

# Eyetracking for two-person tasks with manipulation of a virtual world

JEAN CARLETTA, ROBIN L. HILL, CRAIG NICOL, AND TIM TAYLOR  
*University of Edinburgh, Edinburgh, Scotland*

JAN PETER DE RUITER  
*Max Planck Institute for Psycholinguistics, Nijmegen, The Netherlands*

AND  
ELLEN GURMAN BARD  
*University of Edinburgh, Edinburgh, Scotland*

Eyetracking facilities are typically restricted to monitoring a single person viewing static images or pre-recorded video. In the present article, we describe a system that makes it possible to study visual attention in coordination with other activity during joint action. The software links two eyetracking systems in parallel and provides an on-screen task. By locating eye movements against dynamic screen regions, it permits automatic tracking of moving on-screen objects. Using existing SR technology, the system can also cross-project each participant's eyetrack and mouse location onto the other's on-screen work space. Keeping a complete record of eyetrack and on-screen events in the same format as subsequent human coding, the system permits the analysis of multiple modalities. The software offers new approaches to spontaneous multimodal communication: joint action and joint attention. These capacities are demonstrated using an experimental paradigm for cooperative on-screen assembly of a two-dimensional model. The software is available under an open source license.

Monitoring eye movements has become an invaluable method for psychologists who are studying many aspects of cognitive processing, including reading, language processing, language production, memory, and visual attention (Cherubini, Nüssli, & Dillenbourg, 2008; Duchowski, 2003; Griffin, 2004; Griffin & Oppenheimer, 2006; Meyer & Dobel, 2003; Meyer, van der Meulen, & Brooks, 2004; Rayner, 1998; Spivey & Geng, 2001; Trueswell & Tanenhaus, 2005; G. Underwood, 2005; Van Gompel, Fischer, Murray, & Hill, 2007). Although recent technological advances have made eyetracking hardware increasingly robust and suitable for more active scenarios (Land, 2006, 2007), current software can register gaze only in terms of predefined, static regions of the screen. To take eyetracking to its full potential, we need to know what people are attending to as they work in a continuously changing visual context and how their gaze relates to their other actions and to the actions of others. Although present limitations simplify data collection and analysis and call forth considerable ingenuity on the part of experimenters, they invite us to underestimate the real complexity of fluid situations in which people actually observe, decide, and act. At present, we are only beginning to understand how people handle multiple sources of external information, or multiple communication modalities. Despite growing interest in the interactive processes involved in human dialogue (Pickering

& Garrod, 2004), the interaction between language and visual perception (Henderson & Ferreira, 2004), how visual attention is directed by participants in collaborative tasks (Bangerter, 2004; Clark, 2003), and the use of eye movements to investigate problem solving (Charness, Reingold, Pomplun, & Stampe, 2001; Grant & Spivey, 2003; J. Underwood, 2005), the generation of suitably rich, multimodal data sets has up until now been difficult.

There is certainly a need for research of such breadth. Complex multimodal signals are available in face-to-face dialogue (Clark & Krych, 2004), but we do not know how often interlocutors actually take up such signals and exploit them online. Examining each modality separately will give us an indication of its potential utility but not necessarily of its actual utility in context. Single-modality studies may leave us with the impression that, for joint action in dialogue or shared physical tasks, all instances of all sources of information influence all players. In some cases, we tend to underestimate the cost of processing a signal. For example, we know that some indication of the direction of an interlocutor's gaze is important to the creation of virtual copresence, and that it has many potential uses (Cherubini et al., 2008; Kraut, Gergle, & Fussell, 2002; Monk & Gale, 2002; Velichkovsky, 1995; Vertegaal & Ding, 2002). Controlled studies with very simple stimuli, however, have shown that processing the gaze of another is not a straightforward

---

J. Carletta, j.carletta@ed.ac.uk

---



ward bottom-up process. Instead, it interacts with what the viewer supposes the gazer might be looking at (Lobmaier, Fischer, & Schwaninger, 2006). In genuine situated interaction, there are many sources of such expectations, and all require some processing on the part of an interlocutor. General studies of reasoning and decision making (Gigerenzer, Todd, & ABC Research Group, 1999) have suggested that people have astute strategies for circumventing processing bottlenecks in the presence of superabundant information. It would be surprising if they did not streamline their interactions in the same way. To know how, we need to record individuals' dividing their attention between the fluid actions of others, their own attempts at communication, and the equally fluid results of a joint activity.

Given two eyetrackers and two screens, four breakthroughs are required before the technology can be usefully adapted to study cooperation and attention in joint tasks. First, one central virtual world or game must drive both participants' screens, so that both can see and manipulate objects in the same world. Second, it must be possible to record whether participants are looking at objects that are moving across the screen. Third, it must be possible for each participant to see indications of the other's intentions, as they might in real face-to-face behavior. For on-screen activities, those intentions would be indicated by icons representing their partner's gaze and mouse location. Finally, to give a full account of the interactions between players, the eyetracking, speech, and physical actions of the 2 participants must be recorded synchronously. Finding solutions to these problems would open up eyetracking methodology not only to studies of joint action but also to studies of related parts of the human repertoire—for example, joint viewing without action, competitive rather than collaborative action, or learning what to attend to in the acquisition of a joint skill. In the present article, we will describe solutions to these problems using the SR Research Eyelink II platform. The resulting software is available under an open source license from <http://wcms.inf.ed.ac.uk/jast>. We will demonstrate the benefits of the software within an experimental paradigm in which 2 participants jointly construct a figure to match a given model.

## Review

Commercial and proprietary eyetracking software reports where a participant is looking, but only in absolute terms or by using predefined static screen regions that do not change over the course of a trial. The generic software for the Eyelink II tracker supplied with the system (SR Research, n.d.) provides a good basis for going forward. It can be used to generate raw eye positions and will calculate blinks, fixations, and saccades, associating fixations with static, user-defined regions of the screen. It does not support the definition of dynamic screen regions that would be required to determine whether the participant is looking at an object in motion. Nor does it support the analysis of eye data in conjunction with alternative data streams such as language, video, and audio data, or eye stream data from two separate machines. It will, however, take messages generated by the experimental software and add them to the rest of the synchronized data output, and

it can pass messages to other Eyelink trackers on the same network. We use this facility to implement communication between our eyetrackers. SR Research, the makers of the Eyelink II, provided us with sample code that displays the eye cursor from one machine on the display of another by encoding it as a string on the first machine, sending the string as a message to the second machine, parsing the string back into an eye position on the second machine, and using that information to change the display (Brennan, Chen, Dickinson, Neider, & Zelinsky, 2008).

Despite the lack of commercial support for dual eyetracking and analysis against other ways of recording experimental data, some laboratories are beginning to implement their own software. In one advance, a series of eye movements (scanpaths) produced by experts are later used to guide the attention of novices (Stein & Brennan, 2004). In another, two eyetrackers are used in parallel with static displays but without cross-communication between one person's tracker and the other's screen (Hadelich & Crocker, 2006; Richardson, Dale, & Kirkham, 2007; Richardson, Dale, & Tomlinson, 2009). In a third, one participant's scanpath is simulated by using an automatic process to display a moving icon while genuinely eyetracking the other participant as he or she views the display (Bard, Anderson, et al., 2007). In a fourth, Brennan et al. (2008) projected genuine eye position from one machine onto the screen of another while participants shared a visual search task over a static display. Steptoe et al. (2008) and Murray and Roberts (2006) keyed the gaze of each avatar in an immersive virtual environment to actual tracked gaze of the participant represented by the avatar, but without a shared dynamic visual task.

