

Article

DOI: 10.1111/j.1468-0394.2011.00584.x

Supporting activity modelling from activity traces

Olivier L. Georgeon,¹ Alain Mille,¹ Thierry Bellet,²
Benoit Mathern¹ and Frank E. Ritter³

(1) *Université de Lyon, 86 Rue Pasteur 69007 Lyon, France*

Email: olivier.georgeon@liris.cnrs.fr

(2) *Institut Français des Sciences et Technologies des Transports, de l'Aménagement et des Réseaux, 25, Avenue François Mitterrand, 69500 Bron, France*

(3) *The Pennsylvania State University, University Park, PA 16802, USA*

Abstract: *We present a new method and tool for activity modelling through qualitative sequential data analysis. In particular, we address the question of constructing a symbolic abstract representation of an activity from an activity trace. We use knowledge engineering techniques to help the analyst build an ontology of the activity, that is, a set of symbols and hierarchical semantics that supports the construction of activity models. The ontology construction is pragmatic, evolutionist and driven by the analyst in accordance with their modelling goals and their research questions. Our tool helps the analyst define transformation rules to process the raw trace into abstract traces based on the ontology. The analyst visualizes the abstract traces and iteratively tests the ontology, the transformation rules and the visualization format to confirm the models of activity. With this tool and this method, we found innovative ways to represent a car-driving activity at different levels of abstraction from activity traces collected from an instrumented vehicle. As examples, we report two new strategies of lane changing on motorways that we have found and modelled with this approach.*

Keywords: sequence mining, timeline analysis, activity trace, knowledge-based system, activity modelling

1. Introduction

We introduce here new principles based on knowledge engineering techniques for designing systems to help *analysts* create models of activity from *activity traces*. We illustrate these principles with a software tool that we have implemented, and with an example modelling analysis that we have performed using this tool.

By *activity trace* we mean a set of multiple streams of quantitative or symbolic data that record (at least partially) an activity performed by a *subject*. The *analysts* may be psychologists seeking to build theories of the subject's cognition, ergonomists seeking to design better user

interfaces, analysts seeking to predict the subject's behaviour in specific conditions, trainers seeking to improve training techniques or even the subjects themselves seeking to improve their understanding of their actions. In each case, the created models of activity constitute micro-theories proposed by the analysts to describe, explain and try to predict how the subject performs the activity.

The principles and the tool that we introduce here address three needs for helping analysts construct models of activity from activity traces. The first need is for helping the analyst learn previously unknown aspects or details of the subject's activity from the activity traces. The

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
(BWUK EXSY 584 Webpdf:=02/11/2011 04:46:26 51355 Bytes 15 PAGES n operator=) 2/11/2011 4:46:31 PM

1 second need is for helping the analyst construct
2 meaningful symbolic representations of
3 interesting aspects of the activity. These representations,
4 associated with the explanations proposed by the analyst,
5 constitute the models of activity. The third need is for
6 helping the analyst test and support the created models
7 of activity with regard to the activity trace.

8 Activity traces have also been called *protocols*
9 (Ericsson & Simon, 1993) or simply *sequential*
10 *data* (Sanderson & Fisher, 1994). We prefer the
11 term *activity trace* because this term conveys the
12 idea that the trace is intended to be interpreted by
13 somebody (designated here as the analyst). We think
14 of a trace as a footprint that helps who sees it
15 understand what happened. Our activity traces
16 yet differ from mere footprints in that they are not
17 accidentally produced or unprocessed but they
18 rather result from the analyst's choices and set-up.

19 Many software tools have been implemented
20 for activity trace analysis. A recent review (Hilbert
21 & Redmiles, 2000) notes 40 of them. These tools
22 cannot autonomously generate a comprehensive
23 explanation of human behaviour but they interactively
24 support analysis. This analysis consists of identifying,
25 categorising, labelling and transforming pieces of
26 data and information in the activity trace. We
27 summarize this process by the notion of *abstraction*.
28 The analysts use their expertise and knowledge to
29 formulate a whole set of tiny hypotheses and
30 choices concerning how to collect the data, how
31 to filter it, how to cluster and label it and how
32 to display and to report it so that it responds to
33 the analysis purpose.

34 Although most of the existing tools acknowledge
35 the central role of the analysts and the importance
36 of their knowledge and expertise in the analysing
37 process, these tools still lack knowledge representation
38 mechanisms to support the management of the
39 analysts' knowledge. For instance, MacShapa
40 (Sanderson *et al.*, 1994) does help analysts label
41 and cluster the behavioural data. It also acknowledges
42 the usage of these labels as symbols to describe the
43 activity. It, however, does not help the analyst
44 formulate and manage the symbolic inferences
45 she can make from these symbols.

We formulated the hypothesis that aspects of
knowledge engineering can help design software
systems that address the three needs identified
above. We use ontology management facilities
and rule engines to capture the hypotheses and
choices made by the analyst. When interactively
used by the analyst or a group of analysts, these
facilities help formalize the way analysts find
interesting symbolic patterns and infer models
from them. Once this knowledge is formalized,
the system uses it to automatically compute new
representations of the activity from the activity
traces. The system also helps analysts organize
and store the different concepts and rules that
summarize different studies and help capitalize
on these studies.

To explain our principles and demonstrate
the system and its design, we have organized
this paper as follows: Section 2 presents the
principles of activity trace modelling, based on
a pragmatic and evolutionist approach. Section
3 presents the prototype software tool that we
have implemented from these principles, its
technical features, its architecture and its user
interface. Section 4 presents an example study
in which we have used this tool to create models
of lane change on motorways from activity traces
generated with an instrumented car. The method
is then summarized in the conclusion.

2. Modelling activity traces

The notion of activity traces is widely used in
the human behaviour literature, and we cannot
attribute its origin to a specific author. Only
more specific-related notions can be identified,
such as pattern languages, as reviewed by
Dear den and Finlay (2006), or grammar
representations (Olson *et al.*, 1994). Despite
the wide usage of the term activity trace, we
could not find a definition of it, which led us
to propose the following definition:

An activity trace is a meaningful inscription,
from the viewpoint of an analyst, of the flow
of what has happened, from the viewpoint of
a subject.

With this definition, we want to highlight
that an activity trace always implies two
viewpoints,

1 situated in two different moments. It implies
2 the subject's viewpoint, when he or she was
3 performing the activity, and the analyst's view-
4 point, when he or she is analysing the activity
5 trace. Indeed, an activity trace cannot be an
6 inscription of all that happened (if that had any
7 sense), because an activity only concerns what
8 relates to the subject's perspective, goals and
9 intentions. Thus, inevitably, the analyst has to
10 make assumptions about what is meaningful to
11 the subject when the analyst sets up the tracing
12 mechanism. In addition, the activity trace de-
13 pends on what activity aspects interest the analyst
14 and what makes sense to her according to her
15 previous knowledge and to her analysis goals.

16 Because an activity trace depends on the ana-
17 lyst's knowledge and assumptions, it can only be
18 modelled in an iterative way, each iteration pro-
19 ducing new knowledge leading to new hypotheses
20 for the next iteration. Ericsson and Simon (1993)
21 described this iterative nature of analysing human
22 behaviour: 'In designing our data-gathering
23 schemes, we make minimal essential theoretical
24 commitments, then try to use the data to test
25 stronger theories' (p. 274). Moreover, none of the
26 iterations can produce knowledge that could be
27 proven to be true in an absolute sense, but only
28 knowledge that is more efficient and useful with
29 regard to the analyst's goals and that is more
30 convincing to the analyst's community than the
31 knowledge from the previous iteration. More
32 broadly, this conception of knowledge relates to
33 a pragmatic epistemology (James, 1907) and an
34 evolutionist epistemology (Popper, 1972).

