. Selection of "with respect to all AOIs" will mean that gaze density is computed with respect to gaze data in all AOIs during the applicable time period. Other dialog items are as described in sections 14.2.1.

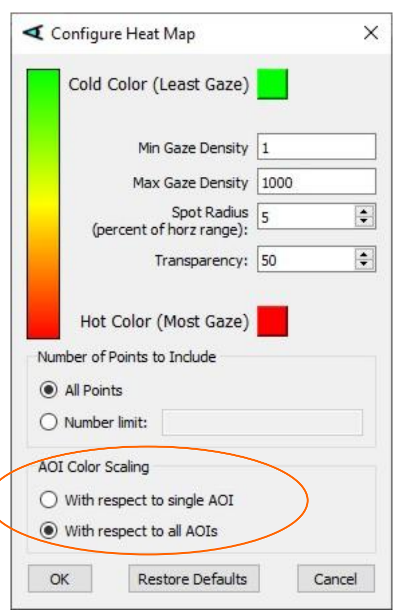

16.10 Swarm video with shared stimulus videos and MAOIs

"Swarm Video over Shared Video" and "Swarm Video over Moving AOIs" are available selections under the Group menu.

"Swarm Video over Shared Video" is applicable if the same stimulus video is associated with multiple segments or events (for example, when the same stimulus video is shown to multiple participants). More specifically, it is usually applicable if data was gathered with table mounted optics, or head mounted optics using *ET3Space* or *Stimulas Tracking*, as multiple subjects watched the same video presentation on a display screen. (See section 1.2 for description of *Stimulus Tracking* and *ET3Space*). Gaze points for multiple events (usually multiple participants) are shown as multiple dots, each a different color, moving about over the shared stimulus video. See section 16.10.1 for detailed instructions.

"Swarm Video over Moving AOIs" is usually applicable when data was gathered by a head mounted eye tracker using only a head mounted scene camera. Note that in this case the scene video is different for every subject (or every event). Objects in the environment move about on head mounted scene camera image as subjects move their heads. Even if all subjects moved about in the same environment, they all move differently, and the scene video recorded from the head mounted camera will be different each time. Moving Areas of Interest must be defined separately for each event.

In this case it is not possible to swarm data from multiple events over a video because there is not a common video. The system does however calculate where gaze is with respect to the bounding box of each moving AOI.

For example, at a given time from an event start a particular visual object in the environment, say a soft drink bottle, might be at a completely different position in the subject 1 video and subject 2 video. In each case, however, the system will know exactly where gaze was with respect to edges of the moving AOI defining the soft drink bottle outline in each video. If we have a static image showing the soft drink bottle and make a static AOI to define its outline, we can show the position of both subject 1 gaze and subject 2 gaze with respect to the static AOI on this image. This can be done for multiple Areas of Interest. Detailed instructions are in section 16.10.2.

16.10.1 Swarm Video over Shared Video

Before Selecting "Swarm Video over Shared Video" be sure to have an appropriate shared video configured for all events to be incuded. If data was gathered using *ETVision* with *ET3Space* (see section 20), use "Configure Video Data" selection from a file node or from each relevant Segment node or event node, to configure the shared video file as described in section 16.1. If using *ETVision* data with the *Stimulus Tracking* feature (section 17), use the "Configure Stimulus" selection from each relevant "Monitor" node to configure the shared video file, as described in section 17.5.

When "Swarm Video over Shared Video" is selected from the Group menu, a selection chart appears labeled "Select events that share same video". The user can select the events by checking individual event boxes, or higher-level boxes, in a tree diagram. Checking a Segment node, checks all of the events under it. Checking the Participants node selects all events in the project, etc. There are also

"Select all" and "Unselect all" buttons. The Advanced button brings up a dialog that allows selection of events based on various criteria.

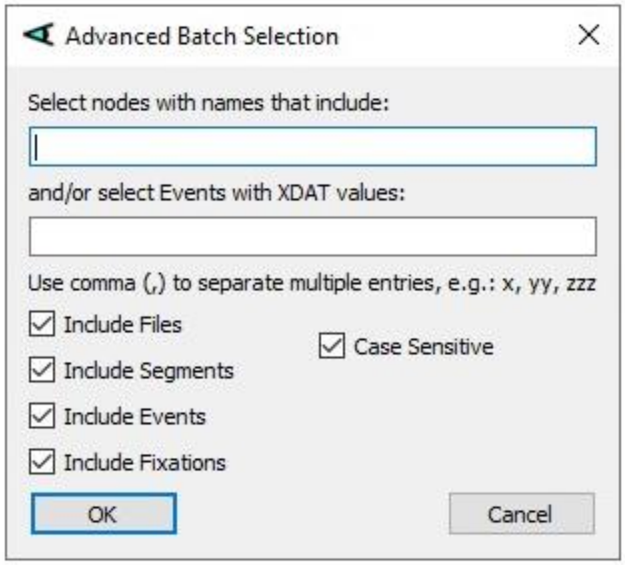

The events can be automatically selected based on the name of the node, the first XDAT value in the event, or the configured background associated with the event. So, for example, it is possible to select all events for which the initial XDAT value is 1, etc. Either check the box corresponding to each event that will be included, or use the "Advanced" button to specify display shows point of gaze for each selected event as different colored dot which moves about over the stimulus video.

When data from many events are combined, it looks like a "swarm of bees" flying over the background image, and can provide a visual illustration of whether all subjects followed a similar gaze pattern (dots stay tightly grouped) or a variety of different patterns (dots tend to spread out over display).

The "Swarm Video over Shared Video" is very similar to "Swarm Video over Background" discussed in section 15.1, but displays the gaze data with respect to video scene images rather than static background images. The viewer controls are the same as those discussed in section 16.9.

16.10.2 Swarm Video over Moving AOIs

It is assumed that, for each event to be combined in a "swarm", the same set of visual objects were defined by an MAOI set on the stimulus video, and that Fixation Sequence has been computed for each of these events. Before selecting "Swarm Video over Moving AOIs" proceed as follows to create a static image and static AOIs containing the visual objects of interest.

From the Main menu select **Configure→Background Images**; then, from the Configure Backgrounds tab, select **Background→Create single background**. Select a static image showing the visual objects of interest. Note that the "Extract Background Image from Video" button can be used to select a frame from the scene video for use as the static image. Enter desired name for "Current Background" and "Save and Close". (Note: in this case it is <u>not</u> necessary to "define attachment points").

From the Main menu select **Configure→Areas of Interest (moving in background)**. A tab labeled "Configure MAOIs in static background" will appear in the display area. Next to "Background:", at the top of the display, select the background created as described in the previous paragraph. Next to AOI set::", at the top of the display, use the pull down menu to select the MAOI set used for one of the scene videos. An "Areas of Interest:" list will appear to right of the display listing all the AOI names used in the selected MAOI set. Draw an AOI around one of the visual objects in the background. Follow the instructions in section 8.1 to draw either a rectangular or polygon area. The only difference is that the AOI properties box does not allow typing an AOI name. Rather, it has a pull down menu next to "Name:" that contains the list of MAOI file areas. One of these must be selected. The AOI must be the same type (rectangle or polygon) as the area with corresponding name in the MAOI set.

Repeat the procedure to draw all the applicable AOIs on the static background, and be sure to click "Save and Close" when finished.

Now select "Swarm Video over Moving AOIs" from the Group menu. Use the selection chart to select the desired **Fixation Sequence** Nodes. As with other Group displays, the "Advanced" button brings up a dialog that allows selection of events based on various criteria. Click OK to close the selection chart.

A video viewer will appear on a tab labeled "Swarm Video over Areas of Interest".

As with other swarm displays, the video viewer will show the gaze point from each event as a different colored dot. The viewer controls and selections are also the same as for other swarm displays. The difference is that gaze points will *only* be shown when they were within one of the MAOIs.

The viewer has the usual controls for "play", "pause", etc. At the upper right, next to "Background:", be sure the static image created in step 4 is selected. The image will show the Areas of Interest created in step 5.

A key on the right, labeled "Participants", shows the color assigned to each event being displayed. Un-checking the box next to one of the event labels will cause the data for that event to not be displayed.