## Data Capture

Dual eyetracking could be implemented using any pair of head-free or head-mounted eyetrackers, as long as they can pass messages to each other. Our own implementation uses two head-mounted Eyelink II eyetrackers. Eyetracking studies start with a procedure to calibrate the equipment that determines the correspondence between tracker readings and screen positions, and they usually incorporate periodic briefer recalibrations to correct for any small drift in the readings. The Eyelink II outputs data in its proprietary binary format, "EDF," which can either be analyzed with software supplied by the company or be converted to a time-stamped ASCII format that contains one line per event. The output normally contains a 500-Hz data stream of eye locations with additional information about the calibration and drift correction results used in determining those locations, plus an online parsed representation of the eye movements in terms of blinks, fixations, and saccades. In addition to this data originating from the eyetracker itself, the eyetracker output will contain any messages that have been passed to the eyetracker from the experimental software, stamped with the time they were received. An Eyelink II natively uses two computers that are connected via a local network. The host machine drives the eyetracking hardware by running the data capture and calibration routines, and the display machine runs the experimental software. Our installation uses two Pentium 4 3.0-GHz display machines

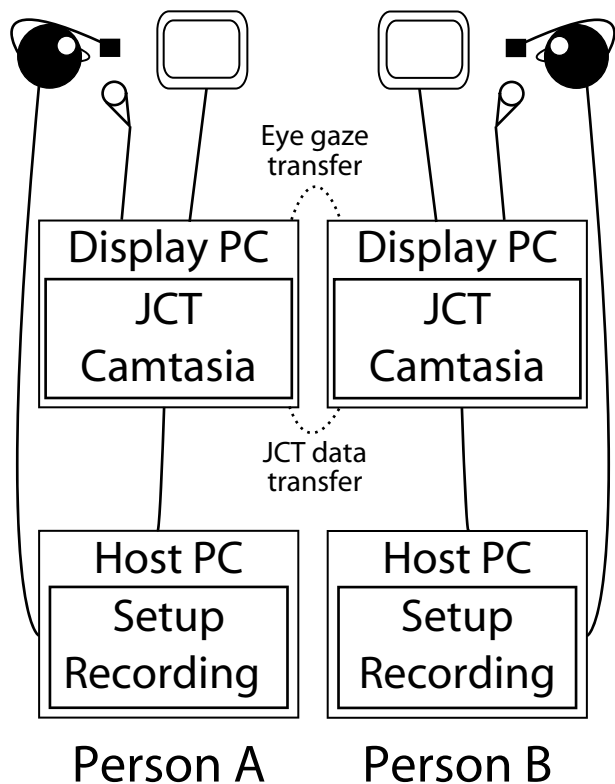


Figure 1. Hardware in the joint construction task (JCT) experimental setup.

running Windows XP with 1-GB DDR RAM, a 128-MB graphics card, a 21-in. CRT monitor, a Gigabit Ethernet card, and a Soundblaster Audigy 2 sound card, and two Pentium 4 2.8-GHz host machines running Windows XP and ROMDOS7.1 for the Eyelink II control software, with 512-MB DDR RAM and a Gigabit Ethernet card.

Our arrangement for dual eyetracking is shown in Figure 1. Here, because there are two systems, four computers are networked together. In addition to running the experimental software, the display machines perform audio and screen capture using Camtasia (TechSmith, n.d.) and close-talking headset microphones. Audio capture is needed if we are to analyze participant speech; screen capture provides verification and some insurance against failure in the rest of the data recording by at least yielding a data version that can easily be inspected, although the resulting videos have no role in our present data analysis.

As usual, the display machines are networked to their respective hosts so that they can control data recording and insert messages into the output data stream but, in addition, the display machines pass messages to each other. These are used to keep the displays synchronized. For instance, if Participant A moves his or her eyes, his mouse, or some on-screen object, the experimental software running on A's display machine will send a message to that effect to the experimental software on Participant B's display machine, which will then update the graphics to show A's gaze cursor, A's mouse cursor, and the shared object in their new locations. Coordinating the game state is sim-

ply a matter of passing sufficient messages. There is, of course, a lag between the time when a message is generated on one participant's machine and the time when it is registered on the other's. In our testing, 99% of the lags recorded were less than 20 msec. In pilot experiments, debriefed participants did not notice any lag. Even with the small additional time needed to act on a message to change the display, this degree of lag is tolerable for studies of joint activity. During any trial, the experimental software should loop through checking for local changes in the game and passing messages until both of the participants have signaled that the trial has finished. In our experimental software, each loop takes around 42 msec. Since eye positions move faster than do other screen objects, we pass them twice during each loop. Whenever a display machine originates or receives a message, it sends a copy of that message to its host machine for insertion into its output data stream. As a result, the output of both eyetracking systems contains a complete record of the experiment, although the eye positions are sampled more coarsely on the opposing than on the native machine.

### Example

**The experimental paradigm.** In our experimental paradigm, 2 participants play a series of construction games. Their task is to reproduce a static two-dimensional model by selecting the correct parts from an adequate set and joining them correctly. Either participant can select and move (left mouse button) or rotate (right mouse button) any part not being grasped by the other player. Two parts join together permanently if brought into contact while each is "grasped" by a different player. Parts break if both participants select them at the same time, if they are moved out of the model construction area, or if they come into contact with an unselected part. Any of these "errors" may be committed deliberately to break an inadequate construction. New parts can be drawn from templates as required.

These rules are contrived to elicit cooperative behavior: No individual can complete the task without the other's help. The rules can easily be reconfigured, however, or alternative sets of rules can be implemented. Figure 2A shows an annotated version of the initial participant display from our implementation of the paradigm with labels for the cursors and static screen regions. Figure 2B shows a later stage in a trial in which a different model (top right) is built. Initial parts are always just below the model. A broken part counter appears at the top left, a timer in the middle at the top, and new part templates across the bottom of the screen. After each trial, the participants can be shown a score reflecting accuracy of their figure against the given model.

This experimental paradigm is designed to provide a joint collaborative task. The general framework admits a number of independent variables, such as the complexity of the model, the number of primary parts supplied, the difficulties that are presented by the packing of the initial parts, and whether or not the participants have access to each other's speech, gaze cursor, or mouse position. These variables can all be altered independently from trial to trial. The paradigm is suitable for a range of research interests in the area of joint activity. Performance