35 These pragmatic and evolutionist aspects are
36 crucial when defining a methodology and a tool
37 for activity trace modelling. By fully acknowl-
38 edging these aspects, we have designed a tool that
39 facilitates and accelerates the evolutionist model-
40 ling process. The tool helps formulate a series of
41 micro-hypotheses of possibly useful symbols, pos-
42 sibly useful transformation rules to transform the
43 low-level data into higher-level data and possibly
44 useful representations of the activity trace based
45 on the micro-hypotheses. If the obtained repre-
46 sentation does not help the analyst understand the
47 activity better, then she rejects these micro-
48 hypotheses; if it helps, then she keeps them.

The tool is designed to shorten this formula-
tion/usage/validation-or-rejection loop. This
process leads to the construction of a set of
micro-hypotheses that are validated by the ana-
lyst. This set constitutes a formalization of the
analyst's knowledge about how to understand
the activity. The tool stores this knowledge,
helps the analyst keep track of it and helps the
analysts' community discuss and question it.

The next section details how the tool reaches
this goal by helping the analyst define sequences
at the right level of abstraction and simulta-
neously identify interesting subsequences, define
them precisely and query the whole trace in
search for their occurrences.

2.1. Collecting a symbolic trace

A raw activity trace can be made of any kind of
data describing a subject's activity flow and
intended for an analyst's usage. In a broad
sense, it can range from video or audio record-
ing to computer logs. The only common point is
that the data are temporally organized, meaning
that each data piece is associated with a time-
stamp referring to a common time base. The
first abstraction step consists of converting these
raw traces into sequences of symbols. We refer
to this step as the *discretization* of the raw trace
into a symbolic trace. The symbols in the sym-
bolic trace have to be meaningful to the analyst,
and they are chosen on a pragmatic and evolu-
tionist basis, in compliance with a pragmatic
and evolutionist epistemology introduced above
(introduction of Section 2). The discretization
process can be manual, semi-automatic or auto-
matic. The definition of the symbols may evolve
in parallel with the implementation of the dis-
cretization process. This is because the later
interpretation of the symbols may differ from
the meaning initially intended by the analyst
when she specifies the discretization algorithm.

For example, in a study of car driving, we
have used the classical mathematical curve ana-
lysis method to generate symbols of interest
from numerical values of the vehicle speed, the
steering wheel angle and the pedal positions. In
this case, the symbols correspond to threshold

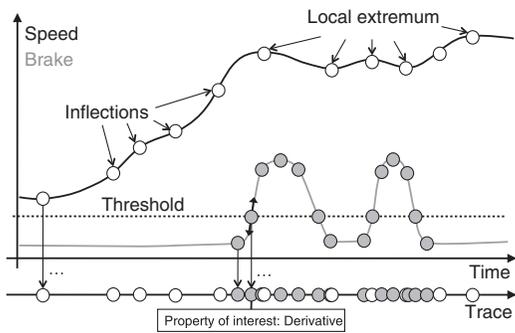


Figure 1: Discretization of analogical traces.

crossing, local extremum and inflexion points. Figure 1 illustrates this discretization process.

In Figure 1, the curves represent the vehicle speed in km/h and the brake pedal position in percentage of range. Symbols of interest are shown on these curves as circles (threshold crossings, inflexion points, local extremums). The symbols are merged into the symbolic trace that is represented at the bottom of Figure 1. The figure also shows a derivative value as an example property of interest associated with a symbol generated by a brake pedal threshold crossing. The analyst specifies the way to generate these symbols so that they correspond to meaningful events that describe the activity. In this example, the threshold crossing indicates the extent of the braking action and the derivative value indicates the abruptness of this action. Notably, while creating these symbols, the analyst claims the *existence* of the events that these symbols represent. By so doing, the analyst defines an ontology of the activity.

Our experience has taught us that the system must maintain a connection between the raw trace and the symbolic trace, and provide parallel displays of both of them. The analyst needs to tune many parameters of the discretization algorithms, like the threshold values or noise filters. The analyst validates the chosen symbols and algorithms by comparing the symbolic trace to the raw trace and ensuring that the symbolic trace represents what is happening. While she defines and validates these symbols, the analyst also supports her claim

that the events represented by these symbols 'exist'. This support arises because the method to generate these symbols from the recorded data is formally specified and explained by the analyst. In this example, after we fully specified the discretization algorithm in accordance to our specific modelling goals, the discretization algorithm could then compute the symbols fully automatically.

2.2. Modelling the symbolic trace

At the symbolic level, analysts most often want to focus on relations between events. Indeed, events are not meaningful by themselves, but they become meaningful in the context where they relate to each other (Sanderson & Fisher, 1994). Examples of such relations include a 'sequence following within a certain period of time', 'co-occurrence within a certain period of time' and 'causality with regard to a certain explanative theory'. Building and understanding these relations between events is a part of the analysing process. By definition, a set of elements connected through relations is a graph. Therefore, we model the symbolic traces with a graph structure.

More precisely, our trace graph structure has two parts: a sequence and an ontology. The sequence is a part of the graph that is made of event instances and of relation instances between event instances. The ontology is a part of the graph that is made of event classes and of relations between event classes. Figure 2 illustrates this graph structure with a simplified example taken from the car-driving study.

The sequence is represented in the bottom part of Figure 2 and the ontology in the top part. In the sequence graph, event instances are represented as circles and triangles. Relation instances between event instances are represented as solid arrows between these circles and triangles. Example properties of event instances are represented with grey dashed arrows pointing to their value at the bottom of the figure (e.g. the duration of an eye movement and acceleration value associated with an inflexion point of the speed curve).

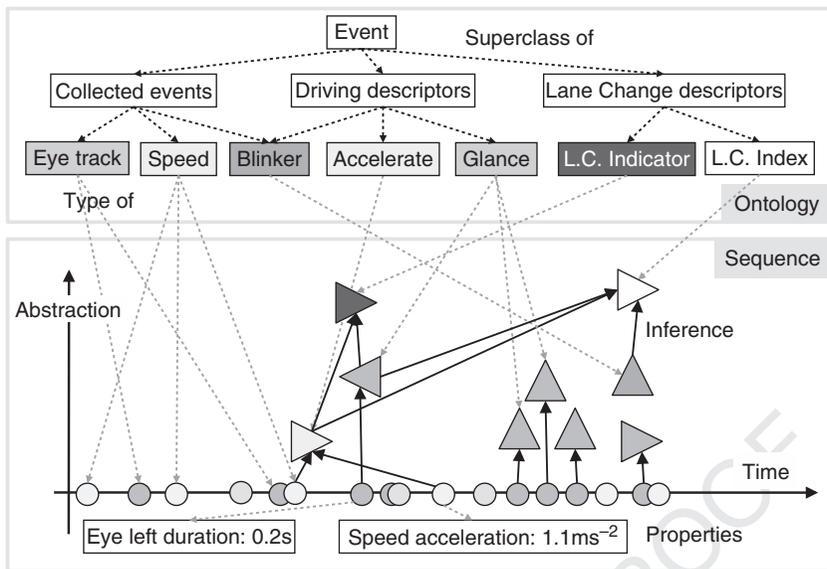


Figure 2: Activity trace modelled in a knowledge engineering system.