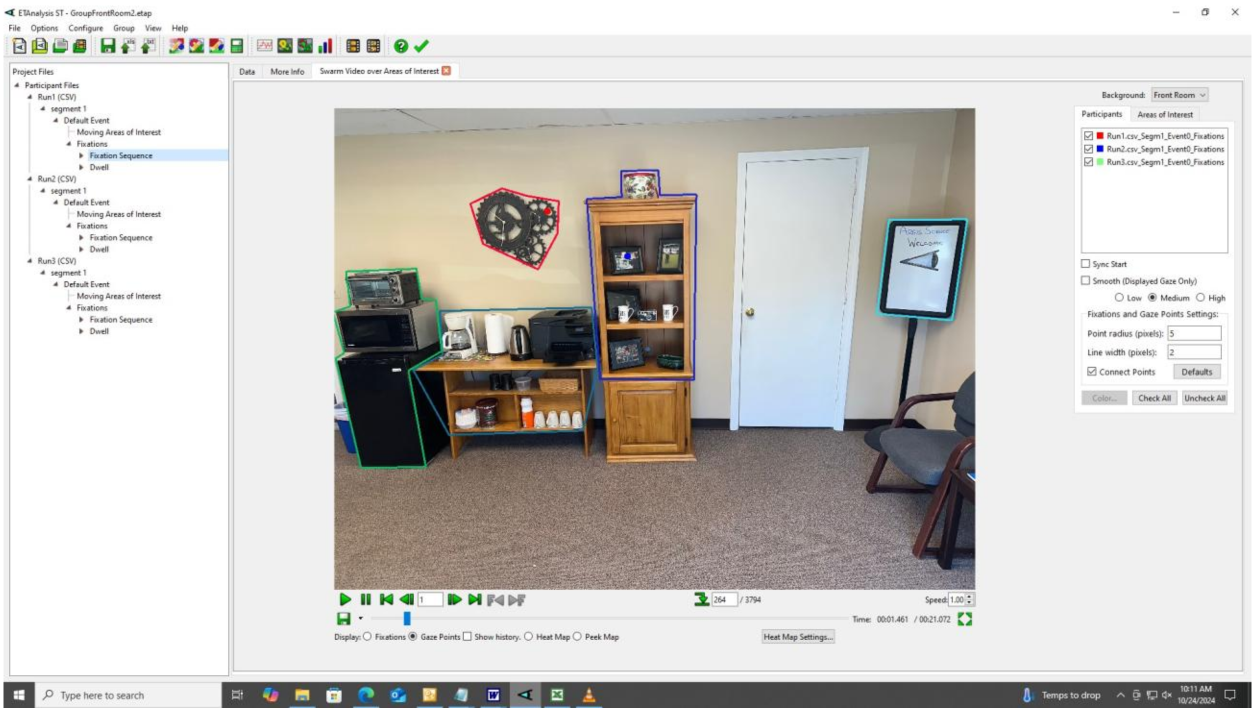

A smoothing filter can be applied to the data by checking the “Smooth” checkbox, just under the “Participants” key. Three levels of smoothing are available, as determined by the low, medium, and high radio buttons. Smoothing will apply only to the display. It will not change the data table or computed statistics.

Use the “Gaze Point Setting” controls, below the “Smooth” selection, to adjust the size of the gaze dots and the width of connecting lines (if “Show History” and “Connect Points” are both selected).

Use the radio buttons and check boxes below the video controls to select the display type. The “Show History” check box applies only to the “Gaze Point” selection. If not checked, each frame will show only the gaze positions for that data frame. If “Show History” is checked, all previous gaze points will also be shown. In other words, once a gaze point is displayed it will remain as subsequent points are added to the display. A Heat map or peek map, rather than gaze points, can be shown for each subject by selecting the corresponding radio button. Note that when Heat Map or Peek Map is selected, it becomes impossible to distinguish between the data from the different events.

The Viewer display includes the usual controls for “play”, “pause”, single step forward or back, advance to specified frame, and playback speed. A slide bar can be dragged to advance or back up through the video.

If “Fixations” are selected the “Play” button becomes inactive and “F<” and “>F” buttons become active. The viewer does not play in real time, but rather advances to start of the previous or next fixation each time the “F<” or “>F” button is clicked. For all other display type selections the “Play” button is active and the viewer plays in real time.

There is also a pull down menu at the lower left of the main display  with selections to record the video display as a wmv file, or capture the current frame as bit map image. A “full screen” button at the lower right will toggle to a full screen display that omits the project tree diagram, or back to the standard display.

If “Gaze Points”, “Heat Map” or “Peak Map” has been selected, the “record” ( ) pull down menu includes a selection labeled “Save Gaze”. This selection allows the user to record a text file with a list of all gaze coordinates with respect to the background image pixel coordinate frame.

```
0,Array-2.csv_Segm_1_Event_0,
1,Array-3.csv_Segm_1_Event_0,

sample_#,time_secs,GazeX0,GazeY0,GazeX1,GazeY1,
1,0.000,313.365,225.240,323.365,235.240,
2,0.006,313.210,225.075,323.210,235.075,
3,0.011,312.845,225.360,324.845,237.360,
```

A short sample from such a file is shown above. It begins with a list of events, In this case 2 events numbered 0 and 1. Each event is listed on a separate row with the file name, segment number, and event number. Next is a row of comma separated column labels, followed by a row of comma separated values for each data sample. Time values are from the beginning of each event. The horizontal and vertical gaze positions are labeled “GazeX0” and “GazeY0” for data from the first event, and “GazeX1” and “GazeY1” for the second event. The units for the gaze values are background image pixel values. If, for example, the background used is a 1280 x 720 pixel image, a gaze point at the center would have X,Y coordinates (640, 360).

If “Fixations” are selected, the “record” menu selection is “Save Fixations” instead of “Save Gaze”, and fixation coordinates rather than gaze coordinates are recorded to the text file. In this case the column labels are “fixation_ #,FixPosX0,FixPosY0,FixPosX1,FixPosY1, …” etc.

17 *Stimulus Tracking* Feature (Requires ST License)

If gaze data is collected from participants as they view a single display monitor, and if data is collected using a head mounted eye tracker with only a head mounted scene camera (no separate head tracker, and therefore, no *ET3Space* function), then Stimulus Tracking can significantly improve the efficiency and ease of data analysis. It can make analysis of this type of data almost as convenient and efficient as if the data had been gathered from an eye tracker with desktop mounted optics.

This advanced option requires an additional license.

The Stimulus Tracking feature tracks the position of the display monitor in the head mounted scene camera image, and translates gaze data to positions relative to the display image. The gaze data can then be displayed relative to a recording (or image) of the monitor content rather than the head mounted scene camera image. The procedure has the following steps:

- Track the image of the display monitor (edges of monitor bezel) on head mounted scene camera video.
- Parse data to define "Event" periods that correspond to video recordings or static images that were displayed on the monitor.
- Import and "Configure" the stimulus recordings and/or images to specify the boundary of the image with respect to the physical monitor boundary.
- If desired, define areas of interest within the monitor stimulus recordings or images.
- View gaze position superimposed on recorded display videos or images
- Do fixation and saccade analysis, and compute fixation sequence and dwell statistics with respect to areas of interest.

17.1 Corner Markers

Although not required, four markers may be placed on the corners of the participant display monitor (or similar rectangular display region), and may sometimes improve the ability of the software to accurately track the display. Two sets of these stickers are provided by Argus Science, with the Stimulus Tracking package, and are shown in the following image.

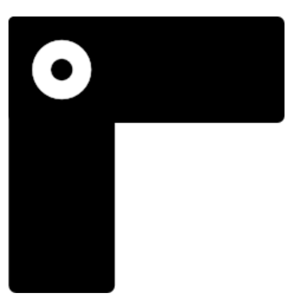

Proper placement of these stickers is shown in the following image. They should not obstruct the view of any stimulus being presented to participants.

17.2 Making a display recording

The Argus Science *ETRemote* application, running on the display computer, can be used to record the monitor display as the participant is tracked, and the *ETAnalysis* Stimulus Tracking feature can be used to display and analyze gaze with respect to that recording. The *ETRemote* manual explains the procedure for recording the display computer screen video synchronously with gaze data recorded by the *ETVision* eyetracker. Alternately user created software or another third part application can be used to record display video, but there must be a means to start and end this recording at the same time the eyetracker gaze recording begins and ends or at the same time that external data marks are placed on the eyetracker data.

If participants are interacting with the computer such that the display on the monitor is indeterminate (e.g., contains scrolling or video game responses which may be different for each participant), then it is essential that the display be recorded synchronously with eye data, as described in the previous paragraph. If participants are watching pre-generated videos or looking at static images, then these original video or image files can be used in place of screen capture video; however, in this case the eye tracker gaze data must contain marks to indicate the beginning and ending of periods during which each video or image was being displayed to the participant. The data file must be parsed into "Events" such that each event corresponds to the period during which a single image or video presentation was being displayed. See section 6 for an explanation of event "parsing". Note that a data segment can be parsed into multiple events either before or after the "Track Monitor" step described in the next section.

17.3 Track Monitor

In order to analyze gaze with respect to stimuli presented on a computer monitor, the *ETAnalysis* Stimulus Tracking feature must track the position of the computer monitor through each frame of the participant's scene video. This can be done from an **Event** node or from the **Segment** node.

If data will be parsed into multiple events, it is often most efficient to wait and parse data into multiple events after completing the "monitor tracking" step. The exception is if there are long periods of recorded data during which the subject was not looking at the stimulus monitor or which, for other reasons, will not be part of the eventual set of events. In these cases it may be more efficient to parse into events before the "Track Monitor" step so that unneeded data sections do not have to be "tracked".

17.3.1 Initiate automatic Monitor Tracking from an event node

To start the Monitor Tracking process either right click the **Event** node and select "Track Computer Monitor" from the context menu (A), or highlight the node and click the Monitor Tracking shortcut Icon (B).

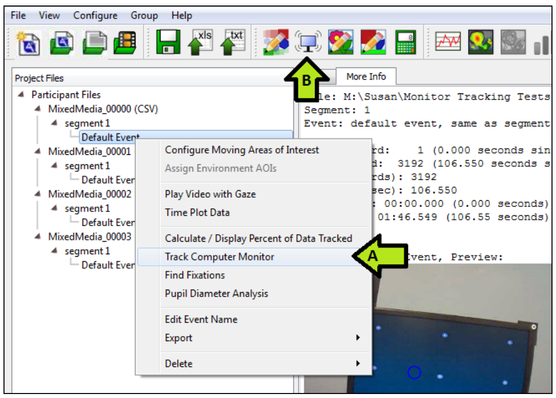

A "Tracking Monitor" tab will open in the display window, and will first be used to approve or correct the initial estimation of the monitor corner positions. If corner markers have been used, as described in section 17.1, the corner vertices should be the circle center positions at each corner of the monitor bezel. If corner markers are not present, any recognizable features near the corners of the monitor bezel can be used, but it is suggested that these be the inside corners of the monitor bezel.