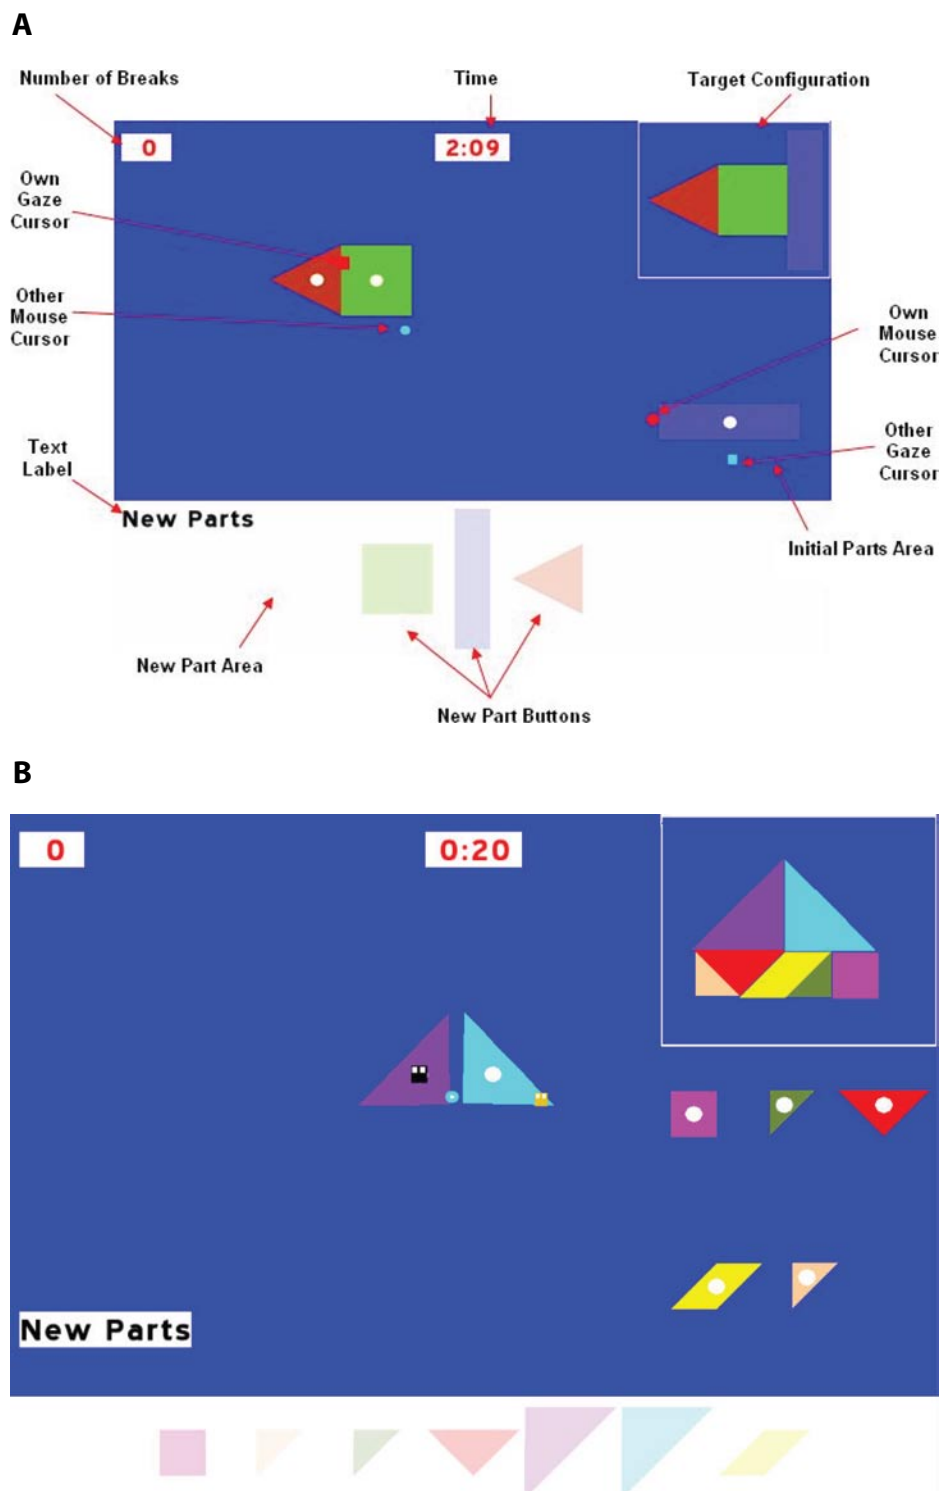


Figure 2. Annotated screen shots of two variants of the JAST joint construction task. (A) Initial layout of a construction task. (B) Screen shot 20 sec into a trial in a tangram construction task.

can be automatically measured in terms of time taken, breakages, and accuracy of the constructed figure against the target model (maximum percentage of pixels that are colored correctly when the constructed figure is laid over the model and rotated). In addition, the paradigm allows for an analogous task for individuals that serves as a useful control. In the individual “one-player” version, two parts join together when they touch, even if only one is selected.

**The experimental software.** The JAST joint construction task, or JCT, is software for running experiments that fit this experimental paradigm. It is implemented under Windows XP in Microsoft Visual C++.Net and draws on a number of open source software libraries. These include Simple DirectMedia Layer (SDL) support for computer graphics (Simple DirectMedia Layer Project, n.d.); add-ons available for download along with SDL and from the SGE project (Lindström, 1999) that provide further support for things like audio, rotation, collision detection, and text; Apache’s XML parser, Xerces (Apache XML Project, 1999) and the Simple Sockets Library for network communication (Campbell & McRoberts, 2005), which we use instead of the Eyelink II networking support specifically so that the software can also be run without eyetracking. When the experimental software is run without an eyetracker, rather than passing messages to a host machine, the display machine stores the time-stamped messages in an ASCII file that follows the same structure that the eyetracker itself produces, so that the file can be used in the same onward processing. Even without the richness of eye-movement data, collaborative behavior can therefore still be investigated with this software using only two standard networked PCs.

The configuration of an experiment is defined by a set of extensible markup language (XML) files. Each experiment consists of a number of trials using one model per trial. Models are built of polygon parts. Curved surfaces are simulated by a polygon with a high number of vertices. Our software includes a utility to generate these curved parts. The initial and target configuration for each model, and descriptions of the polygon parts used, are stored in a “stimulus set” XML file, and the experiment configuration links to a number of these files, which are presented in order. For each trial, the experiment configuration file specifies the stimulus set, whether to show the clock, whether this clock counts up or down, the time limit, whether the position of the partner’s eye and mouse should be visible, and any text or graphics that should be shown between trials. It also specifies the machines on which the experiment will run, whether the experiment is to be run for individuals or pairs of participants, the size and location of the static screen regions, and a number of experimenter-composed text messages to display at certain points in the experiment. When run, the software performs the eyetracker calibration and then presents the trials in the order specified, performing drift correction between trials.

In this implementation of the experimental paradigm, the messages passed between the eyetrackers, and therefore stored in the output, are the following.

1. Markers showing the time when each trial starts and ends, along with the performance scores for the number of breakages and the accuracy of the final construction.

2. Markers near the beginnings of trials showing when a special audio tone and graphical symbol were displayed, to make it easier to synchronize the data against audio and video recordings.

3. Participant eye positions.

4. Sufficient information to reconstruct the model-building state, including joins, breakages, part creation, and any changes in the locations or rotations of individual parts.

Our video reconstruction utility takes the ASCII data from one of the eyetrackers and uses it to create a video that shows the task with the eye and mouse positions of the 2 participants superimposed. The videos produced from the two ASCII data files for the same experiment are the same, apart from subtle differences in timing: Each video shows the data at the time that it arrived at, or was produced by, the individual eyetracker used as its source. The utility, which is again written in Microsoft Visual C++.Net, uses FFmpeg (Anonymous, n.d.) and libraries from the JCT experimental software to produce MPEG2 format videos with a choice of screen resolutions and color depths.

**Interpreting participant behavior.** Capturing the data with the methods described creates a complete and faithful record of what happened during the experiment, but using primitives below the level of abstraction required for analysis. For instance, the messages describe the absolute positions of eyes, mice, and all parts throughout the task, but not in terms of when a participant is looking at a particular part, even though knowing this is essential to understanding how gaze, speech, and action relate. Therefore, the next step is to add this higher level of analysis to the data set. To do this, we transfer the data out of the ASCII format exported by the Eyelink II and into an XML format. XML format allows us to use existing parsers and makes it easy to check that the data conform to our expectations by validating data files against a document-type definition specifying what sorts of events and objects they should contain.

The resulting XML file contains most of the information gathered so far in the experiment. Drawing data from the output of both eyetrackers and from the experiment configuration file, it includes all of the messages that were passed, parsed eye data for both participants, plus a list of the parts, their movements, and their join events, with part properties—such as shape and initial locations—taken from the experiment configuration file. Although it would not be difficult to include the full 500-Hz eye position data in the XML representation, it is now omitted for several reasons. First, this level of detail is better suited for the kind of parsing provided by the eyetracker software itself. Second, including it would increase file size substantially. Third, it is unnecessary for our intended analysis.

In addition to reencoding these kinds of data from the experimental software, the utility that produces the XML file adds a number of data interpretations:

*Look events.* During these, a participant looked at a part, composite, or static region (typically referred to as regions of interest, or ROIs). Where gaze is on a mov-

able object (a dynamic ROI, or DROI), these events cover both the entire time that the eye position is within a configurable distance of the moving screen region associated with whatever is being looked at, and whether the eye is currently engaged in a fixation, smooth pursuit, or a saccade. Since the XML file contains parsed eye movements, the class of eye activity can be established by later analysis, if required.

*Hover events.* During these, a participant located the mouse over a part or composite without “grasping” or moving the part by buttonpress. Hover events also cover the entire time that the mouse cursor is within a configurable distance of the moving screen region associated with the part.

*Construction history for each trial.* This is a description of how the pair constructed the figure, which is given as a set of binary trees where the leaves are the parts, each with a unique identifier. Each node uniquely identifies a composite that is created by joining any other composites or parts that are the children of the node.

*Composition phases.* These divide the time taken to create each composite into phases. The final phase covers the final action of docking the last piece of the composite, and two earlier phases simply divide the rest of the composition time in half. Like the construction history, this representation can be used to study construction strategy or to find subtasks of particular difficulty.

*Per-trial breakage scores.* These are calculated over the XML format for use as a parity check against those coming directly from the experimental software.

The JastAnalyzer software that performs this processing requires the same environment as does the JCT experimental software. It works by passing the ASCII data back through libraries from the JCT software in order to interpret the task state. Because we exploit the resulting XML data for a number of different purposes, we call this our “general data format,” or GDF. It is straightforward, given the data format, to produce scripts that show useful analyses, such as the lag between when one participant looks at an object and when the other participant follows. It is also easy to produce per-trial summary statistics, such as how many times a participant looked at each of the parts and what percentage of the time he or she spent looking at the clock.