Our tool displays the sequence in a similar form as shown in Figure 2; the exact form is shown in Figure 4. This display uses two axes: the time axis and the 'abstraction' axis. That is, the events' time-code attributes determine their x coordinates and the analyst specifies their y coordinates when she configures these event's class in the ontology. The analyst can use the y coordinate to express different meanings; our recommendation is to use it to express an idea of abstraction level related to a specific analysis. In this example, the lowest level (circles) represents the events obtained from the discretization process; the intermediary level represents events that describe the activity in usual driving terms: accelerate, glance, turn signal on/off (blinker); and the higher level represents events describing lane-change behaviour: indicator of intention to change lane and index of lane change. This example expresses the analyst's assumption that the conjunction of an acceleration and a glance to the left rear mirror can generate an indicator of the driver's intention to change lane (L.C. Indicator).

In the ontology graph, the nodes represent event classes and the edges represent the relation 'sub-class of' (dashed black arrows). The analyst defines the ontology during the modelling pro-

cess. For example, the 'Collected events' class includes all the event classes that come from the discretization process. The 'Driving descriptors' class gathers the intermediary event classes that describe the activity in usual driving terms. The 'Lane change descriptors' class gathers the most abstract event classes describing lane changes. The dotted grey arrows in the figure represent the relation 'type of' going from event classes defined in the ontology to event instances in the sequence.

Again, like most ontologies, this ontology is made by the analyst on a pragmatic basis. It is likely that two analysts will create two different ontologies. While the software cannot demonstrate that one is better than the other, it does help the analysts formalize and discuss them. This discussion leads to the construction of a language for describing the activity that represents an agreement about the terms that can be used to describe the activity. In addition, the ontology also supports the analysts' agreement about how the trace should be visualized, because the visualization properties of the symbols are stored in the ontology.

This trace formalism enables the analyst to conduct a hierarchical analysis of the activity.

Event instances form a hierarchy where higher-level events are temporal patterns of lower-level events. The ontology also defines a hierarchy because some event classes are sub-classes of others. Notably, these two hierarchies are different because lower-level event instances do not necessarily belong to a sub-class of the class of higher-level event instances.

3. System implementation

We have implemented a prototype system based on an assemblage of open-source knowledge engineering tools: an ontology editor, an inference engine, visualization facilities and documentation facilities. This system is named ABSTRACT (Analysis of Behavior and Situation for menTal Representation Assessment and Cognitive acTivity modelling). Figure 3 illustrates this assemblage.

The system can be split into three levels: a lower level, at the bottom of the figure, which is the *collection system*; a core level, in the centre of

the figure, which is the *symbolic trace system* itself; and a higher level, on the top of the figure, which is a *documentation level*.

3.1. The collection system

The collection system integrates tools to help the analyst prepare the symbolic trace. We call these tools collection agents. Collection agents may be automatic when specified once by the analyst or may require the analyst's intervention. Automatic collection agents can be tools for preprocessing sensor data or computer logs, as in the example of Section 2.1. Semi-automatic collection agents can be tools for helping the analyst take notes, record interviews or transcript video data. As noted, this discretization cannot be done blindly, but must be driven by the analyst. Hence, this level requires visualization facilities. We use Microsoft Excel with specific Visual Basic Application (VBA) macros as a visualization tool for the collection system. In this visualization, each event of the symbolic trace

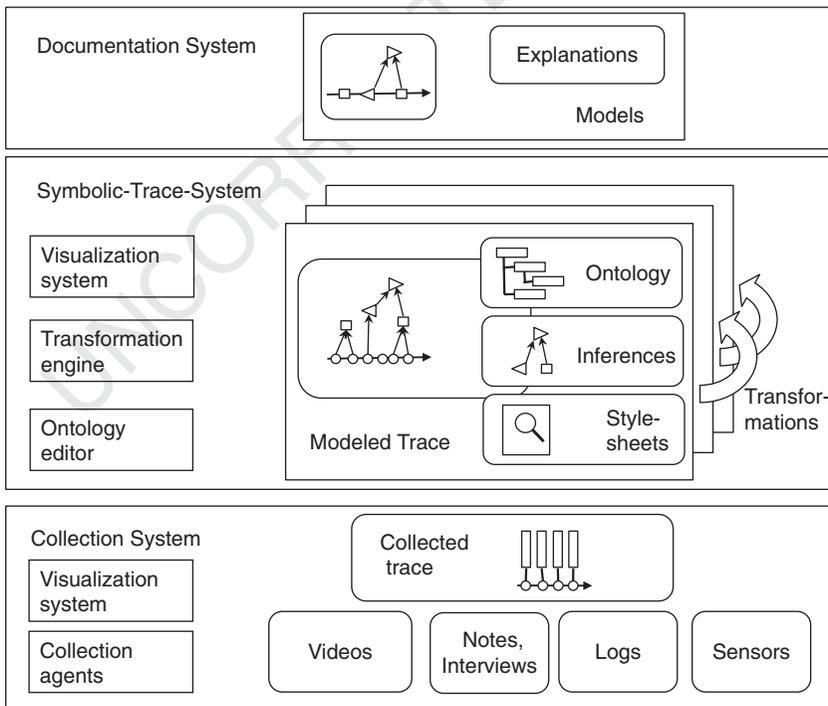


Figure 3: The architecture of the ABSTRACT activity analysis system.

1 is displayed as a line in the spreadsheet. The
2 lines are coloured according to the event's type
3 and the event's properties are organized in
4 different columns. We have implemented a spe-
5 cific video player and analogical data player that
6 triggers VBA macros that automatically scrolls
7 down the spreadsheet in synchronization. Some
8 of these facilities are also available in commer-
9 cial quantitative data analysis tools such as
10 MacShapa (Sanderson *et al.*, 1994) and NVivo.
11 These facilities allow the analyst to check and
12 validate or reject the symbolic trace, that is,
13 refine the discretization algorithm and its di-
14 verse parameters until she gets a satisfying and
15 meaningful symbolic trace including appropri-
16 ate properties of interest.

17 3.2. *The symbolic trace system*

18 The symbolic trace system is the knowledge
19 engineering system itself. At this level, the traces
20 are modelled as described in Section 2.2. In
21 addition, they are associated with a set of
22 inference rules and a set of *style sheets*. A style
23 sheet is a specification for displaying the trace
24 on the screen. It specifies how semantic prop-
25 erties of the events that are defined in the ontology
26 should be converted into visualization prop-
27 erties, such as shape and position. Style sheets also
28 implement particular time scales, and particular
29 filters to display only the interesting aspect of
30 the trace for a particular analysis. They corre-
31 spond to different ways of looking at the trace
32 according to different modelling goals.

33 The inference rules produce inferred symbols
34 from patterns of previously existing symbols.
35 The principle is to query the graph in the search
36 for sub-graphs that match certain patterns, and
37 to attach new nodes and arcs to the matching
38 sub-graphs. These new nodes represent the in-
39 ferred symbols and the new arcs represent the
40 *inference* relations. The usage of this inference
41 mechanism is further described in Section 3.5.

42 Technically, the *sequence* part of the activity
43 traces is encoded as resource description frame-
44 work (RDF) graphs. We choose RDF because it
45 is the most widely used specification for graph
46 encoding. We use XML as a serialization of

RDF to store sequences, because XML makes
RDF graphs easy to share with other applica-
tions. The ontology is encoded as RDF schema
(RDFS), because RDFS is the simplest ontology
language based on RDF. We use Protégé as a
graphical ontology editor. That is, an installa-
tion of Protégé is embedded in our tool, and the
analyst uses it to define the ontology of his
traces. The graphical displays of our traces are
encoded under the scalable vector graphic
(SVG) specification. Because Firefox natively
supports SVG, we use it as a visualization tool,
and we have implemented most of the tool as a
web application in PHP. We use extensible
stylesheet language (XSL) as a transformation
language for transforming RDF traces into
their SVG graphical representation. We use
SPARQL as a query language for graphs, as we
will explain in Section 3.5.