The initial estimation of the monitor area-of-interest (AOI) will be outlined on the video frame as shown in the example below (A).

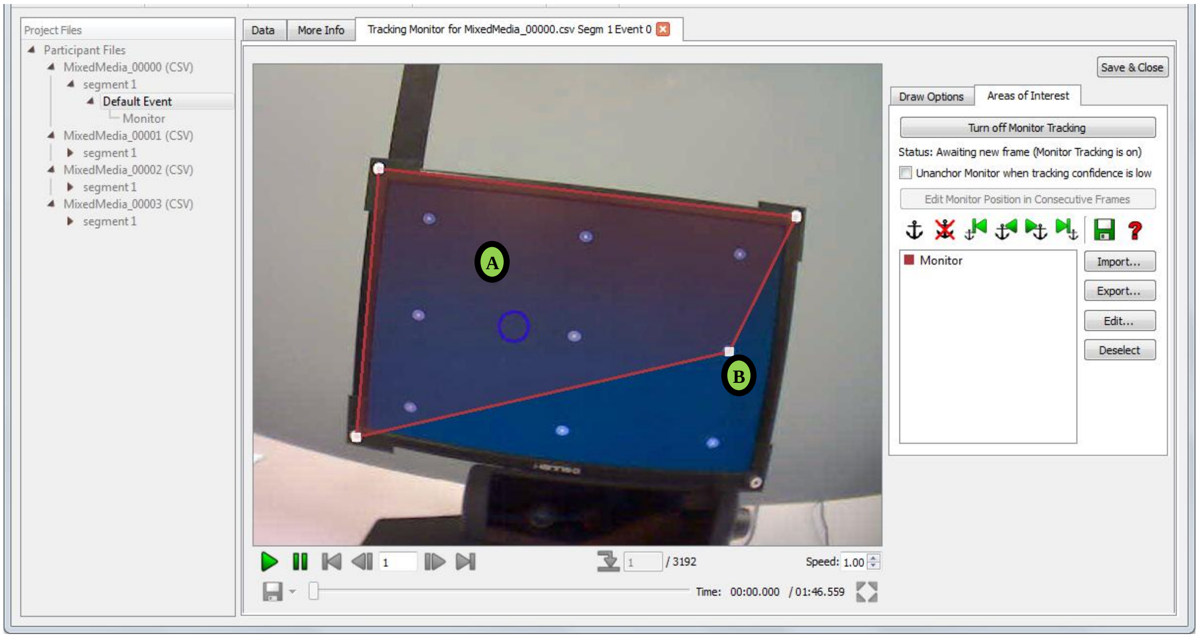

One or more corner vertices of the outline may be inaccurate due to other features in the video looking similar to the monitor targets (B). If any monitor vertex appears to be wrong, move the mouse over that vertex so that the mouse cursor changes to a 4-way arrow symbol, and use the left mouse button to drag the vertex to the correct position. When all four points appear to be in the correct positions, click "Turn on and Start Trac" (button at upper right of dialog) to start tracking. Watch as the monitor is tracked through the video (or step away if it is a long video, and return when the play through is compete). If "Speed" (at the bottom right of the scene image window) is set to 1.0 the system will attempt to recognize the monitor corner positions and correctly place the outline vertex points on every 4th video frame. These will become "anchor" frames, and the vertex positions will be interpolated between these anchor frames. If "Speed" is set to a higher value the number of frames between "anchors" will be increased proportionately, and the tracking process will complete more quickly. Conversely, if the "Speed" value is decreased, the number of frames between anchors will decrease proportionately and the process will take longer. For example, if "Speed" is set to 0.25, every frame will be used as an anchor. If the monitor outline cannot be identified on a frame that would normally be an anchor frame, the outline vertices are not changed. The program advances to what should be the next anchor frame and tries again.

If at any point a corner of the monitor AOI veers off its proper target, pause the video, adjust that corner vertex (by dragging it with the left mouse button) and hit play to restart tracking. These manual edits may also be done at any time after tracking has completed. See section 17.3.3 for more information about editing the tracking results. In most cases, only a small number of manual adjustments, if any, will be necessary. When finished, click "Save and Close" (upper right corner of dialog).

A **Monitor** node will appear under each processed event on the project tree diagram.

17.3.2 Initiate auto Monitor Tracking from a Segment node

If the only segment event is the "default event" (the whole segment), selecting "Track Monitors" from the **Segment** node context menu is exactly the same dialog and monitor tracking procedure as selecting "Track Computer Monitor" from an **Event** node (described in the previous section).

If the segment has already been parsed into multiple events, selecting "Track Monitors" from the segment node brings up the slightly different dialog shown below. Initially it will show the first frame of the first event in the segment.

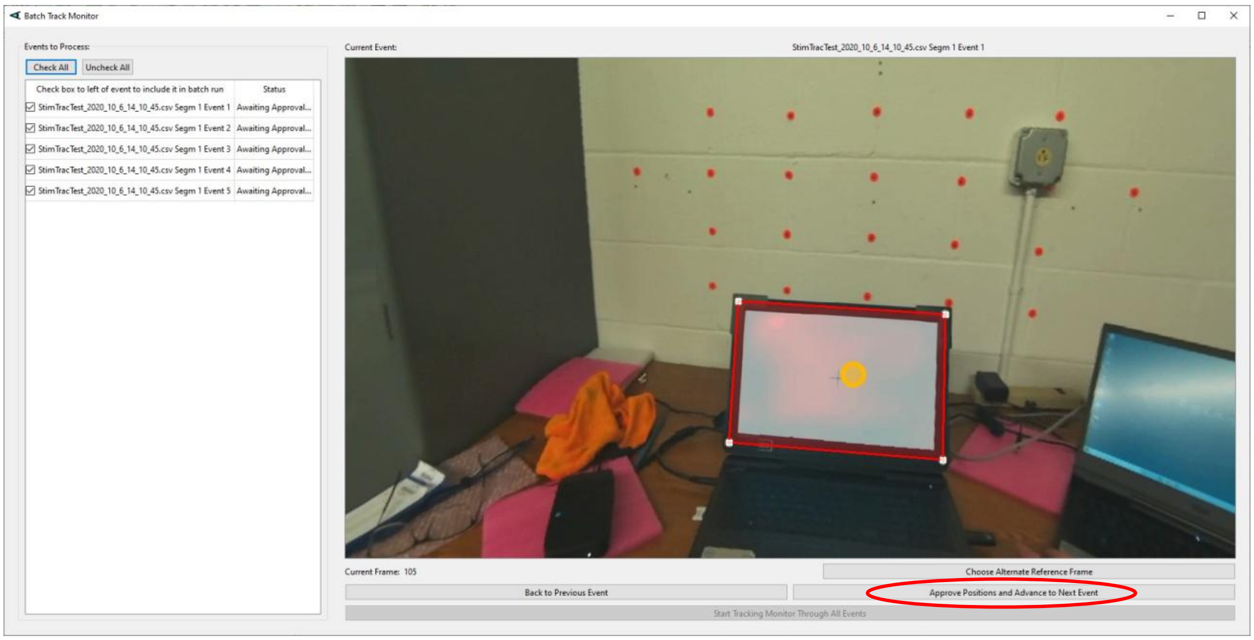

If any monitor vertex appears to be wrong, correct the vertices by dragging with the left mouse button just as described in the previous section, then click "Approve Positions and Advance to Next Event". The fist frame of the second event will be displayed. Correct the vertices if necessary, and repeat for each event in the segment. After "approving" the outline for the last event, the "Approve Positions…." button will become inactive (grayed out), and the "Start Tracking Monitor Through All Events" button will become active. Click this button to start the Monitor Tracking process. The "Speed" setting will control the frequency of "anchor" frames as described in the previous section.

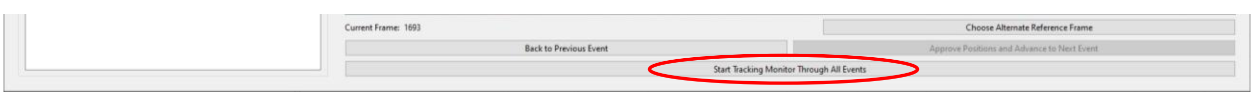

The system will track the monitor through each event. A "Batch Track Monitor" dialog will appear to show which event is being tracked.

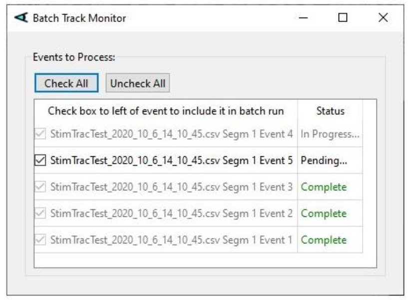

Once this process has begun, it will not be possible to pause and make manual corrections. If manual corrections are required it will be necessary to edit the monitor tracking results from the event nodes as described in the next section. When the batch tracking process is complete, a **Monitor** node will be present below each event node.

17.3.3 View or Edit Monitor Tracking results in Scene Video

Right-click a **Monitor** node, and choose “View Monitor in Scene Video” (A) to view the monitor video.