**Export to data display and analysis packages.** Thus far, we have seen how the automatically collected JCT data are represented. For more complicated data analysis, we export from our GDF into the format for existing data display and analysis packages. These allow the data to be inspected graphically, “playing” them synchronized to the video and audio records. Such packages are useful both for checking that the data are as expected and for browsing with an eye to understanding participant behavior. But the real strength of such packages lies in the ability to create or import other sources of information for the same data, such as orthographic transcription, linguistic annotations relating to discourse phenomena, and video coding. Some packages also provide data search facilities that can go beyond the simple analyses generated with our GDF-based scripts—for example, investigations of the relationship between task behavior and speech. Our

software includes export to two such packages: ELAN (MPI, n.d.) and the NITE XML Toolkit, or NXT (Language Technology Group, n.d.). Both can display behavioral data in synchrony with one or more corresponding audio and video signals, and both support additional data annotation. ELAN has strengths in displaying the time course of annotations that can be represented as tiers of mutually exclusive, timestamped codes, whereas NXT has strengths in supporting the creation and search of annotations that relate to each other both temporally and structurally, as is usual for linguistic annotations built on top of orthographic transcription.

Both ELAN and NXT use XML in their data formats. Because our utilities for export to these programs are based on XSLT stylesheets (World Wide Web Consortium, 1999)—the standard technique for transducing XML data from one format to another—they will run on any platform, as long as a stylesheet processor is installed.

## EXAMPLE EXPERIMENT AND ANALYSES AFFORDED

### Method

**Task.** We can illustrate the utility of the JCT systems via a designed corpus of time-aligned multimodal data (eye movements, actions, and speech) that were produced while pairs of individuals assembled a series of tangrams collaboratively. Produced as part of the Joint-Action Science and Technology project ([www.jast-project.eu](http://www.jast-project.eu)), this corpus was used to explore factors that might benefit human–robot interactions by studying human–human interactions in collaborative practical tasks. In this example, we devised 16 target tangrams, none of which resembled any obvious nameable entity. To engineer referring expressions, we designed each part to represent a unique color–shape combination, with each color represented only once, and each shape at most twice. The same initial set of seven pieces was available at the beginning of each construction trial. All had to be used.

Because trials in our paradigm can take over 4 min to complete, drift correction is important. Midtrial interruptions for calibration are undesirable for collaborative problem solving. Instead, the software package offers a manual correction utility for use on the reconstructed videos, which enables optional offline adjustments to be made.

**Design.** In order to investigate the relative worth of speech and gaze feedback in joint action, communication modalities were varied factorially: Participants could speak to each other and see the other person’s eye position; participants could speak but could not see where the other person was looking; participants could not speak but could see where their collaborator was looking; participants could neither speak nor see the gaze location of their partner. The condition order was rotated between dyads, but each pair of participants built four models under each condition. Additionally, because leadership and dominance factors can modulate how people interact, half of the dyads were assigned specific roles: one person was the project manager, whereas the other was the project assistant. The other half were merely instructed to collaborate. All were asked to reproduce the model tangram as quickly and as accurately as possible while minimizing breakages. To determine the usefulness of a verbal channel of communication during early strategy development, half of the dyads encountered the speaking conditions in the first half of the experiment, and the other half in the second.

### Results

In the following sections, we will indicate how such a corpus may be exploited. In the first two cases, we will illustrate types of automatically produced data that could

be used to explore some current questions. In the third section, we will show how these automatically recorded events can be combined with human coding to provide multimodal analyses. We will cite other work that further exploits the rich data of similar corpora.

**Fine-grained action structures.** Often resources are provided for joint actors on the assumption that whatever is available to an actor is used. Sometimes (see Bard, Anderson, et al., 2007) this proves not to be the case. The data that were recorded for the present experiment permitted very fine-grained analyses of action/gaze sequences that would allow us to discover without further coding who consults what information at critical phases of their actions. Figure 2B shows Participant B's screen 20 sec into an experimental trial. In this figure, Participant A has grasped the triangle on the right, and B has grasped the triangle on the left. Two seconds after the point captured in Figure 2B, the game software records that the players successfully joined the two parts together. Because our software captures continuous eyetracks, and because it records the precise time at which parts are joined, we can use the construction and eyetrack records to discover that in the 10 sec preceding the join, Participant A's gaze, which is shown here as a small circular cursor, moved rapidly between the two moving triangles and briefly rested in the workspace position at which the pieces were ultimately joined. Participant B's gaze, which is not displayed on this snapshot of B's screen but would appear on the reconstructed video, also traveled between the two triangles, but with two separate excursions to fixate the target model at the top right corner of the screen. Thus, the players are consulting different aspects of the available information while achieving a common goal. This particular goal—a combination of two triangles to construct a larger shape—was successfully achieved. Because all partial constructions have unique identifiers that are composed of their ancestry in terms of smaller constructions or primary parts, it is possible to distinguish successful acts of construction, which are not followed by breakage and a restart, from those that are quickly broken. It would then be possible to discover whether, as in this example, successful constructions were characterized by the differentiation of visual labor in their final delicate phases, whereas unsuccessful constructions were not. It would also be possible to discover whether the distribution of critical phase visual labor in each dyad could predict overall performance (automatically recorded per trial durations, breakage counts, and accuracy scores). Since the composition history automatically divides the final half of each interval between construction points from two earlier periods, we might also look for differentiation or alignment of gaze during those earlier phases in which players must develop a strategy for achieving the next subgoal.

**Trialwise measures: Actions entraining gaze.** Our software can also summarize events during a given trial. We illustrate how such figures might be used to determine whether gaze is largely entrained by the ongoing physical actions. If so, there may be little opportunity for role or strategy to determine attention here. We have several measures to use.

In the following examples, figures for Participant A are followed by those for Participant B in parentheses. In the trial seen in Figure 2B, 71% (72%) of the time is spent looking at the model, and 27% (20%) at the construction area, with the rest spent looking in other areas or off-screen. The player's position overlaps with a tangram part (moving or stationary) in the construction area 24% (11%) of the time, divided over 213 (173) different occasions, of which 56 (41) involve stable fixations. There are 29 occasions (of the 213 [173] gaze-part overlaps) during which both participants' eyetracks overlap with the same object, totaling 21.5 sec across the entire trial, but only 14 occasions during which 1 player's gaze overlaps with an object that the other player is currently manipulating. Thus, we have *prima facie* evidence that, for this dialogue, the common task on yoked screens does not draw the majority of players' visual activity to the same objects at the same time. Nor does an object that 1 player moves automatically entrain the other player's gaze. Making individual tests on pilot trials for a new type of joint task would allow the experimenter to determine whether the task had the necessary gaze-drawing properties for the experimental purpose.

In capturing this information about dyadic interaction, our experimental setup reveals a level of detail that is unprecedented in previous studies. The data from our eyetracker software suffice for studies of how the experimental variables interact and influence measures of task success, such as speed and accuracy. They also enable the analyst, for instance, to calculate the lag between when one person looks at a part and when the partner does, or to determine who initiates actions, as would be required for measuring dominance.

**Hand coding.** Although the system automates some tasks, it is clear that it cannot meet every experimental purpose. Export to NITE or ELAN allows further coding to suit the main experimental goals. As a simple demonstration, we follow the data of this study through transcription and reference coding. We use ChannelTrans (International Computer Science Institute, n.d.) to create orthographic transcription in which the conversational turns are time stamped against the rest of the data. By importing the transcription into the NITE XML Toolkit, we can identify and code the referring expressions used during the task. Figure A1 in the Appendix illustrates the use of an NXT coding tool to code each referring expression with the system's identifier for its referent. The result is a time-aligned multimodal account of gaze, reference, action, and trial conditions.

This allows us, for instance, to count the number of linguistic expressions referring to each part and to relate these references to the gaze and mouse data. As an example of the sorts of measures that this system makes easy to calculate, over the eight speech condition trials, for one of the participant dyads, there were 267 instances of speech and 1,262 instances of "looking" at DROIs, which lasted over a 45-msec duration. On 78 occasions, a player looked at a part while uttering a speech segment containing a reference to it. On 95 occasions, one player looked at a part while the other was uttering a speech segment referring to it. Figure 3 shows the complete data flow used to obtain this analysis.