3.3. *The documentation system*

Analysts using our symbolic trace system ex-
pressed the need for a higher system layer provid-
ing a way to both index and attach documentation
to episodes of interest. We implemented this by
associating the symbolic trace system with a Wiki.
As we have made the choice of implementing
ABSTRACT as a web application, analysts can
reference each episode of interest by their URL,
and easily paste this URL into a Wiki page.
Moreover, some new Wiki implementations, like
Semantic MediaWiki,¹ include semantic facilities.
We are still investigating how these semantic
facilities can be used to merge the ontology editor
with the documentation system into a single
semantic documentation system.

3.4. *System usage*

The user interface is accessible as a web page
in any browser that supports SVG, such as
Firefox. This interface is illustrated in Figure 4.
It has four tabs: the *Open* tab that allows the
analyst to select a trace in a list; the *Info* tab that
displays general information about the selected
trace, such as its creation date and its version,

¹<http://semantic-mediawiki.org>

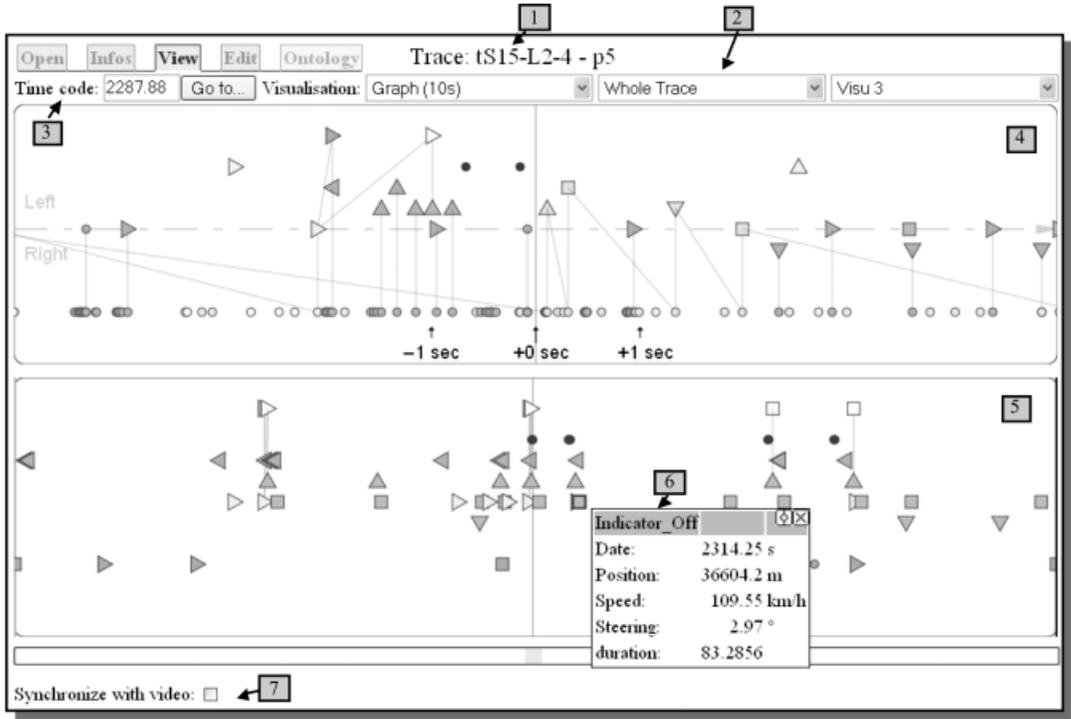


Figure 4: ABSTRACT user interface.

and the management of stylesheets; the *View* tab that displays graphical visualizations of the trace; and the *Edit* tab that allows the analyst to write queries to transform the trace. The interface also provides a link to the ontology editor, Protégé.

The *View* tab shown in Figure 4 provides the following functionalities (noted with numbered boxes):

1. Unique ID of the analysis, including aspects of the trace file and analyses done.
2. Selection of different visualization style sheets in drop-down lists. Different visualizations can be displayed simultaneously on the screen and their time code is synchronized.
3. Time code: this value corresponds to the cursor position in the visualization modules (vertical red line). The analyst can enter a time code and click *Go to* to focus on it.
4. Visualization sample with a time span of 10 s. The analyst can scroll the trace left and

right with the mouse. This visualization example corresponds to the simplified description given in Figure 2.

5. Visualization sample of an entire trace (20 min), with only the high-level symbols displayed. These trace examples come from our car-driving study and are further explained in Section 4.
6. The analyst can show the symbols' properties by clicking on the symbols.
7. The system can synchronize with a video player. When this box is checked, the system gives the time-code control to the video player and automatically follows it.

3.5. The transformation mechanism

The analyst uses the *Edit* tab to write queries that infer higher-level symbols from patterns of lower-level symbols. For instance, Table 1 illustrates a query to infer the *Lane change indicator symbol* shown in Figure 2, which indicates that a

Table 1: Simplified inference query to infer a lane change

```
CONSTRUCT
(?r1, infer, Indicator_Symbol)
(?r2, infer, Indicator_Symbol)
(Indicator_Symbol, type, Lane_Change_Indicator)
WHERE
(?r1, type, Accelerate)
(?r2, type, Left_Mirror_Glance)
(?r1, time-code, ?d1)
(?r2, time-code, ?d2)
(?r2, Acceleration_Value, ?v1)
FILTER
(?v1 > 1)
(sequence(?d1,?d2,1)) # in order, within 1 second
```

lane change is about to happen. In this example, the analyst wants to test the hypothesis that this indicator can be inferred from a conjunction of an *accelerate* event with an acceleration value $\geq 1 \text{ m/s}^2$, followed by a *glance* event pointing to the left mirror (generated by an eye-tracker), both within 1 s of each other.

The graph elements, either from the sequence or the ontology, are handled as triples [node, edge, node]. A query consists of a selection clause (WHERE) and a CONSTRUCT clause. The selection clause specifies a pattern of triples that should match the graph, and the CONSTRUCT clause specifies a pattern of triples that should be added to the graph wherever a pattern matches the selection clause. In addition, matching patterns can be restricted by a FILTER clause. The syntax shown in Table 1 has been simplified for clarity.²

In this query, ?r1, ?r2, ?d1, ?d2 and ?v1 represent variables. Each of them must match the same graph element each time they appear in the query. The *sequence* function tests that the time codes ?d1 and ?d2 occur in order and within the parameter of 1 s of each other.

In our implementation, the analyst has to know SPARQL to specify queries on the trace. To make it simpler, however, we have imple-

mented a template mechanism that prepares skeletons of queries. We have also added some customized functions in our implementation of SPARQL, such as the *sequence* function presented above. These functions facilitate the specification of queries that compare time codes of events, and make it easier to specify temporal constraints. In so doing, we are implementing semantics of time, for example, the semantics of the relation of co-occurrence and of sequential ordering. In the future, we plan to add options to let the analyst specify these queries from a graphical interface based on the visualization of the trace.