If the monitor track is incorrect for some frames and needs to be manually edited, instead of selecting “View Monitor in Scene Video”, select “Edit/Track Monitor in Scene Video” (B). Pause the video on any frame and use the left mouse button to drag incorrect vertices (C) as necessary. Moving vertices on an anchor frame modifies that anchor frame, and moving vertices on a frame between anchor points makes that frame an anchor frame. Remember that vertex point positions are interpolated between anchors. Use the next anchor or previous anchor buttons (near center right of dialog) to advance the video to the next anchor or move it back to the previous anchor.

To delete the “anchor frame" status from all frames in a continuous section of the video, pause the video on the first anchor frame to be “un-anchored” and click the button labeled “Edit Monitor Position in Consecutive Frames" (D). The button label will change to "Mark End Frame" (E). Advance to the last frame for which “anchor” status is to be retained, and click the "Mark End Frame" button. The video will jump back to the anchor frame just before the first “un-anchored” frame. All frames between this and the “End Frame” will now be un-anchored. Advance to any frame or frames in this section and adjust one or more vertex positions to create a new "anchor" at that frame.

Usually, it will be necessary to adjust monitor recognition in only a small number of frames or video sections. If a very large number of frames would need to be manually corrected, an alternate auto tracking method can be tried instead, as described in the next section.

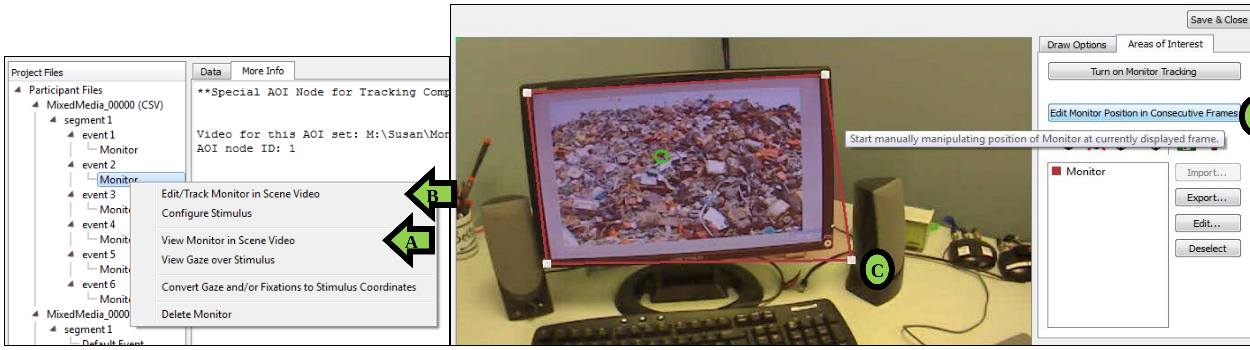

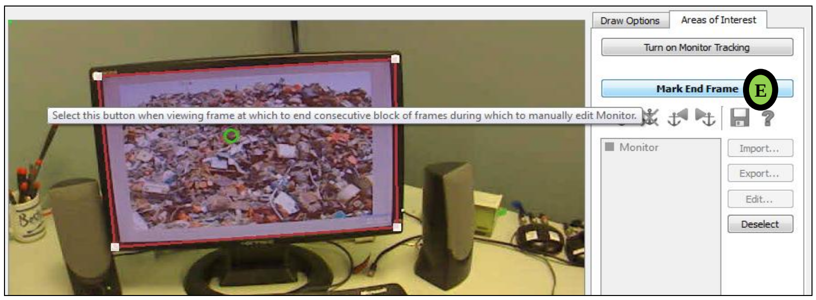

17.3.4 Alternate auto Monitor Tracking method

The auto Monitor Tracking method previously described (sections 17.3.1 and 17.3.2) attempts to track the monitor outline, and tries to ignore the monitor contents, which are often changing. The alternate method described in this section attempts to track the entire area within the monitor boundary. In some cases this may produce better results with fewer manual corrections required. If the primary method seems to result in a very large number of frames that would need to be manually corrected, the alternate can be tired instead.

First delete the monitor node created with the primary method by right clicking the Monitor node and selecting “Delete Monitor” from the drop down context menu. Right click the Event node, and select “Track Computer Monitor”. When the “Tracking Monitor” tab opens, correct the initial outline, as described in section 17.3.1. Rather than clicking the “Turn on Start Stim Trac” button, right click within the outline, and select “Auto Detect Monitor” from the context menu.

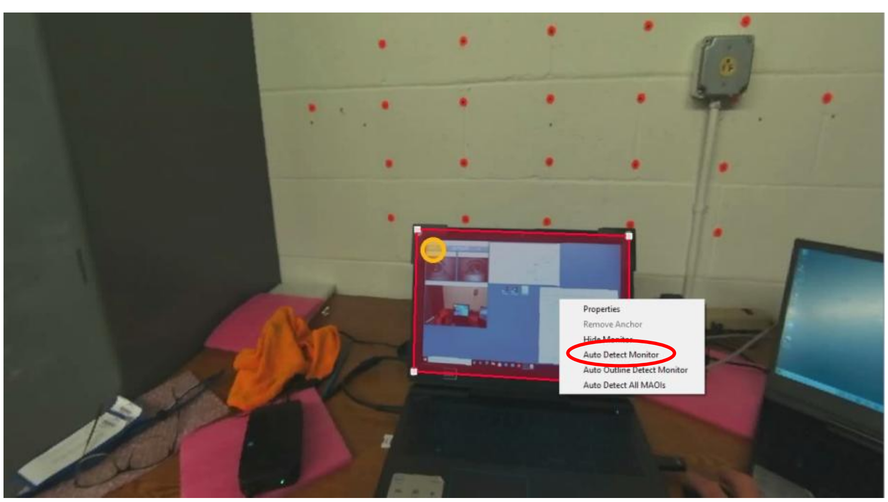

Just as for “Auto Detect MAOIs” (section 16.7.5) the program will use the image within the defined area as a “template” to try to find the same area on frames throughout the video file. A “Detecting MAOI…” window will pop up to show the area being detected as the video advances from the beginning to the end of the video. Note that the context menu includes an item labeled “Auto Outline Detect Monitor”. This selection will try to match the image near the area boundary. It will yield the same result as the “Start Monitor Tracking” method discussed in the previous sections, except that when the monitor outline cannot be found at all, the MAOI will be “hidden” until it is again identified.

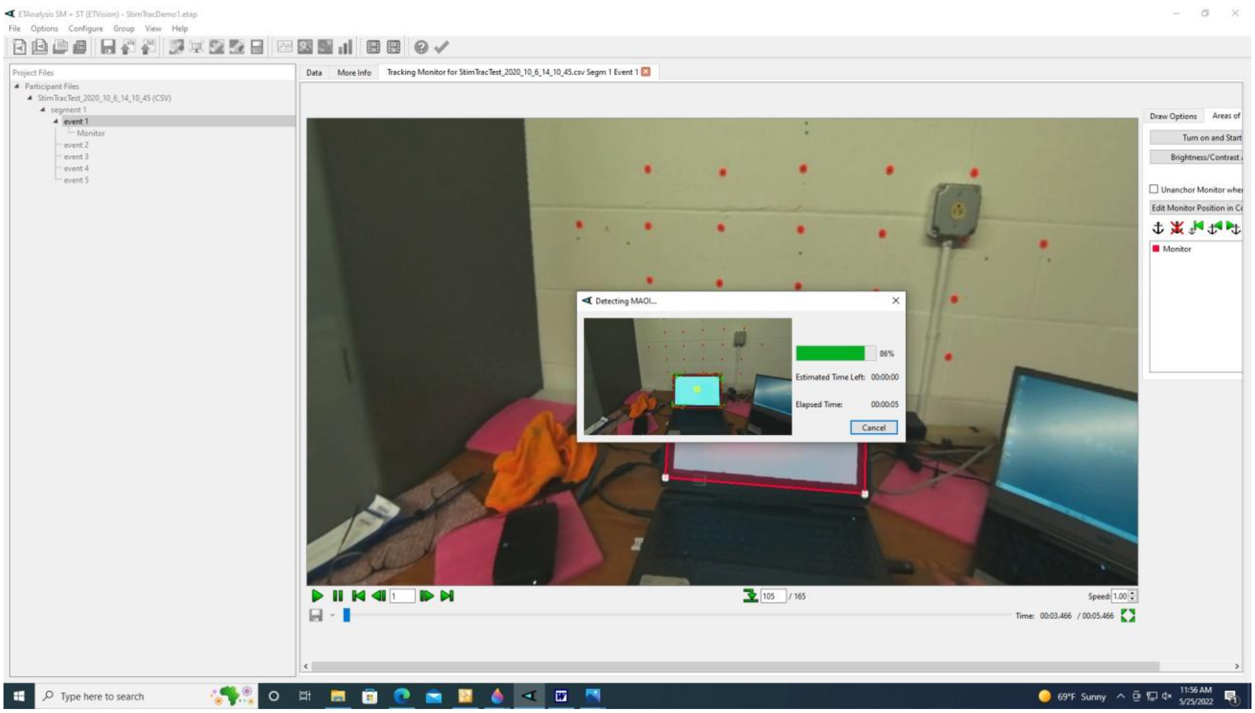

The pop up window also shows a progress bar and the estimated time remaining to advance through the video. The process cannot be interrupted for manual corrections. When the process reaches the end of the video a pop up “Finished” dialog appears, with an “OK” button. Click OK to complete the process.