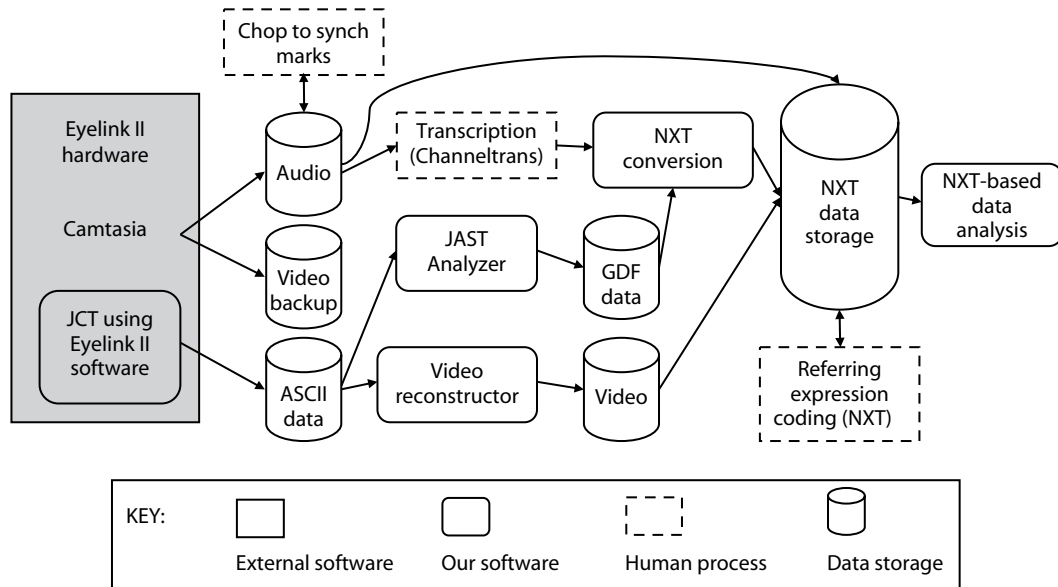


Figure 3. Data flow used to obtain the analysis for the example experiment.

Using data like these, we have shown how mouse movements coincide with particular distributions of referring expressions (Foster et al., 2008) and how the roles assigned to the dyad influence audience design in mouse gestures (Bard, Hill, & Foster, 2008).

**Empirical investigation: Determining whether dialogue or gaze projection during joint action leads to greater visual alignment.** If dialogue encourages alignment all the way from linguistic form to intellectual content, as Pickering and Garrod (2004) proposed, and if there is a shared visual environment in which to establish common ground (Clark, 1996; Clark & Brennan, 1991; Clark, Schreuder, & Buttrick, 1983; Lockridge & Brennan, 2002), then the opportunity for 2 people to speak should assure that they look at their common visual environment in a more coordinated way. Similarly, having a directly projected visual cue that indicates where another person is looking (perhaps analogous to using a laser pointer) offers an obvious focus for visual alignment. Thus, we would expect conditions allowing players to speak as they construct tangrams to have more aligned attention than those in which they play silently. And we would expect conditions in which each one's gaze is cross-projected on the other's screen to have more aligned attention than those with no gaze cursors. Using the aforementioned study, we can test these hypotheses via analyses of viewing behavior, visual alignment, and joint action, which were not previously possible.

First, we will examine a new dependent variable that can be automatically extracted from the data: the time spent looking at the tangram parts and the partially built tangrams. These are particularly complex DROIs, because either player is free to move, rotate, or even destroy any of the parts at any time. In this paradigm, however, the parts are the tools, and the tangrams are the goals of the game. They must be handled and guided with care. In this ex-

ample, we began by examining the time that participants spent looking at any of these ROIs, dynamic or otherwise, as a percentage of the total time spent performing the task using a  $2$  (speech, no speech)  $\times$   $2$  (gaze visible, invisible)  $\times$   $2$  (dyad member: A, B) ANOVA with dyads as cases. Since we know that, in monologue conditions, what people hear someone say will influence how they view a static display (Altmann & Kamide, 1999, 2004, 2007), we would expect the speaking conditions to direct players' gazes to the parts under discussion quite efficiently. In fact, the proportion of time spent inspecting any of the available ROIs was significantly lower when participants could speak to each other (17.41%) than when they were unable to speak (21.78%) [ $F(1,31) = 53.49, p < .001$ ]. We would also expect the direction of one player's gaze to help focus the other's attention on objects in play. We know that gaze projection provides such help for static ROIs (Stein & Brennan, 2004). In fact, there was no discernable difference in players' DROI viewing times as a consequence of being able to see their partners' eye positions [ $F(1,31) < 1$ ]: With and without cross-projected gaze, players tracked DROIs about 20% of the time overall. So players were not looking at the working parts as much when they were speaking, and seeing which piece their partners were looking at did not alter the proportion of time spent looking at task-critical objects.

Second, we exploited a measure that is contingent on being able to track both sets of eye movements against potentially moving or shifted targets: how often both players are looking at the same thing at exactly the same time. Here, a  $2$  (speech)  $\times$   $2$  (gaze) ANOVA indicates that the ability to speak in fact reduced the number of instances of aligned gaze by nearly 24% [29.7 with speech vs. 39.1 without speech;  $F(1,31) = 15.76, p < .001$ ]. Reduction in the frequency of visual alignment is not an artifact of trial duration. In fact, trials involving speech were an aver-



age of 7 sec longer than those without speech [95.99 sec vs. 89.02 sec, respectively;  $F(1,31) = 5.11, p < .05$ ] and might be expected to increase the number of eye–eye overlaps, even if just by chance. Again, however, the ability to see precisely where the other person was looking had no influence [ $F(1,31) < 1$ ], and there was no interaction.

The analyses that reveal overlapping gaze can also generate the latency between the point at which one player begins to look at an ROI and the point at which the other's gaze arrives: the eye–eye lag (Hadelich & Crocker, 2006). Eye–eye lag was shorter when participants could speak to each other [197 msec with speech as compared with 230 msec without speech;  $F(1,31) = 6.44, p < .05$ ]. Perhaps surprisingly, the projection of a collaborator's gaze position onto the screen failed to make any difference to these lag times [212 msec when gaze position was visible vs. 214 msec when it was not;  $F(1,31) < 1$ ], and again, there was no interaction between the variables.

Taken together, these results indicate that in a joint collaborative construction task, the facility for members of a working dyad to speak to each other reduces the proportion of time spent looking at the construction pieces and reduces the number of times both partners look at the same thing at the same time. Visual alignment appears to be inhibited when a dyad could engage in speech, but when it did occur, there was a shorter delay in coordination since the lag between one person looking at a potential target and the partner then moving his or her eyes onto the same target was smaller. In contrast, a direct indicator of a collaborator's gaze position did not appear to have any effect (facilitatory or inhibitory) on the measures examined. Contrary to expectations, therefore, being able to discuss a yoked visual workspace does not automatically yoke gaze. And there is no evidence that one person will track what another person is looking at more often when his or her eye position is explicitly highlighted. The combined influence and shape of any interaction between available modalities is obviously of critical importance to the understanding and modeling of multimodal communication and, as our present example demonstrates, the paradigm expounded here offers an ideal method of enhancing research into this topic.

Finally, gaze coordination can be further examined using cross-recurrence analysis. This technique has been used in the investigation of language and visual attention by Richardson and colleagues (Richardson & Dale, 2005; Richardson et al., 2007; Richardson et al., 2009) to demonstrate that visual coordination can manifest as being in phase or non-random without necessarily being simultaneous. However, these studies used only shared static images or replayed a monologue to participants. Bard and colleagues (Bard, Hill, & Arai, 2009; Bard et al., 2008; Bard, Hill, Nicol, & Carletta, 2007) have taken this technique further to examine the time course of visual alignment during the JCT.

## DISCUSSION

Our software implements a new method for collecting eyetracking data from a pair of interacting participants, in-

cluding showing the same game on the screens of the eyetrackers with indicators of the partner's gaze and mouse positions. The method allows us to construct a record of participant interaction that gives us the fine-scale timing of gaze, mouse, and task behaviors relative to each other and to moving objects on the screen.