The *Edit* tab allows the analyst to visualize the resulting trace in a similar way as the *View* tab, but it also allows her to reject the trace if she is not satisfied by the result. This feature helps the analyst refine the query in search for the best symbols and inference rules she can get to describe the activity from the actual data. Our system returns the number of times the pattern has matched in the trace. This number indicates the number of new symbols added. The system also provides an export function to a text file that can be imported into other tools like Microsoft Excel for further statistical computations. Queries are saved as independent files and the tool helps the user reference them. The database of queries associated with the ontology constitutes a representation of the analyst's understanding about how to make sense of the activity trace.

4. Example activity model

We report here an example activity modelling analysis taken from a car-driving study (Henning *et al.*, 2007). Another example analysis – in a study of non-state political violence – is reported by Georgeon *et al.* (2010). Figure 5 shows a 10 s section of a car-driving activity trace focusing on a lane change on a motorway. The legend is presented in Figure 6.

In Figure 5, the 'Button' is an index signal from the experimenter recorded during the experiment, the *start thinking* event comes from

²The complete SPARQL syntax can be found in the SPARQL documentation, <http://www.w3.org/TR/rdf-sparql-query/>

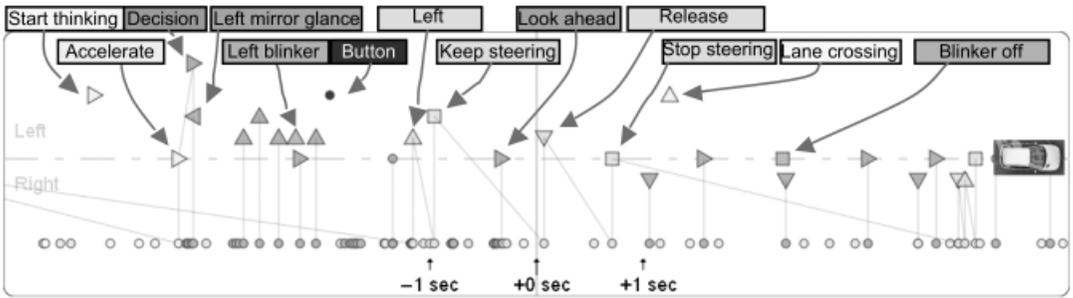


Figure 5: Example of lane change on a motorway (screenshot with text labels added at the top).

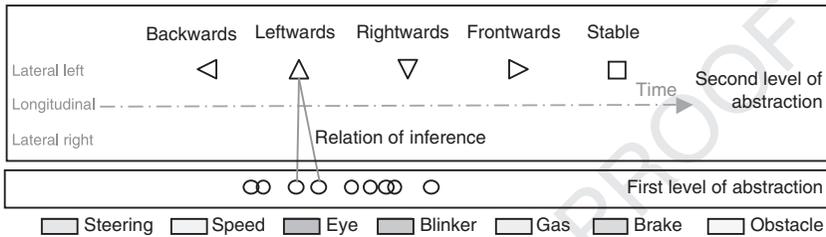


Figure 6: Legend of car driving symbolic trace.

a verbal signal given by the driver in a video-based post-experiment interview, and the *lane crossing* event is the moment when the left front wheel crosses the lane, manually encoded from the video. The representation of the driving episode given in Figure 5 is then automatically generated with the inference rules defined by the analyst.

Using the trace querying facilities of ABSTRACT on this driving data set, we could identify two categories of lane changes that we explain by the performance of two strategies (Figures 7 and 8). In these descriptions, the lower part is a representative trace episode from our database, while the upper part is drawn by hand as an abstract description of the strategies. This upper part also shows the car trajectory in the lanes, respecting the scale ratio length/width: about 300 m of 4 m wide lanes.

The strategy displayed in Figure 7 is characterized by beginning in a situation where the subject is impeded by a slow vehicle. In this case, the subject starts accelerating [1] almost at the same time as he looks at his left mirror [2]. Then, if

there is no vehicle coming from behind, he starts looking at the left lane [3], he switches his blinker on [4] and he performs the lane change [5]. In this situation, the acceleration associated with a glance to the left mirror appears as a good predictor of the lane change. It occurs more than 1 s before the subject switches the blinker on.

In the situation of Figure 8, no slow vehicle impedes the subject and he performs the lane change 'on the fly'. In this case, we can find no behavioural sign of his intention to change lane before the blinker is switched on [1]. Nevertheless, the blinker appears to be a sufficient predictor in that case, because it is switched in anticipation of the lane change, several seconds before the manoeuvre: looking to the left lane [2], looking to the left mirror [3] and starting steering [4].

In parallel to searching and identifying these categories of situations and strategies, we define symbols to represent them and inference rules to generate these symbols. Finally, we have named the first strategy *Lane_change_delayed* and represented it with white triangles, and the second strategy *Lane_change_anticipated* represented

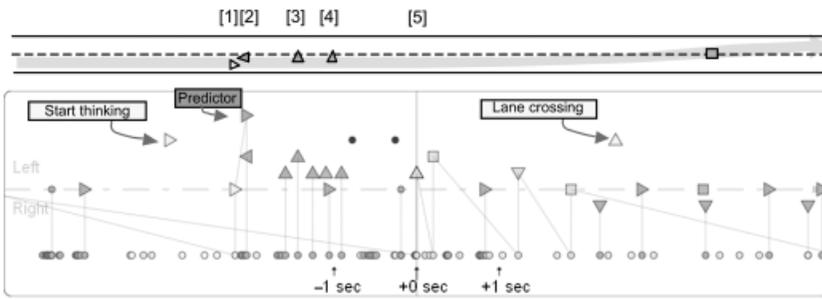


Figure 7: Lane change with acceleration (*Lane_change_delayed*).

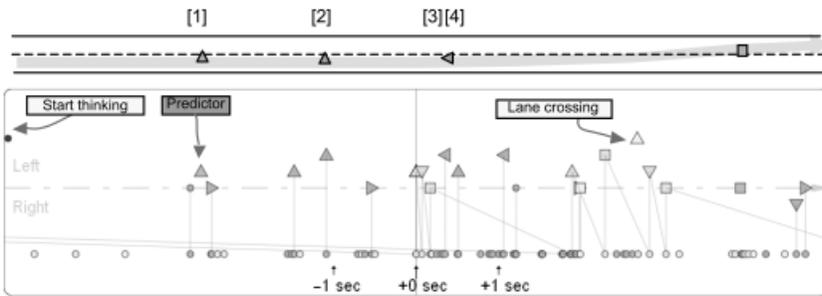


Figure 8: Lane change anticipated without acceleration (*Lane_change_anticipated*).

with white squares. Figure 9 shows that these two types of event occur four times in a representative 20-min motorway ride of a subject.

Figure 9 displays only the symbols that are useful to see the lane changes: the index button pressed by the experimenter (blue circles), the blinker (orange triangles up and down), the accelerations (ochre triangles to the right), the left mirror glances (grey triangles to the left), main junctions on motorways given by the GPS position (green square and triangles on the bottom of the display) and the lane change category symbols (white triangles and squares). In this example, one lane change (marked by the vertical cursor line in the figure) was not categorized. The uncategorized cases would require further study to understand their specificity. Once the analyst has finished her analysis, she can export the abstract traces into a spreadsheet to compute and report the statistics of the occurrences of events of interest.

Despite the extensive existing studies of car driving (e.g. Groeger, 2000) and on lane change

manoeuvres (e.g. Salvucci & Liu, 2002), we could find no representation of the driving activity that could compare to ours in terms of comprehensiveness of the data and capacity to support higher-level understanding. In the case of lane changes, this innovative description helped us discover different strategies that have not been reported in the literature before. In so doing, this study shows how our principles and tool have addressed the needs set out in the introduction: our prototype tool helped us generate comprehensive symbolic representations of the activity at an appropriate abstraction level to discover previously unknown knowledge about the activity. It also helped us explain, report and back up our models with the collected traces.