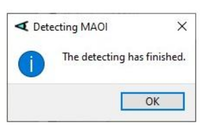

As with the primary method, the program does not attempt to match the area image content template on every video frame. If the video viewer speed is set to 1.0, it attempts to find monitor outline on every 4th frame, setting these frames as “anchor frames” and interpolating area of interest position between the anchor frames. The matching interval is always 4 times the viewer speed setting; so if the viewer speed is set to 2.0, every 8th frame will be an anchor, and if the speed is set to 0.5 every 2nd frame will be an anchor, etc. Note that the more frames used as anchors, the longer the auto-detect process will take.

Once the process is complete, a Monitor node will appear under the event node. The result can be manually corrected as described in section 17.3.3. Note that once complete, the monitor tracking results are stored the same way whether computed by either method, and there is no way to tell which method was used.

17.4 Parse Segment into multiple events if appropriate

If it is appropriate to parse the data segment into multiple events, and this was not done before the "Monitor Track" step, it can be done at this point in the process. Presumably, there is only a Default Event node under the **Segment** node, and a **Monitor** node under that. Select "Parse data" from the **Segment** node, and follow instructions from section 6. Each new **Event** node should appear with a **Monitor** node underneath it.

> *Tip: If the stimulus is a video that began and ended at the beginning and end of a segment (usually the case if using a screen capture video), and if the data segment will be parsed into multiple events, it may be most efficient to wait unit after the "Configure Stimulus" procedure (next section) to parse the segment into multiple events. The "Configure Stimulus" procedure will then need to be done only once for the segment, and the video will automatically be divided properly when the segment is parsed into multiple events.*

17.5 Import and Configure Stimulus images and/or recordings

Each image and/or video presented to participants (or screen capture videos, when applicable) must be imported to the *ETAnalysis* ST project and then configured. Backgrounds and videos can be imported globally such that they are available to the entire project, as described in the next section (section 17.5.1). Alternately the backgrounds and/or videos can be imported separately from the Configure Stimulus dialog available at each event.

In either case, the "Configure Stimulus" dialog must then be used to map the background or video to the monitor outline that was determined for each event in the "Monitor Track" step. In other words the program must determine how the background or stimulus video background corresponds to the monitor boundary outline. This is described in section 17.5.2.

17.5.1 Add Stimulus Files to project

To add static backgrounds to the project, go to the "Configure" menu and choose "Background Image(s)" or click the toolbar button. Once in the background configuration tab, select the "Create New Background" button (C) and then the "Browse" button (D) to select the first image file.

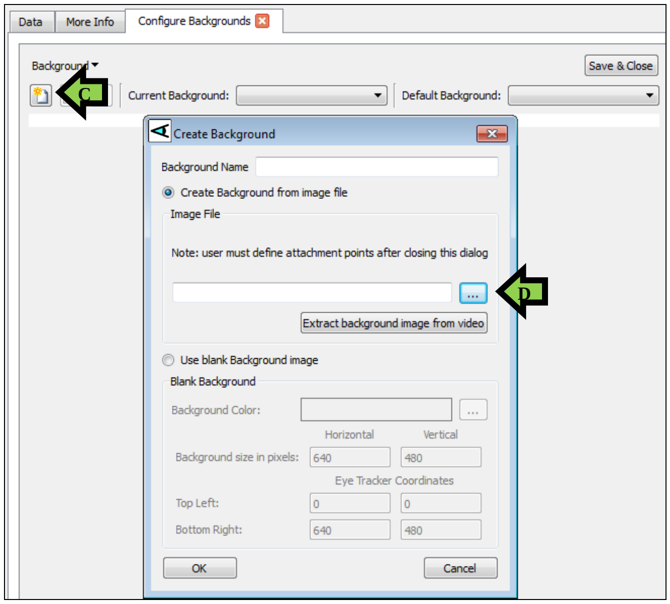

Then choose "Create multiple backgrounds" (E) and add remaining images (F). Hold down the control key to select multiple images in the File Open dialog. Then "Save and Close" (G) the Configure Backgrounds tab.

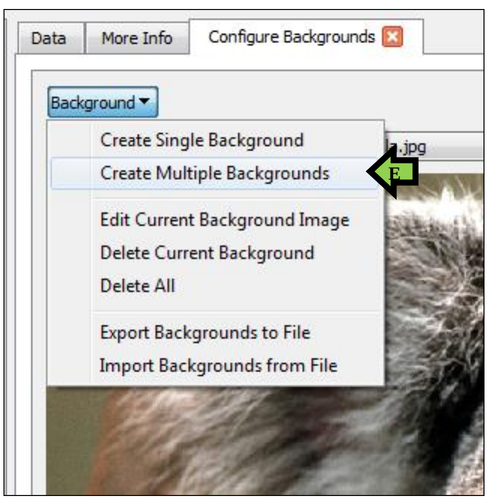

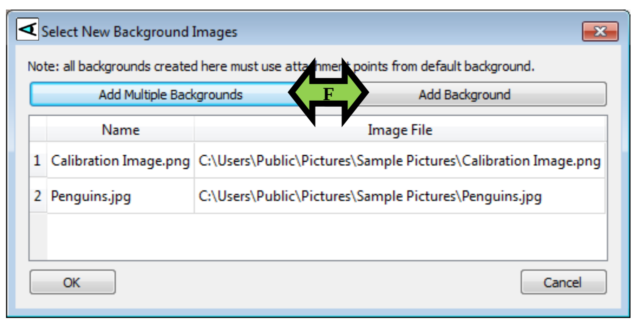

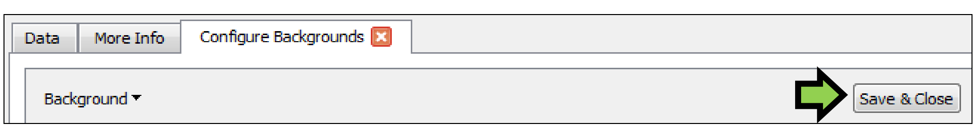

To add "canned" videos to the project, open each presented video by selecting "Open Environment Video…" from the File menu or clicking the  toolbar button. Multiple files can be selected by holding down the control or shift key or clicking and dragging across desired files.

If using screen-capture videos, each participant will be configured to a separate video and these can just be selected in the "Configure Monitor Stimulus" dialog discussed in the next section.

17.5.2 For each event, Configure Stimulus to specify the boundary of the image with respect to the tracked boundary of the monitor

In order to analyze gaze data with respect to stimuli, *ETAnalysis* ST needs to know which stimulus was presented during each event and where the stimulus was located with respect to the tracked corners of the monitor. This step does need to be performed for each event because the corners of the stimulus must be accurately defined for each event. Right-click an event Monitor node and choose "Configure Monitor Stimulus". A "Configure Monitor Stimulus" dialog will appear, and will show the first head mounted scene camera video frame from that event. Note that the image window on the dialog will show only a portion of the scene camera image, but the image window scroll bars can be used to display any portion of the image.

If the stimulus is a video that began and ended at the beginning and end of a segment (usually the case if using a screen capture video), and if the data segment will be parsed into multiple events, do the "Configure Stimulus" procedure at the Default Event node before parsing into multiple events. The procedure will then need to be done only once for the segment, and the video will automatically be divided properly when the segment is parsed into multiple events.

In the "Configure Stimulus" dialog, choose the background or video file. If these files have already been added to the project (as described in the previous section), choose the stimulus file from the corresponding pull down menu (A). Otherwise, choose the "Load background image" (B) or Video File "Browse" button (C, see Note) to select the file.

For a video stimulus that has already been loaded as an "environment video" (see previous section), set the "Use Environment Video" radio button and select the video file from the drop down menu. Otherwise, set the "Use File" radio button and browse to the desired video file. Choose whether the video should be synchronized to the event or segment (D) (i.e., does the video start and stop with the entire file – segment – or with the parsed event); often, if it's a canned video it will be synchronized with the event beginning and end, and if it's a screen-capture video it will often be synchronized with the beginning and end of the segment.

Note, if using canned videos, as opposed to screen capture videos, it is usually most efficient to add these to the project before this step (see previous section) and NOT to add them via the Video File "Browse" button on the Configure Stimulus dialog. Adding such videos globally, as described in the previous section, greatly simplifies configuring moving AOIs in the stimulus video and analyzing these moving AOIs in each event.

The "Use as project defaults" checkbox (E) pertains just to this file information and, when checked, will set the default for the next configuration to the same filename and video sync options.