We have demonstrated our techniques using a pair of Eyelink II eyetrackers and experimental software that implements a paradigm in which 2 participants jointly construct a model. Although our experimental paradigm is designed for studying a particular kind of joint action, the basic methods demonstrated in the software can be used to advance work in other areas. The software shows how to use one task to drive two eyetracker screens, record gaze against moving screen objects, show a partner's gaze and mouse icons, and synchronize the eyetracking, speech, and game records of 2 participants. Inventive message passing is the key, since messages can be both passed between the eyetrackers and stored in the output record, providing all of the data communication required. Tracking the relationship between gaze and moving objects, for instance, is simply a case of having the experimental software note the movement as messages in the eyetracker output, so that it can be checked against the gaze track analytically later on. Audio and screen capture can be synchronized to the eyetracker output by having the experimental software provide audible and visual synchronization marks, noting the time of these marks relative to the eyetracker output, again, as messages. Coordinating two screens requires each copy of the experimental software to record its state in outgoing messages and to respond to any changes reported to it in incoming ones. In theory, our techniques will work with any pair of eyetrackers that can handle message passing, including pairs of eyetrackers from different manufacturers. The message passing itself could become a bottleneck in some experimental setups, especially if it was used to study groups, but we found it adequate for our purpose. Manufacturers could best support the methodological advances we describe by ensuring that they include message-passing functionality that is both efficient and well documented.

These advances are useful not just for our work, but also for work in other areas in which more limited eyetracking methods are currently in use. Work on visual attention could benefit from the ability to register dynamic screen regions. Multiple object tracking (Bard et al., 2009; Bard, Hill, et al., 2007; Pylyshyn, 2006; Pylyshyn & Annan, 2006; Wolfe, Place, & Horowitz, 2007), subitization (Alston & Humphreys, 2004), and feature binding (Brockmole & Franconeri, 2009; Luck & Beach, 1998), for instance, would be less constrained if the objects were treated in the way we suggest.

It is joint action that has the most to gain, however, because it requires all of the advances made. The ability to import the data that comes out of the experimental paradigm into tools used for multimodal research, such as ELAN and NXT, offers particular promise for the study of joint action, especially where language is involved. These tools will enable the basic data to be browsed, combined

with orthographic transcription, enriched with hand annotation that interprets the participant behaviors, and searched effectively. Using these techniques will allow a fuller analysis of what the participants are doing, saying, and looking at than can any of the current techniques. The effect of copresence on joint action (Fussell & Kraut, 2004; Fussell et al., 2004; Horton & Keysar, 1996; Kraut, Fussell, & Siegel, 2003; Kraut et al., 2002) can also easily be manipulated by altering the proximity of the two eyetrackers—the limit only depending on the network connections. Similarly, the effect of perspective on language use (Hanna & Tanenhaus, 2004), including the use of deictic expressions such as “this” or “that,” can be investigated, as well as the differences in eye movements during comprehension (Spivey, Tanenhaus, Eberhard, & Sedivy, 2002) versus production (Horton & Keysar, 1996), and whether speakers engage in “audience design” or instead operate along more egocentric lines (Bard, Anderson, et al., 2007; Bell, 1984). Functional roles, such as instruction giver or follower, can be assigned to members of a dyad to determine how this might influence gaze behavior and the formation of referring expressions (Engelhardt, Bailey, & Ferreira, 2006).

The experimental environment opens up many possibilities for eyetracking dynamic images and active scene viewing, topics that are still surprisingly underresearched. In particular, the smooth pursuit of objects (Barnes, 2008; Burke & Barnes, 2006) that is controlled either by the person him- or herself or by his or her partner in a purposeful, nonrandom fashion can be investigated. Data can be output in their raw, sample-by-sample form, permitting customized smooth pursuit detection algorithms to be implemented as total gaze durations (accumulated time that the eye position coincided with an on-screen object, irrespective of eye stability), or as a series of fixations automatically identified by the SR Research (n.d.) software. Since most paradigms involve only static images, eye movements are almost exclusively reported in terms of saccades or fixations, ignoring the classification of pursuit movements. The EyeLink II also offers either monocular or binocular output, both of which can be utilized by our software.

The eyes gather information that is required for motor actions and are therefore proactively engaged in the process of anticipating actions and predicting behavior (Land & Furneaux, 1997). These skills are developed surprisingly early, with goal-directed eye movements being interpretable by the age of 1 (Falck-Ytter, Gredeback, & von Hofsten, 2006). More recently, Gesierich, Bruzzo, Ottoboni, and Finos (2008) studied gaze behavior and the use of anticipatory eye movements during a computerized block-stacking task. The phenomenon of the eyes preempting upcoming words in the spoken language stream is, as we have noted, the basis of the well-established “visual world” paradigm in psycholinguistics (Altmann & Kamide, 1999, 2004, 2007).

In summary, we have successfully developed a combined, versatile hardware and software architecture that enables groundbreaking and in-depth analysis of spontaneous, multimodal communication during joint action. Its flexibility and naturalistic game-based format help nar-

row the distinction between field and laboratory research while combining vision, language, and action. The system allows for more or less visual information recording (monocular, binocular data, or none at all) as well as for varying symmetrical and asymmetrical cross-projections of action or attention. Although optimized for joint action, it can operate in a stand-alone solo mode. This system therefore offers a new means of investigating a broad range of topics, including vision and active viewing, language and dialogue, joint action, copresence, strategy development and problem solving, and hand–eye coordination.

#### AUTHOR NOTE

J.P.d.R. is now at the Department of Linguistics, University of Bielefeld, Germany. R.L.H. is now at the Department of Psychology, School of Philosophy, Psychology, and Language Sciences, University of Edinburgh. The present work was supported by the EU FP6 IST Cognitive Systems Integrated Project Joint Action Science and Technology (JAST) FP6-003747-IP. The authors are grateful to Joseph Eddy and Jonathan Kilgour for assistance with programming. Correspondence concerning this article should be addressed to J. Carletta, Human Communication Research Centre, Informatics Forum, University of Edinburgh, 10 Crichton Street, Edinburgh EH8 9AB, Scotland (e-mail: j.carletta@ed.ac.uk).

#### REFERENCES

- ALSTON, L., & HUMPHREYS, G. W. (2004). Subitization and attentional engagement by transient stimuli. *Spatial Vision*, *17*, 17-50. doi:10.1163/15685680432277825
- ALTMANN, G. T. M., & KAMIDE, Y. (1999). Incremental interpretation at verbs: Restricting the domain of subsequent reference. *Cognition*, *73*, 247-264. doi:10.1016/S0010-0277(99)00059-1
- ALTMANN, G. T. M., & KAMIDE, Y. (2004). Now you see it, now you don't: Mediating the mapping between language and the visual world. In J. Henderson & F. Ferreira (Eds.), *The interface of language, vision and action: Eye movements and the visual world* (pp. 347-386). New York: Psychology Press.
- ALTMANN, G. T. M., & KAMIDE, Y. (2007). The real-time mediation of visual attention by language and world knowledge: Linking anticipatory (and other) eye movements to linguistic processing. *Journal of Memory & Language*, *57*, 502-518. doi:10.1016/j.jml.2006.12.004
- ANONYMOUS (n.d.). FFMPEG Multimedia System. Retrieved from <http://ffmpeg.mplayerhq.hu/>; last accessed February 2, 2010.
- APACHE XML PROJECT (1999). Xerces C++ parser. Retrieved February 2, 2010, from <http://xerces.apache.org/xerces-c/>.
- BANGERTER, A. (2004). Using pointing and describing to achieve joint focus of attention in dialogue. *Psychological Science*, *15*, 415-419. doi:10.1111/j.0956-7976.2004.00694.x
- BARD, E. G., ANDERSON, A. H., CHEN, Y., NICHOLSON, H. B. M., HARVARD, C., & DALZEL-JOB, S. (2007). Let's you do that: Sharing the cognitive burdens of dialogue. *Journal of Memory & Language*, *57*, 616-641. doi:10.1016/j.jml.2006.12.003
- BARD, E. G., HILL, R., & ARAI, M. (2009, July). *Referring and gaze alignment: Accessibility is alive and well in situated dialogue*. Paper presented at the Annual Meeting of the Cognitive Science Society, Amsterdam.
- BARD, E. G., HILL, R., & FOSTER, M. E. (2008, July). *What tunes accessibility of referring expressions in task-related dialogue?* Paper presented at the Annual Meeting of the Cognitive Science Society, Washington, DC.
- BARD, E. G., HILL, R., NICOL, C., & CARLETTA, J. (2007, August). *Look here: Does dialogue align gaze in dynamic joint action?* Paper presented at AMLaP2007, Turku, Finland.
- BARNES, G. (2008). Cognitive processes involved in smooth pursuit eye movements. *Brain & Cognition*, *68*, 309-326. doi:10.1016/j.bandc.2008.08.020
- BELL, A. (1984). Language style as audience design. *Language in Society*, *13*, 145-204.