5. Related work

In this section, we situate our work in relation to two research areas: the area of qualitative data

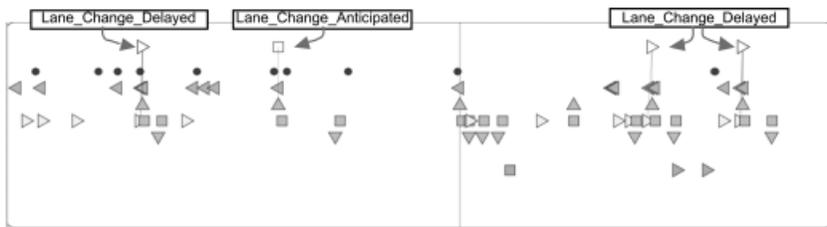


Figure 9: Categorization of lane changes between Lyon and airport (20 min long).

analysis and the area of trace-based reasoning (TBR).

In the area of qualitative data analysis, many software tools address the need for supporting the transcription of raw data into sequences of encoded events, for example, Dismal (Ritter & Larkin, 1994; Ritter & Wood, 2005), NVivo, INTERACT, InfoScope, MORAE, MarShapa (Sanderson *et al.*, 1994) and MacVisSTA. As such, these tools relate to our collection system as described in Section 3.1. Among these tools, MacVisSTA (Rose *et al.*, 2004) particularly relates to this aspect of our work in that it supports merging multi-modal data into a common timeline.

At the symbolic level – also called the transcript level in some studies – tools like Theme (Magnusson, 2000) support the automatic discovery of temporal patterns based on statistical properties. We consider our tool complementary to these tools because our tool helps find the symbolic pattern based on the meaning they have to the analyst rather than on their statistical properties. In the car-driving example, our symbols of interest are not particularly frequent or infrequent nor do they obey pre-assumed statistical laws.

HyperRESEARCH appears to be the only qualitative data analysis tool that supports the validation of hypothetic theories through rule-based *expert system* techniques (Hesse-Biber *et al.*, 2001). We find HyperRESEARCH's underlying principles for theory building very similar to ours. It, however, does not focus on temporal semantics and does not offer symbolic timeline visualization facilities to support activity modelling. Its rule engine also does not exploit elaborated hierarchical semantics defined in an ontology, as opposed to our solution based on SPARQL and RDFs.

The other related research area, TBR (Cordier *et al.*, 2009), comes from the domain of knowledge representation, and more precisely from case-based reasoning (CBR) (Aamodt & Plaza, 1994). CBR consists of helping users solve new problems by adapting solutions that have helped them solve previous problems. TBR extends CBR to retrieve useful cases as episodes from the stream of an activity trace (Mille, 2006). Our work relates to TBR because both address the question of modelling and representing a stream of activity for future usage. In particular, we pull lessons from the work of Settouti *et al.* (2009) that implemented a trace-based system to support the management and the transformation of traces. TBR has been used to implement companions that provide assistance to the user based on previous usage (e.g. Cram *et al.*, 2008) or to support reflexive learning by providing the users with a dynamic display of his passed activity (Ollagnier-Beldame, 2006). As opposed to these works, our work helps a user (the analyst) understand the activity of another user (the subject). In so doing, our work brings principles of qualitative data analysis to TBR and brings TBR techniques to address problems of quantitative data analysis.

6. Conclusion

We have defined the principles of a methodology and a software tool to help an analyst create models of activity from activity traces in an iterative and interactive fashion. These principles relate to the notion of *abductive reasoning* in that they consist of helping analysts form, organize and test micro-hypotheses to explain and represent the activity. Specifically, they help

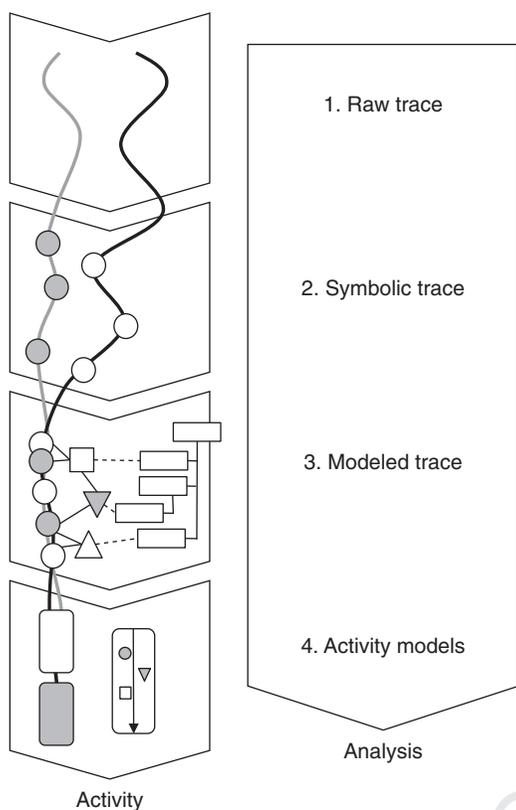


Figure 10: Activity modelling from activity traces

investigate what concepts and semantics describe the activity the best and represent these concepts in an ontology and apply this semantics through a rule engine. These principles are summarized in Figure 10.

In Figure 10, the collected raw data are represented as curves along the activity axis. The first analysis step consists of collecting this raw trace. The second step consists of producing a symbolic trace (symbols represented by circles) through the discretization of the raw trace. The third step consists of modelling the symbolic trace by inferring more abstract symbols (represented as squares and triangles) and organizing these symbols in an ontology (hierarchy of white rectangles). The fourth step consists of producing explained models of activity (round-angled rectangles) that are backed up by the abstract trace. During the

modelling process, the analyst formalizes her understanding of the activity in the form of transformation rules, ontology and documentations that are stored in the system, which allows capitalizing on the analyst's knowledge across studies.

To illustrate these principles, we have implemented a prototype software tool through an assemblage of open-source knowledge engineering software modules. With this tool, we have modelled car-driving activity traces collected with an instrumented vehicle. This analysis allowed us to identify and describe two strategies of lane change on motorways. This example shows that our knowledge engineering approach of activity modelling from activity traces offers answers to the needs for tools to help analysts understand better an observed activity, create models of this activity, report and back up these models with the observational data.

The abstract activity traces that we have constructed constitute a model of the car driver in their own, in that the analysts can use our tool to query these traces to answer new questions they may have about the driving activity. Our future developments will consist of simulating the activity, for instance, in the field of car driving, generating realistic driving behaviour in a driving simulator, based on our abstract activity traces.

Acknowledgements

We would like to acknowledge support for this project from the European Commission through the HumanIST Network of Excellence. Partial support for this report was provided by ONR (contracts N00014-06-1-0164 and N00014-08-1-0481). We also thank Jean-Marc Trémeaux for his participation in the software implementation, and Matthias Henning for his contribution to the car-driving experiment. We appreciate the comments on this report from Jonathan Morgan and Ryan Kaulakis.

References

AAMODT, A. and E. PLAZA (1994) Case-based reasoning: foundational issues, methodological variations, and system approaches, *AI Communications*, 7, 39–59.