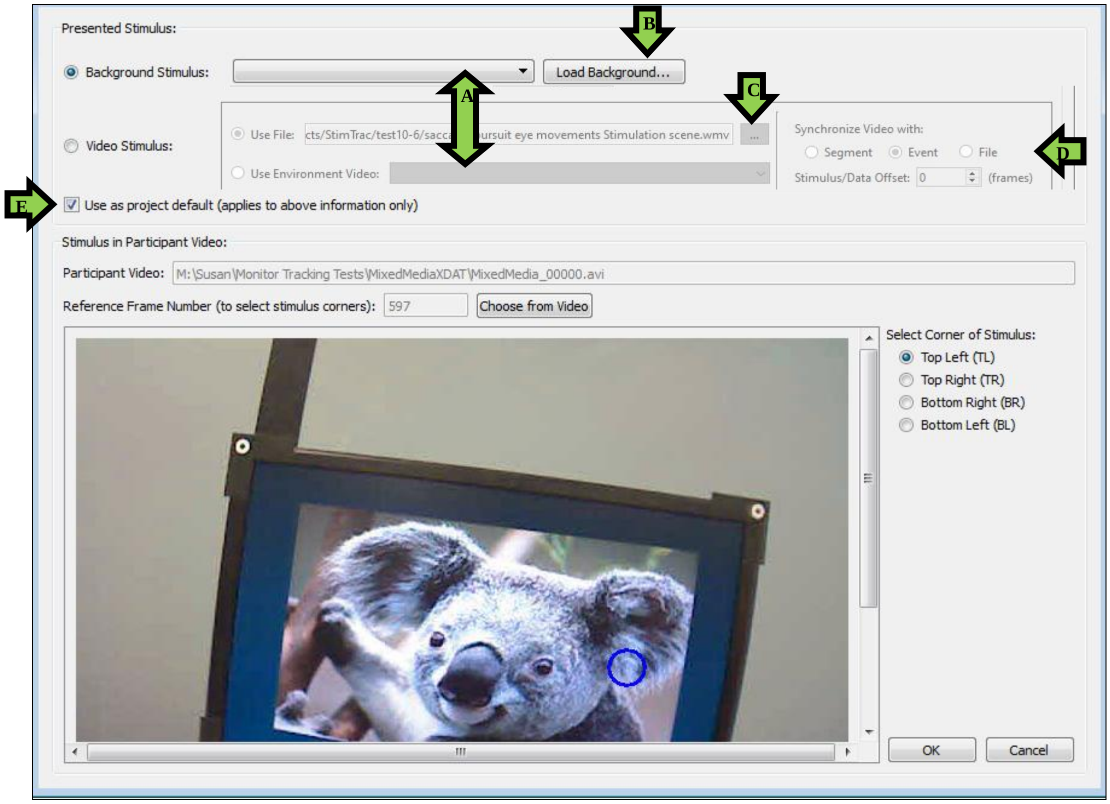

Left click the four corners of the stimulus in a frame of the participant's scene video, as shown by the red crosses on the screen shot below. The display will show a zoomed-in view of the area around the current mouse position to aid in accurate positioning (B). The size of the dialog can be expanded (if there is room on the screen), or scroll to view entire frame. If any of the four corners of the monitor or the stimulus is not visible in the currently displayed frame (by default the first frame of the event is displayed), then choose another frame via the "Choose frame" button (C). Make sure all four corners of the monitor are visible in the selected frame because the monitor position must be accurate in this frame; however, **be sure to select the corners of the <u>stimulus</u> image *which is not always same as the corners of the monitor***. The corner currently being selected should be indicated by the radio button to the right of the frame (D). Once selected, a red crosshair (E) will appear with initials corresponding to that corner (e.g., "TR" for "Top Right"). To reselect a corner, select its radio button (D), to make sure that corner is active, and use the mouse to reselect the corner.

After configuring the stimulus, a preview of the stimulus will appear in the "More Info" tab when the Monitor node is selected.

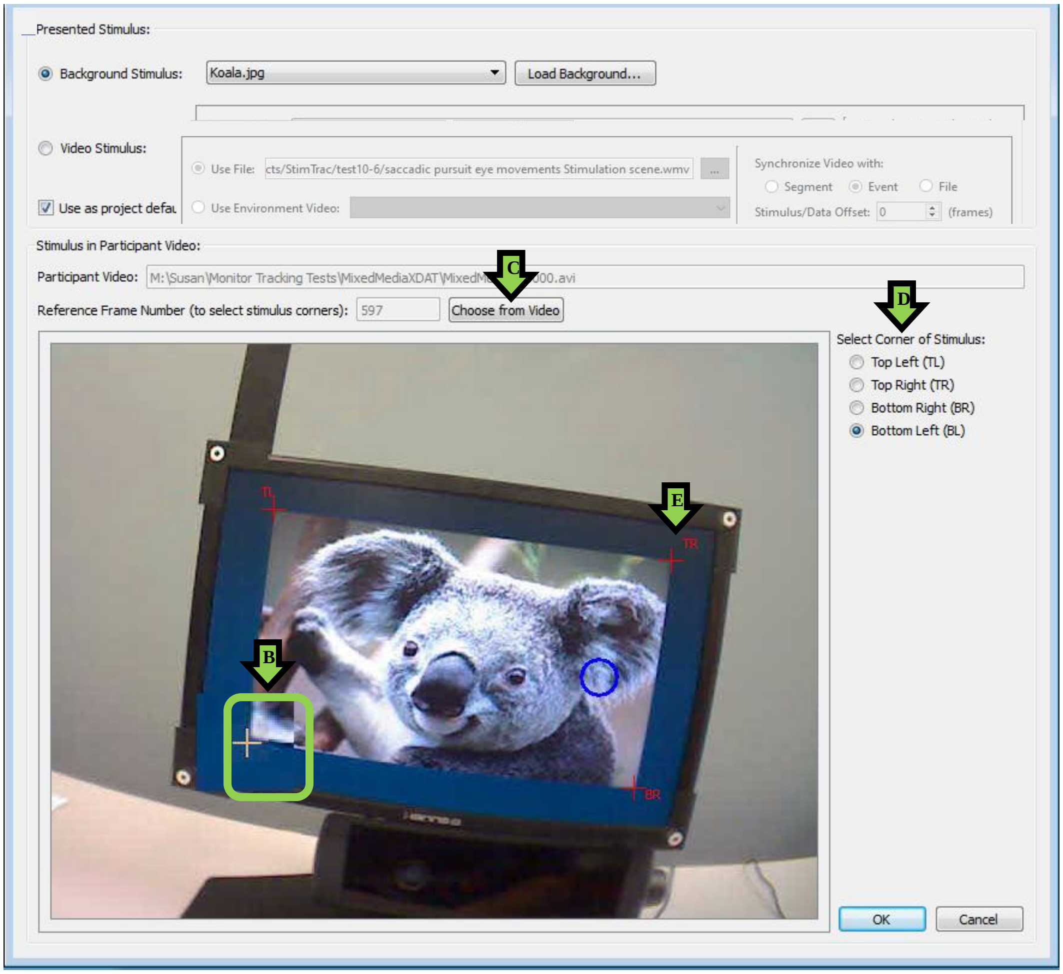

Tip: To expedite selection of stimulus file for each event in the project, it may be most convenient to configure events in order of the stimulus presented. In other words, configure all events where the Koala image is presented, with the "Use as project default" box checked, then configure all events where the "Hidden Man" video is presented, etc. In this case it will be necessary to select the corners of the stimulus and update the stimulus file information only once for each stimulus.

When configuring screen capture videos, make sure to select the corners of the displayed screen image in the scene video image. These corners may or may not line up with the inside edge of the monitor bezel (illustrated by [B] in screen shot, below). The best way to determine where these corners fall is to browse to a video frame where the edges of the screen image are most visible (via button C). If some of the corners are "underneath" the monitor bezel (i.e., cropped out of view), then the calculated gaze in stimulus may be less accurate. It is best to make sure that all corners of the presented screen image are visible when using the screen-capture method for the stimulus video.