- BRENNAN, S. E., CHEN, X., DICKINSON, C., NEIDER, M., & ZELINSKY, G. (2008). Coordinating cognition: The costs and benefits of shared gaze during collaborative search. *Cognition*, **106**, 1465-1477. doi:10.1016/j.cognition.2007.05.012
- BROCKMOLE, J. R., & FRANCONERI, S. L. (Eds.) (2009). Binding [Special issue]. *Visual Cognition*, **17**(1 & 2).
- BURKE, M., & BARNES, G. (2006). Quantitative differences in smooth pursuit and saccadic eye movements. *Experimental Brain Research*, **175**, 596-608. doi:10.1007/s00221-006-0576-6
- CAMPBELL, C., JR., & McROBERTS, T. (2005). The simple sockets library. Retrieved February 2, 2010, from <http://mysite.verizon.net/astronaut/ssl/>.
- CHARNESS, N., REINGOLD, E. M., POMPLUN, M., & STAMPE, D. M. (2001). The perceptual aspect of skilled performance in chess: Evidence from eye movements. *Memory & Cognition*, **29**, 1146-1152.
- CHERUBINI, M., NÜSSLI, M.-A., & DILLENBOURG, P. (2008). Deixis and gaze in collaborative work at a distance (over a shared map): A computational model to detect misunderstandings. In K.-J. Riih a & A. T. Duchowski (Eds.), *ETRA 2008—Proceedings of the Eye Tracking Research and Application Symposium* (pp. 173-180). New York: ACM Press. doi:10.1145/1344471.1344515
- CLARK, H. H. (1996). *Using language*. Cambridge: Cambridge University Press.
- CLARK, H. H. (2003). Pointing and placing. In S. Kita (Ed.), *Pointing: Where language, culture, and cognition meet* (pp. 243-268). Mahwah, NJ: Erlbaum.
- CLARK, H. H., & BRENNAN, S. E. (1991). Grounding in communication. In L. B. Resnick, J. Levine, & S. D. Teasley (Eds.), *Perspectives on socially shared cognition* (pp. 127-149). Washington, DC: American Psychological Association.
- CLARK, H. H., & KRYCH, M. A. (2004). Speaking while monitoring addressees for understanding. *Journal of Memory & Language*, **50**, 62-81. doi:10.1016/j.jml.2003.08.004
- CLARK, H. H., SCHREUDER, R., & BUTTRICK, S. (1983). Common ground and the understanding of demonstrative reference. *Journal of Verbal Learning & Verbal Behavior*, **22**, 245-258.
- DUCHOWSKI, A. T. (2003). *Eye tracking methodology: Theory and practice*. London: Springer.
- ENGELHARDT, P. E., BAILEY, K. G. D., & FERREIRA, F. (2006). Do speakers and listeners observe the Gricean Maxim of Quantity? *Journal of Memory & Language*, **54**, 554-573. doi:10.1016/j.jml.2005.12.009
- FALCK-YTTER, T., GREDEBACK, G., & VON HOFSTEN, C. (2006). Infants predict other people's action goals. *Nature Neuroscience*, **9**, 878-879. doi:10.1038/nn1729
- FOSTER, M. E., BARD, E. G., GUHE, M., HILL, R. L., OBERLANDER, J., & KNOLL, A. (2008). The roles of haptic-ostensive referring expressions in cooperative task-based human-robot dialogue. In T. Fong, K. Dautenhahn, M. Scheutz, & Y. Demiris (Eds.), *Proceedings of the 3rd ACM/IEEE International Conference on Human-Robot Interaction* (pp. 295-302). New York: ACM Press. doi:10.1145/1349822.1349861
- FUSSELL, S., & KRAUT, R. (2004). Visual copresence and conversational coordination. *Behavioral & Brain Sciences*, **27**, 196-197. doi:10.1017/S0140525X04290057
- FUSSELL, S., SETLOCK, L., YANG, J., OU, J. Z., MAUER, E., & KRAMER, A. (2004). Gestures over video streams to support remote collaboration on physical tasks. *Human-Computer Interaction*, **19**, 273-309. doi:10.1207/s15327051hci1903\_3
- GESIERICH, B., BRUZZO, A., OTTOBONI, G., & FINOS, L. (2008). Human gaze behaviour during action execution and observation. *Acta Psychologica*, **128**, 324-330. doi:10.1016/j.actpsy.2008.03.006
- GIGERENZER, G., TODD, P. M., & ABC RESEARCH GROUP (1999). *Simple heuristics that make us smart*. Oxford: Oxford University Press.
- GRANT, E. R., & SPIVEY, M. J. (2003). Eye movements and problem solving: Guiding attention guides thought. *Psychological Science*, **14**, 462-466. doi:10.1111/1467-9280.02454
- GRIFFIN, Z. M. (2004). Why look? Reasons for eye movements related to language production. In J. M. Henderson & F. Ferreira (Eds.), *The integration of language, vision, and action: Eye movements and the visual world*. New York: Psychology Press.
- GRIFFIN, Z. M., & OPPENHEIMER, D. M. (2006). Speakers gaze at objects while preparing intentionally inaccurate labels for them. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **32**, 943-948. doi:10.1037/0278-7393.32.4.943
- HADDELICH, K., & CROCKER, M. W. (2006, March). *Gaze alignment of interlocutors in conversational dialogues*. Paper presented at the 19th Annual CUNY Conference on Human Sentence Processing, New York.
- HANNA, J., & TANENHAUS, M. K. (2004). Pragmatic effects on reference resolution in a collaborative task: Evidence from eye movements. *Cognitive Science*, **28**, 105-115. doi:10.1207/s15516709cog2801\_5
- HENDERSON, J. M., & FERREIRA, F. (Eds.) (2004). *The interface of language, vision, and action: Eye movements and the visual worlds*. New York: Psychology Press.
- HORTON, W. S., & KEYSAR, B. (1996). When do speakers take into account common ground? *Cognition*, **59**, 91-117. doi:10.1016/0010-0277(96)81418-1
- INTERNATIONAL COMPUTER SCIENCE INSTITUTE (n.d.). Extensions to Transcriber for Meeting Recorder Transcription. Retrieved from [www.icsi.berkeley.edu/Speech/mr/channeltrans.html](http://www.icsi.berkeley.edu/Speech/mr/channeltrans.html); last accessed February 2, 2010.
- KRAUT, R., FUSSELL, S., & SIEGEL, J. (2003). Visual information as a conversational resource in collaborative physical tasks. *Human-Computer Interaction*, **18**, 13-49. doi:10.1207/S15327051HCI1812\_2
- KRAUT, R., GERGLE, D., & FUSSELL, S. (2002, November). *The use of visual information in shared visual spaces: Informing the development of virtual co-presence*. Paper presented at the ACM Conference on Computer Supported Cooperative Work, New Orleans, LA.
- LAND, M. F. (2006). Eye movements and the control of actions in everyday life. *Progress in Retinal & Eye Research*, **25**, 296-324. doi:10.1016/j.preteyeres.2006.01.002
- LAND, M. F. (2007). Fixation strategies during active behavior: A brief history. In R. P. G. van Gompel, M. H. Fischer, W. S. Murray, & R. L. Hill (Eds.), *Eye movements: A window on mind and brain* (pp. 75-95). Oxford: Elsevier.
- LANGUAGE TECHNOLOGY GROUP (n.d.). NITE XML Toolkit Homepages. Retrieved from <http://groups.inf.ed.ac.uk/nxt/>; last accessed February 2, 2010.
- LAND, M. F., & FURNEAUX, S. (1997). The knowledge base of the oculomotor system. *Philosophical Transactions of the Royal Society B*, **352**, 1231-1239.
- LINDSTR M, A. (1999). SGE: SDL Graphics Extension. Retrieved February 2, 2010, from [www.digitalfanatics.org/cal/sge/index.html](http://www.digitalfanatics.org/cal/sge/index.html).
- LOBMAIER, J. S., FISCHER, M. H., & SCHWANINGER, A. (2006). Objects capture perceived gaze direction. *Experimental Psychology*, **53**, 117-122. doi:10.1027/1618-3169.53.2.117
- LOCKRIDGE, C. B., & BRENNAN, S. E. (2002). Addressees' needs influence speakers' early syntactic choices. *Psychonomic Bulletin & Review*, **9**, 550-557.
- LUCK, S. J., & BEACH, N. J. (1998). Visual attention and the binding problem: A neurophysiological perspective. In R. D. Wright (Ed.), *Visual attention* (pp. 455-478). Oxford: Oxford University Press.
- MAX PLANCK INSTITUTE FOR PSYCHOLINGUISTICS (n.d.). Language Archiving Technology: ELAN. Retrieved from [www.lat-mpi.eu/tools/elan/](http://www.lat-mpi.eu/tools/elan/); last accessed February 2, 2010.
- MEYER, A. S., & DOBEL, C. (2003). Application of eye tracking in speech production research. In J. Hy n , R. Radach, & H. Deubel (Eds.), *The mind's eye: Cognitive and applied aspects of eye movement research* (pp. 253-272). Amsterdam: Elsevier.
- MEYER, A. S., VAN DER MEULEN, F., & BROOKS, A. (2004). Eye movements during speech planning: Speaking about present and remembered objects. *Visual Cognition*, **11**, 553-576. doi:10.1080/13506280344000248
- MONK, A., & GALE, C. (2002). A look is worth a thousand words: Full gaze awareness in video-mediated conversation. *Discourse Processes*, **33**, 257-278. doi:10.1207/S15326950DP3303\_4
- MURRAY, N., & ROBERTS, D. (2006, October). *Comparison of head gaze and head and eye gaze within an immersive environment*. Paper presented at the 10th IEEE International Symposium on Distributed Simulation and Real-Time Applications, Los Alamitos, CA. doi:10.1109/DS-RT.2006.13
- PICKERING, M., & GARROD, S. (2004). Towards a mechanistic psychology of dialogue. *Behavioral & Brain Sciences*, **27**, 169-190. doi:10.1017/S0140525X04000056
- PYLYSHYN, Z. W. (2006). Some puzzling findings in multiple object