- 1 CORDIER, A., B. MASCRET and A. MILLE (2009) Extending case-based reasoning with traces, *Grand Challenges for Reasoning from Experiences, Workshop at IJCAI*, Pasadena, CA, pp. 23–32.
- 2
- 3 Q1 CRAM, D., B. FUCHS, Y. PRIÉ and A. MILLE (2008) An approach to user-centric context-aware assistance based on interaction traces, *Modeling and Reasoning in Context, Human Centered Processes*, Delft, The Netherlands.
- 4
- 5
- 6 Q2 DEARDEN, A. and J. FINLAY (2006) Pattern languages in HCI: a critical review, *Human-Computer Interaction*, **21**, 49–102.
- 7
- 8
- 9 ERICSSON, K.A. and H.A. SIMON (1993) *Protocol Analysis: Verbal Reports as Data*, Cambridge, MA: MIT Press.
- 10
- 11 GEORGEON, O., J. MORGAN, J. HORGAN and K. BRADDOCK (2010) Process modeling for the study of non-state political violence, *19th Annual Conference on Behavior Representation in Modeling Simulation*, Charleston, NC, Brims Society, pp. 240–247.
- 12
- 13 Q3 GROEGER, J. (2000) *Understanding Driving*, Hove, UK: Psychology Press.
- 14
- 15
- 16 HENNING, M.J., O. GEORGEON and J.F. KREMS (2007) The quality of behavioral and environmental indicators used to infer the intention to change lanes, *4th International Driving Symposium on Human Factors in Driver Assessment*, Stevenson, Washington, USA, pp. 231–237.
- 17
- 18 Q4 HESSE-BIBER, S., P. DUPUIS and S. KINDER (2001) Testing hypotheses on qualitative data: the use of HyperRESEARCH computer-assisted software, *Social Science Computer Review*, **18**, 320–328.
- 19
- 20 HILBERT, D.M. and D.F. REDMILES (2000) Extracting usability information from user interface events, *ACM Computing Surveys (CSUR)*, **32**, 384–421.
- 21
- 22 Q5 JAMES, W. (1907) *Pragmatism*.
- 23
- 24 MAGNUSSON, M.S. (2000) Discovering hidden time patterns in behavior: T-Patterns and their detection, *Behavior Research Methods, Instruments and Computers*, **32**, 93–110.
- 25
- 26 MILLE, A. (2006) From case-based reasoning to traces-based reasoning, *Annual Reviews in Control*, **30**, 223–232.
- 27
- 28 OLLAGNIER-BELDAME, M. (2006) Traces d'interactions et processus cognitifs en activité conjointe: Le cas d'une co-rédaction médiée par un artefact numérique, *CNRS – LIRIS*, Lyon, Université Lumière Lyon 2.
- 29
- 30 OLSON, G.M., J.D. HERBSLEB and H.H. REUTER (1994) Characterizing the sequential structure of interactive behaviors through statistical and grammatical techniques, *Human-Computer Interaction*, **9**, 427–472.
- 31
- 32 POPPER, K. (1972) *Objective Knowledge*, Oxford, UK: Oxford University Press.
- 33
- 34 RITTER, F.E. and J.H. LARKIN (1994) Developing process models as summaries of HCI action sequences, *Human-Computer Interaction*, **9**, 345–383.
- 35
- 36 RITTER, F.E. and A.B. WOOD (2005) Dismal: a spreadsheet for sequential data analysis and HCI experimentation, *Behavior Research Methods, Instruments, and Computers*, **37**, 71–81.
- 37
- 38 ROSE, R.T., F. QUEK and Y. SHI (2004) MacVisSTA: a system for multimodal analysis, *International Conference on Multimodal Interfaces*, State College, PA, pp. 259–264.
- 39
- 40 SALVUCCI, D.D. and A. LIU (2002) The time course of a lane change: driver control and eye-movement behavior, *Transportation Research, F*, 123–132.
- 41
- 42 Q6 SANDERSON, P.M. and C.A. FISHER (1994) Exploratory sequential data analysis: foundations, *Human-Computer Interaction*, **9**, 251–317.
- 43
- 44 SANDERSON, P.M., M.D. MCNEESE and B.S. ZAFF (1994) Handling complex real-word data with two cognitive engineering tools: COGENT and MacSHAPA, *Behavior Research Methods, Instruments, and Computers*, **26**, 117–124.
- 45
- 46 SETTOUTI, L.S., Y. PRIÉ, J.-C. MARTY and A. MILLE (2009) A trace-based system for technology-enhanced learning systems personalisation, *the 9th IEEE International Conference on Advanced Learning Technologies*, Riga, pp. 93–97.
- 47
- 48 Q7

The authors

Olivier L. Georgeon

Olivier L. Georgeon is a research associate in the LIRIS department at the Claude Bernard University (Lyon, France). He previously had a 12-year industrial experience as a software engineer, developer and project manager in the domain of automatization of industrial processes. He received a PhD in cognitive psychology from the Université Lumière (Lyon, France) in 2008. His research interests are in *learning through activity*. His work both produces practical applications and addresses epistemological questions concerning learning from and about an activity.

Alain Mille

Alain Mille has been a professor at the Claude Bernard University (Lyon, France) since 2000. He is a scientific director in the LIRIS department (UMR 5205 CNRS), and leads the research group *Supporting Interactions and Learning through Experience*. After a first career as a computer engineer and a project manager (Hospices Civils de Lyon), Alain Mille was

1 asked to build a computer science department in
2 a college of engineering (CPE-Lyon). In this
3 latter position, he created a research team on
4 case-based reasoning. The main topics were
5 decision helping and complex tasks assistance.
6 Alain Mille is on several national and interna-
7 tional conference programme committees, and
8 on the editorial board of *Intellectica* (a cognitive
9 sciences journal). He is specifically interested in
10 dynamic modelling of knowledge during com-
11 puter-mediated activities.

12 13 **Thierry Bellet**

14
15 Thierry Bellet has been a researcher at the
16 Ergonomics and Cognitive Sciences Laboratory
17 of IFSTTAR (Institut Français des Sciences et
18 Technologies des Transports, de l'Aménage-
19 ment et des Réseaux) since 1999. His main areas
20 of research concern human activity analysis and
21 cognitive processes modelling (e.g. situational
22 awareness, decision making, cognitive schemas).
23 His research has produced computational simu-
24 lations of mental activities of the car driver (e.g.
25 COSMODRIVE: COgnitive Simulation MOdel
26 of the DRIVER).

27 28 **Benoit Mathern**

29
30 Benoit Mathern is a PhD student in computer
31 science. He first worked as a computer engineer at

the Ergonomics and Cognitive Sciences Labora-
tory of IFSTTAR (Institut Français des Sciences
et Technologies des Transports, de l'Aménage-
ment et des Réseaux), and then he started his PhD
work in collaboration with the LIRIS department
of the Claude Bernard University (Lyon, France).
His main research interests focus on knowledge
engineering, knowledge discovery and human-
machine interaction. His work is applied to cog-
nitive science in the field of transportation.

32 33 **Frank E. Ritter**

34
35 Frank E. Ritter is one of the founding faculty of
36 the College of IST, an interdisciplinary aca-
37 demic unit at Penn State, to study how people
38 process information using technology. Frank
39 Ritter's current research is in the development,
40 application and methodology of cognitive mod-
41 els, particularly as applied to interface design,
42 predicting the effect of behavioural moderators
43 and understanding learning. He edits the
44 Oxford Series on Cognitive Models and Archi-
45 tectures, is associate editor of *Cognitive Systems*
46 *Research*, editorial board member of *Human*
47 *Factors*, and the *Journal of Educational Psychol-*
48 *ogy* and technical programme co-chair for the
BRIMS 2009, 2010 and 2011 conferences and
the associated special issues in *Computational*
and *Mathematical Organizational Theory*.