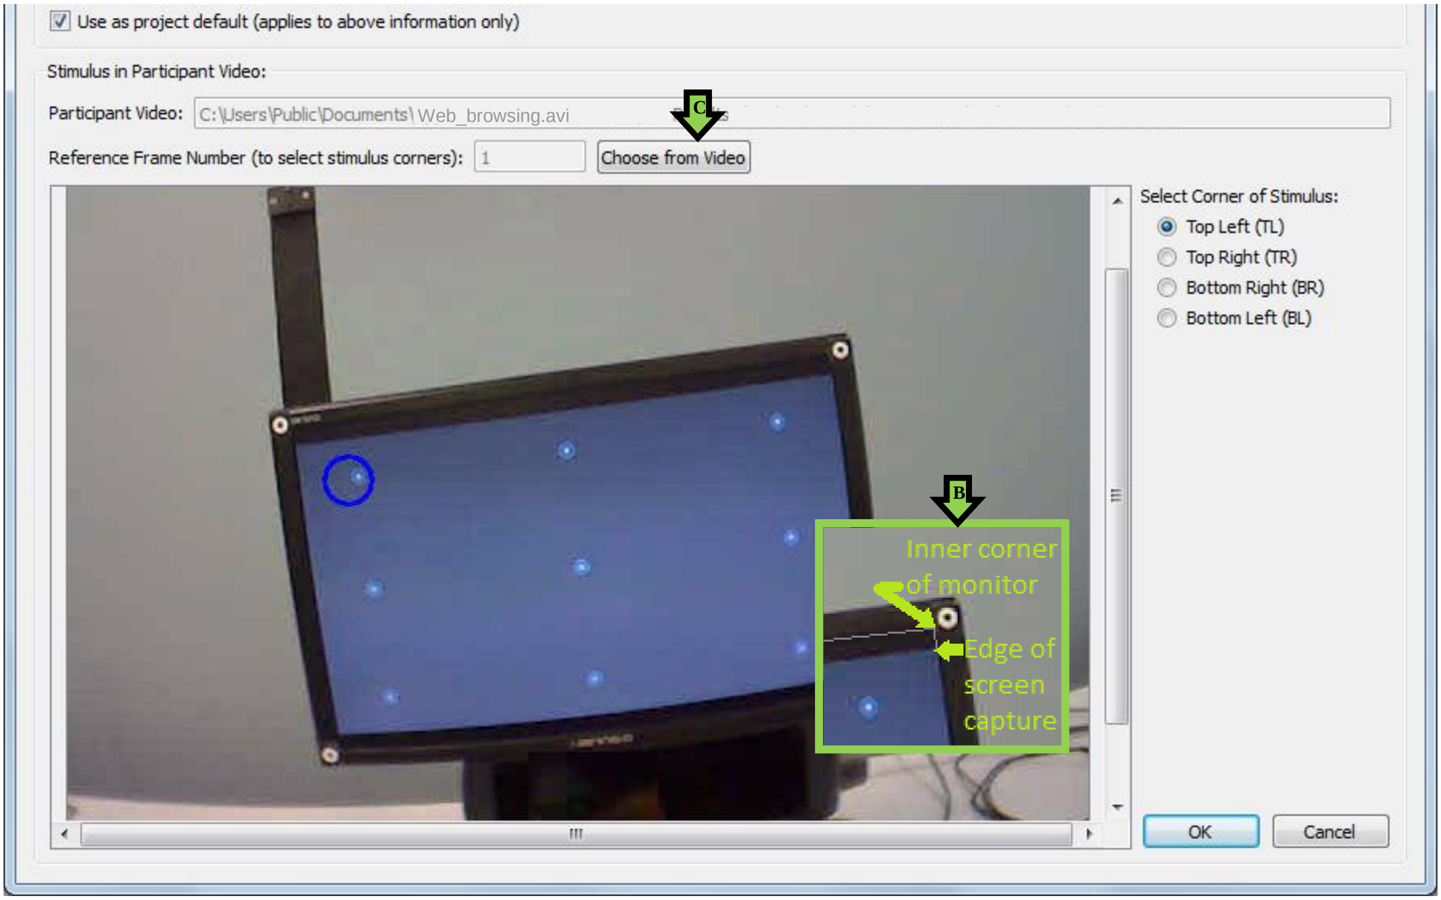

When attempting to synchronize a start and end of a screen capture video with start and end of gaze data recording, or when marking data to indicate when a stimulus display video began and ended, it may sometimes turn out that the correspondence is off by a small number of data samples. For example, a screen capture recording may start just after the first data frame is recorded rather than during the first data frame. If this type of temporal offset is apparent when viewing gaze superimposed on the video recording, or is detected in some other way, it can be corrected using the "Stimulus Data Offset" field on the "Configure Stimulus" dialog. Re-do the "Configure Stimulus" procedure, with "Stimulus Data Offset" set to the appropriate positive or negative offset.

17.6 Configure Areas of Interest in stimulus files

There are three possible scenarios for configuring Areas of Interest for Stimulus Tracking projects: 1) static AOIs in background images, 2) moving AOIs (MAOIs) in canned videos, 3) moving AOIs in screen capture videos. For the first two options, see Sections 8 and 16.7, respectively.

To configure moving AOIs in screen capture videos, first make sure the event has been configured to the screen capture video (previous Section). Then, right-click the **Monitor** node and choose "Configure Moving AOIs in Stimulus Video". Note that this option will only appear after the stimulus has been configured to a video file.

If Moving AOIs have already been configured in the stimulus, the option will read "Edit Moving AOIs in Stimulus Video". Proceed as in Section 16.7. When using screen capture video, Moving AOIs will usually need to be configured for each event (assuming each event uses a different video). If all participants viewed the same video or image files, moving AOIs should need to be configured once per stimulus file and will automatically be applied to each Monitor node that was configured to that stimulus.

It may sometimes be the case that each participant viewed the same set of objects, although they moved about differently on each scene video. Although MAOIs must be created separately for each, it may be important that each MAOI set have the same number of areas, with exactly the same names, so that Group analysis functions can be used. A feature is provided to make this task easier. After the first MAOI set is created, when creating the subsequent MAOI sets, AOI names can be imported from the first set to insure that exactly the same names are used. On the "Configuring MAOIs….." tab, choose "Import Names". Imported AOI names will appear grayed-out in the AOI list until they are created. After drawing a rectangular or polygonal AOI, a dialog box will appear presenting each AOI name of the same type (rectangle or polygon) that has not yet been selected. Chose the appropriate name for the AOI being created.

17.7 Analyze Results

17.7.1 Compute Fixation, Fixation Sequence and Dwell statistics

Gaze data used in Stimulus Tracking projects is always data that has been recorded with head mounted eye tracker optics. This original gaze data specifies gaze with respect to the head, and fixations computed using this data consider fixations to be periods during which the eye is relatively stable with respect to the head (see discussion in section 10). This is the only kind of fixation that is computed in Stimulus Tracking projects.

To compute fixations, right-click an **Event** node (or higher to compute for multiple sub-events at once) and choose "Find Fixations". Follow the instructions in section 10. To compute Fixation Sequence and Dwell statistics select "Find Fixation Sequence (Static AOIs)" or **"**Find Fixation Sequence (Moving AOIs)" from a fixation node or higher, and follow the directions in section 11 and 12, or in section 16.8.1. Note, however, that computation of fixations as periods of gaze stability with respect to MAOIs, is not available from Stimulus Tracking projects.

As in all other project types, AOI Bar Plots can be viewed after fixation sequences have been computed (see Section 14.3), and data can be combined across events as described in Section 15.

17.7.2 View Gaze, Fixations, and Fixation Sequence Statistics, over Stimulus

The participants gaze data, fixations, and fixation sequence statistics can be viewed either over the original head mounted scene camera video, or over the stimulus background or video file.

To view data over the original head mounted scene camera video, select "Play Video with Gaze" from an **Event** node, "Play Video with Fixations" from a **Fixation** node, or "Play Video with Statistics" from a **Fixation Sequence** node. Note that "Play Video with Statistics" will not show AOIs superimposed on the scene video, but in all other respects, these displays are the same as those described in section 16.9.

View the data overlaid on the stimulus video or static image file to get a cleaner, higher resolution view than from viewing gaze over the head mounted scene camera video. View data this way by selecting "Play Gaze over Stimulus" from an **Event** node, "Play Fixations over Stimulus" from a **Fixation** node, or "Play Statistics over Stimulus" from a **Fixation Sequence** node. In this case, "Play Statistics over Stimulus" does include the option to show Areas of Interest drawn on the background or video. See sections 14.3.2 and 16.9 for a description of the various display options and controls.

In addition to the controls described in sections 14.3.2 and 16.9, a smoothing filter can be applied to the data when using "Play Gaze over Stimulus", "Play Fixations over Stimulus", or "Play Statistics over Stimulus". To apply the smoothing filter, check the "Smooth" checkbox located on "Draw Options" tab. Three levels of smoothing are available, as determined by the low, medium, and high radio buttons. Smoothing will apply only to the display. It will not change the data table values or computed statistics.

17.7.3 View and Combine data across multiple events

Use selections under the "Group" menu to view data superimposed from multiple events or to compute statistics over multiple events.

If there are multiple events for which participants viewed the same static image, use the **Group → Heat map** selection to see a 2D heat map plot of combined data from all of the relevant events, as described in section 14.2.1. Use **Group → Fixation 2D plot** to see fixation points from all relevant events plotted over the static background as described by section 14.2.2.2.

Use **Group → Swarm Video over Background** selection to view moving gaze points from all of the events simultaneously superimposed on the background image. Note that this selection can also be used to create a text file showing all gaze points with respect to the stimulus image pixel coordinates. See section 15.1.1 for a more detailed description.

If there are multiple events for which participants viewed the same video presentation, use the **Group → Swarm Video over Shared Video** selection to view the gaze trail from all of the events simultaneously superimposed on the stimulus video. See section 15.1.2 for a more detailed description.

To compute fixation sequence statistics using data from multiple events use **Group → Pool Fixation Data** or **Group → Average FixSeq and Dwell Summaries** as described in section 15.2.

18 *AI Object Training* Feature (Requires AI License)

If *ETAnalysis* is equipped with the *AIObjectTraining* feature, Moving Areas of Interest can be used to provide training data for Artificial Intellegence (AI) object detection. This option requires an additional license (Consult Argus Science).

Follow instructions in the separate "ETVision AI Training" manual to create an AI training video.

Open the training video in *ETAnalysis* as an Environment video (see section 16.3). After creating and saving Moving Areas of Interest (see section 16.7) on the training video, right click the Environment→ Moving Areas of Interest node and select "View Areas of Interest". From the "Save video" pull down menu at the lower left of the display area, select "Save AI Data Set". A "Saving AI Dataset" dialog will appear with default parameters set.

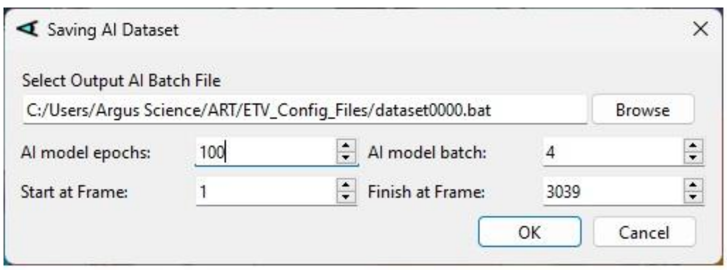

Click the "Browse" button and use the resulting browser to specify a dataset file name and location. The specified location will become the "AI project folder". See "ETVision AI Training" manual for more detail concerning the other parameters. Click "OK" to save "AI Data Set" files to the specified AI project folder.

The "AI Data Set" files can be used to train an AI object detection model as described in the "ETVision AI Training" manual. The resulting "ArgusAIModel" can be used with *ETVision* eyetracker systems to automatically detect objects in the scene camera field of view.

19 Additional *ETAnalysis* Features

19.1 Copy Project Settings from Another Project

Some Project Settings may be relevant to multiple projects and therefore *ETAnalysis* includes a tool to copy these settings between projects. Settings that can be copied include:

- Event Parsing Criteria
- Fixation Calculation Parameters
- Pupil Analysis Settings
- Time Plot Settings
- Eye Tracker units to Degrees Visual Angle
- Background Images and Attachment Points
- Static AOI Sets and/or Moving AOIs in Static Backgrounds
- Background and/or AOI Correspondences
- Fixations in Moving AOIs Preferences (compute with respect to head or AOIs)
- Fixation 2D Plot Drawing Settings
- Heat Map Drawing Settings
- Gaze Trail Drawing Settings
- Fixations in Video Drawing Settings
- Advanced Batch Criteria
- Stimulus Tracking Default Stimulus Settings

To copy any or all of these settings from one project into another, choose "Copy Settings from Another Project…" from the Configure menu. The following small dialog will appear.

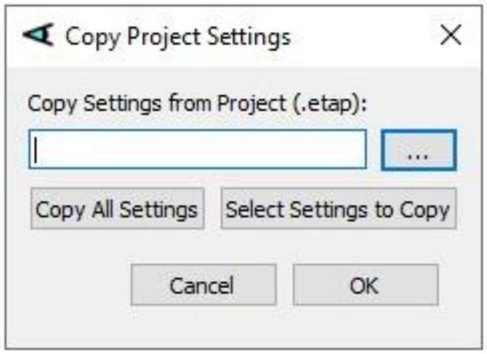

Choose the project file (with extension ".etap") corresponding to the project to be copied. This file will be located in the project folder with the same name. In most cases, it will probably be best to just "copy all settings". Left clicking the "Copy All Settings" button or the "OK" button, will cause all settings to be copied and the dialog will close. To select a subset of all settings, choose "Select Settings to Copy". A dialog will appear with a checkbox for each Settings item relevant to the project type. Check the ones to be copied and click "OK".

19.2 Export data

By using context menu Export selections, any numerical data can be exported to a text file or to an xml file that can be read by Microsoft Excel (version 2008 or higher).

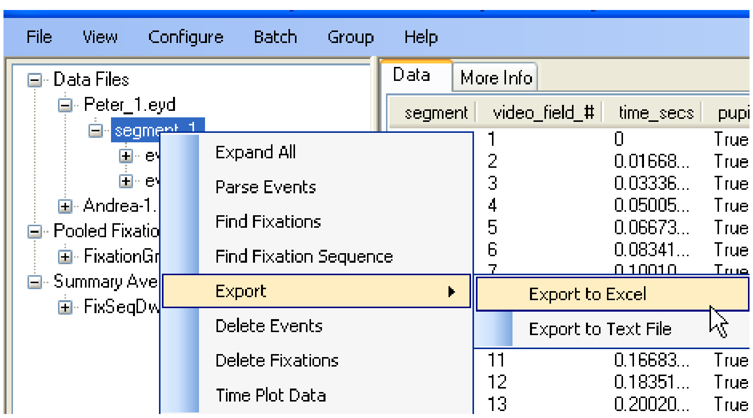

In most cases only data contained in the selected node are exported. For example, selecting Export from an event node will export only raw data for that event (not fixation data or fixation sequence statistics, etc., that may be in sub nodes). Both the contents of the "Data" window and "More Info" window are included. In the case of Excel, these are on separate "sheets".

The exceptions are "file", "fixation sequence", and "dwell" nodes. Exporting from a data file node exports all segments of raw data in the file. File information is included, but not the "More Info" pages from each segment. In the case of Excel, the file information is on one sheet, and each segment is on a separate sheet. Exporting from a "fixation sequence" node does include the statistics sub-nodes underneath it (each on a separate Excel sheet). Similarly, exporting from a Dwell node includes the subordinate statistics nodes.

It may sometimes be desirable to export data from multiple nodes to a single Excel spread sheet or text file list. This can be done from the **File→Export to Excel** or **File→Export to Text File** menu selections. Hovering the mouse over one of these selections shows a list of node types.

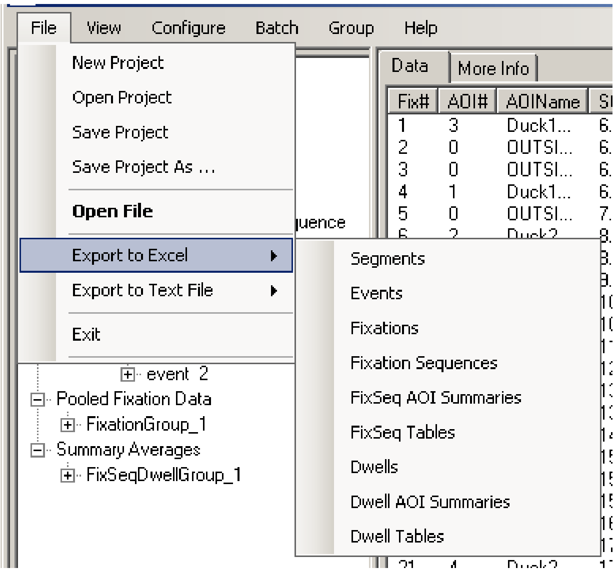

Clicking on one of the node types brings up a tree diagram showing all the nodes of that type in the project. Each node of the tree is a check box.

For example, clicking "Fixations" brings up a diagram like the one shown below. Checking a fixation node selects only that node. Checking a higher-level node selects all of the fixation nodes underneath it. The dialog also has a "Select all" button and an "Unselect all" button.

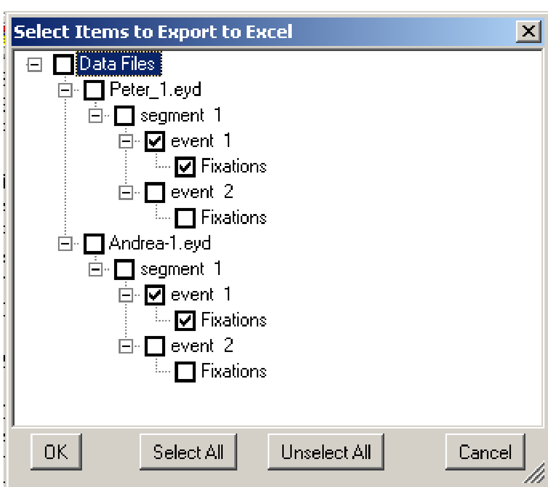

When "OK" is clicked all selected fixation nodes will be exported to Excel, and in this case will all be listed on a single Excel sheet. Columns containing the file name, segment number, and event name are included so that the origin of the data in each row is specified. This works in a similar fashion for whichever type of node was selected. Contents of the "More Info" pages are not included. Only the data lists are exported.

19.3 Save images and video displays

Any graphic image displayed in the *ETAnalysis* display area can be saved as a JPEG, GIF,TIFF, PNG or BMP file (user selectable). Static displays, such as 2D fixations plots, always include a Save Image Icon . Click the Icon and use the resulting Browser to specify the file type, file name, and location for the saved image file. All dynamic displays, such as Play Fixations over Background, appear as video viewers in the *ETAnalysis* display area, and include a Save Video Icon . Clicking the down arrow opens a pull down menu with the choice of recording a video file, or capturing and saving a single frame as an image file.

19.4 Do All Calculations

Click the calculator icon on the *ETAnalysis* shortcut bar to perform all possible fixation, fixation sequence, and pupil diameter computations on all event nodes. Current configuration settings and area of interest correspondences will be used for all nodes. Caution: use this shortcut only if sure that all configurations and correspondences are set as desired for all nodes, and if all possible *ETAnalysis* computations are desired.

19.5 Check for Updates

The *ETAnalysis* program does not update automatically when new versions are released. Users are encouraged to check for updates regularly. There are multiple ways to check whether the current version is out of date: 1) Select from the main toolbar, 2) Choose "Check For Update" from the Help menu, or 3) Open the "About *ETAnalysis*" dialog (also accessed via the Help menu). If an update is available, a link will be provided for downloading the installation file via the Internet.

20 Appendix – *ETVision*, *ET3Space* feature

An optional addition to the Argus Science *ETVision* eyetracker, called *ET3Space*, uses head position and orientation data from one of several third party motion tracking devices to compute gaze with respect to a room fixed coordinate system and to determine point-of-gaze on multiple surfaces in the environment. This optional feature is described in detail in a separate manual. Note that *ET3Space* is an *ETvision* option, not an *ETAnalysis* option. *ETAnalysis* is always equipped to process *ET3Space* data.

Binary data files recorded by *ETVision* with the *ET3Space* feature enabled have a ".ehd" extension. Text (".csv") data files recorded by *ETVision* include text on the first line specifying whether *ET3Space* was enabled ("3Space: Yes") or not enabled ("3Space: No").

On *ETVision* files recorded with *ET3Space* <u>not</u> enabled, gaze is recorded only as a position on the head mounted scene camera video image. On files recorded <u>with</u> *ET3Space* enabled the data also specify which fixed surface in the environment was intersected by gaze as well as the intersection point with respect to a coordinate frame attached to that surface.

When *ET3Space* data is available, *ETAnalysis* uses the *ET3Space* gaze data (gaze position on an environment surface) to compute fixations. For fixation sequence analysis, Areas of Interest must therefore be specified with respect to the same environment surface coordinates.

See the separate *ET3Space* manual for a description of all data items available when the *ET3space* feature is used.