- tracking (MOT): II. Inhibition of moving nontargets. *Visual Cognition*, **14**, 175-198. doi:10.1080/13506280544000200
- PYLYSHYN, Z. W., & ANNAN, V. (2006). Dynamics of target selection in multiple object tracking (MOT). *Spatial Vision*, **19**, 485-504. doi:10.1163/156856806779194017
- RAYNER, K. (1998). Eye movements in reading and information processing: 20 years of research. *Psychological Bulletin*, **124**, 372-422.
- RICHARDSON, D., & DALE, R. (2005). Looking to understand: The coupling between speakers' and listeners' eye movements and its relationship to discourse comprehension. *Cognitive Science*, **29**, 1045-1060. doi:10.1207/s15516709cog0000\_29
- RICHARDSON, D., DALE, R., & KIRKHAM, N. (2007). The art of conversation is coordination: Common ground and the coupling of eye movements during dialogue. *Psychological Science*, **18**, 407-413. doi:10.1111/j.1467-9280.2007.01914.x
- RICHARDSON, D., DALE, R., & TOMLINSON, J. (2009). Conversation, gaze coordination, and beliefs about visual context. *Cognitive Science*, **33**, 1468-1482. doi:10.1111/j.1551-6709.2009.01057.x
- SIMPLE DIRECTMEDIA LAYER PROJECT (n.d.). SDL: Simple DirectMedia Layer. Retrieved February 2, 2010, from www.libsdl.org/.
- SPIVEY, M., & GENG, J. (2001). Oculomotor mechanisms activated by imagery and memory: Eye movements to absent objects. *Psychological Research*, **65**, 235-241. doi:10.1007/s004260100059
- SPIVEY, M., TANENHAUS, M. K., EBERHARD, K. M., & SEDIVY, J. C. (2002). Eye movements and spoken language comprehension: Effects of visual context on syntactic ambiguity resolution. *Cognitive Psychology*, **45**, 447-481. doi:10.1016/S0010-0285(02)00503-0
- SR RESEARCH (n.d.). *Complete eyetracking solutions*. Retrieved February 2, 2010, from www.sr-research.com/EL\_II.html.
- STEIN, R., & BRENNAN, S. E. (2004, October). *Another person's eye gaze as a cue in solving programming problems*. Paper presented at ICMI '04: 6th International Conference on Multimodal Interfaces, State College, PA.
- STEPTOE, W., WOLFF, R., MURGIA, A., GUIMARAES, E., RAE, J., SHARKEY, P., ET AL. (2008, November). *Eye-tracking for avatar eye-gaze and interactional analysis in immersive collaborative virtual environments*. Paper presented at the ACM Conference on Computer Supported Cooperative Work, San Diego, CA.
- TECHSMITH (n.d.). Camtasia Studio. Retrieved February 2, 2010, from www.techsmith.com/camtasia.asp.
- TRUESWELL, J. C., & TANENHAUS, M. K. (Eds.) (2005). *Approaches to studying world-situated language use: Bridging the language-as-product and language-as-action traditions*. Cambridge, MA: MIT Press.
- UNDERWOOD, G. (Ed.) (2005). *Cognitive processes in eye guidance*. Oxford: Oxford University Press.
- UNDERWOOD, J. (2005). Novice and expert performance with a dynamic control task: Scanpaths during a computer game. In G. Underwood (Ed.), *Cognitive processes in eye guidance* (pp. 303-323). Oxford: Oxford University Press.
- VAN GOMPEL, R. P. G., FISCHER, M. H., MURRAY, W. S., & HILL, R. L. (Eds.) (2007). *Eye movements: A window on mind and brain*. Oxford: Elsevier.
- VELICHKOVSKY, B. M. (1995). Communicating attention: Gaze position transfer in cooperative problem solving. *Pragmatics & Cognition*, **3**, 199-222.
- VERTEGAAL, R., & DING, Y. (2002, November). *Explaining effects of eye gaze on mediated group conversations: Amount or synchronization?* Paper presented at the ACM Conference on Computer Supported Cooperation Work, New Orleans, LA.
- WOLFE, J. M., PLACE, S. S., & HOROWITZ, T. S. (2007). Multiple object juggling: Changing what is tracked during extended multiple object tracking. *Psychonomic Bulletin & Review*, **14**, 344-349.
- WORLD WIDE WEB CONSORTIUM (1999). XSL Transformations (XSLT) Version 1.0: W3C Recommendation, 16 November. Retrieved from www.w3.org/TR/xslt; last accessed February 2, 2010.

## APPENDIX

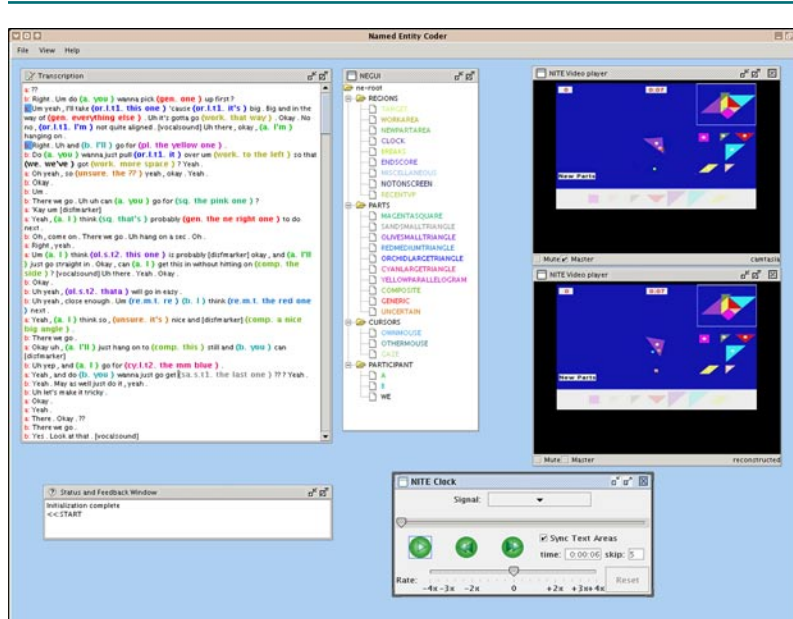


Figure A1. Screen shot of NXT's discourse entity coding tool in use to code referring expressions for the example experiment.