Author Query Form

Journal **EXSY**

Article **584**

Dear Author,

During the copy-editing of your paper, the following queries arose. Please respond to these by marking up your proofs with the necessary changes/additions. Please write your answers clearly on the query sheet if there is insufficient space on the page proofs. If returning the proof by fax do not write too close to the paper's edge. Please remember that illegible mark-ups may delay publication.

Query No.	Description	Author Response
Q1	AUTHOR: Please provide date of workshop for reference Cordier et al. (2009).	
Q2	AUTHOR: Please provide further details for reference Cram et al. (2008).	
Q3	AUTHOR: Please provide date of conference for reference Georgeon et al. (2010).	
Q4	AUTHOR: Please provide the date of symposium for reference Henning et al. (2007).	
Q5	AUTHOR: Please update the reference James (1907) with further details.	
Q6	AUTHOR: Please provide volume number for reference Salvucci and Liu (2002).	
Q7	AUTHOR: Please provide the date of conference for reference Settouti et al. (2009).	

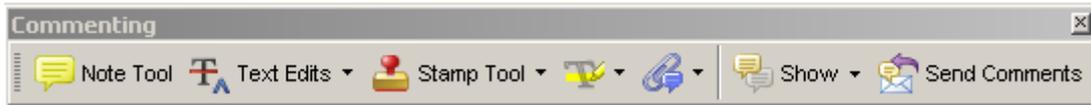


USING E-ANNOTATION TOOLS FOR ELECTRONIC PROOF CORRECTION

Required Software

Adobe Acrobat Professional or Acrobat Reader (version 7.0 or above) is required to e-annotate PDFs. Acrobat 8 Reader is a free download: <http://www.adobe.com/products/acrobat/readstep2.html>

Once you have Acrobat Reader 8 on your PC and open the proof, you will see the Commenting Toolbar (if it does not appear automatically go to Tools>Commenting>Commenting Toolbar). The Commenting Toolbar looks like this:



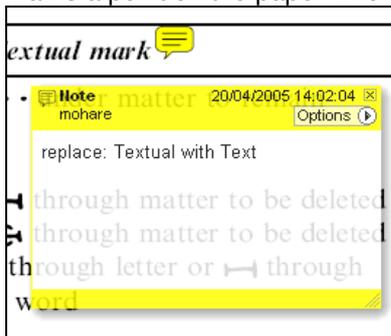
If you experience problems annotating files in Adobe Acrobat Reader 9 then you may need to change a preference setting in order to edit.

In the "Documents" category under "Edit – Preferences", please select the category 'Documents' and change the setting "PDF/A mode:" to "Never".



Note Tool — For making notes at specific points in the text

Marks a point on the paper where a note or question needs to be addressed.

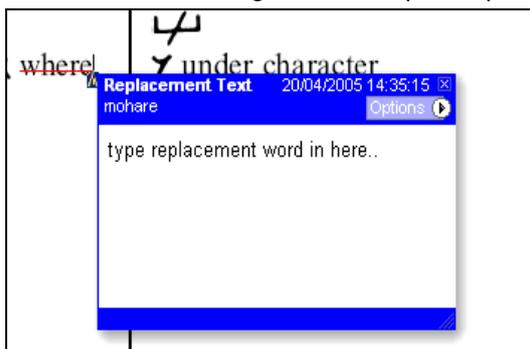


How to use it:

1. Right click into area of either inserted text or relevance to note
2. Select Add Note and a yellow speech bubble symbol and text box will appear
3. Type comment into the text box
4. Click the X in the top right hand corner of the note box to close.

Replacement text tool — For deleting one word/section of text and replacing it

Strikes red line through text and opens up a replacement text box.

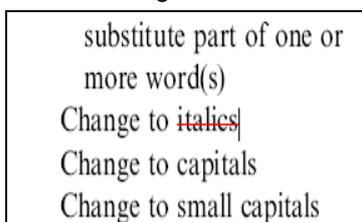


How to use it:

1. Select cursor from toolbar
2. Highlight word or sentence
3. Right click
4. Select Replace Text (Comment) option
5. Type replacement text in blue box
6. Click outside of the blue box to close

Cross out text tool — For deleting text when there is nothing to replace selection

Strikes through text in a red line.



How to use it:

1. Select cursor from toolbar
2. Highlight word or sentence
3. Right click
4. Select Cross Out Text

Approved tool — For approving a proof and that no corrections at all are required.

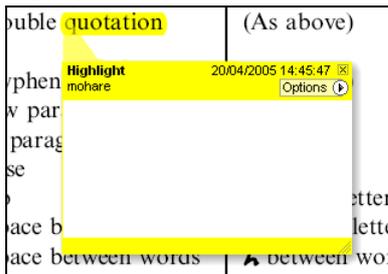


How to use it:

1. Click on the Stamp Tool in the toolbar
2. Select the Approved rubber stamp from the 'standard business' selection
3. Click on the text where you want to rubber stamp to appear (usually first page)

Highlight tool — For highlighting selection that should be changed to bold or italic.

Highlights text in yellow and opens up a text box.

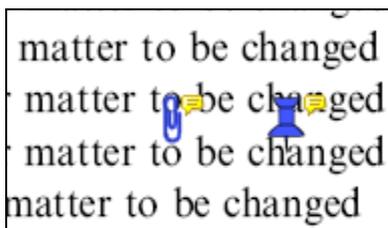


How to use it:

1. Select Highlighter Tool from the commenting toolbar
2. Highlight the desired text
3. Add a note detailing the required change

Attach File Tool — For inserting large amounts of text or replacement figures as a files.

Inserts symbol and speech bubble where a file has been inserted.

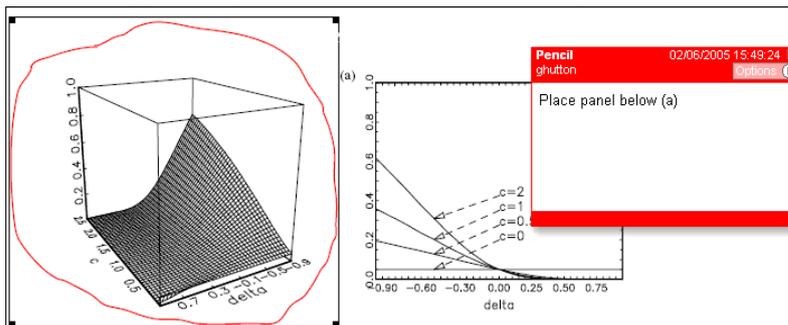


How to use it:

1. Click on paperclip icon in the commenting toolbar
2. Click where you want to insert the attachment
3. Select the saved file from your PC/network
4. Select appearance of icon (paperclip, graph, attachment or tag) and close

Pencil tool — For circling parts of figures or making freeform marks

Creates freeform shapes with a pencil tool. Particularly with graphics within the proof it may be useful to use the Drawing Markups toolbar. These tools allow you to draw circles, lines and comment on these marks.



How to use it:

1. Select Tools > Drawing Markups > Pencil Tool
2. Draw with the cursor
3. Multiple pieces of pencil annotation can be grouped together
4. Once finished, move the cursor over the shape until an arrowhead appears and right click
5. Select Open Pop-Up Note and type in a details of required change
6. Click the X in the top right hand corner of the note box to close.

Help

For further information on how to annotate proofs click on the Help button to activate a list of instructions:

