**QUEEN'S
UNIVERSITY
BELFAST**

A Speech Based Approach to Surveillance Video Retrieval

Qasemizadeh, B., O'Neill, I., Shen, J., Miller, P., Hanna, P., Stewart, D., & Wang, H. (2009). *A Speech Based Approach to Surveillance Video Retrieval*. 336-339. Paper presented at Sixth IEEE International Conference on Advanced Video and Signal Based Surveillance, Genova, Italy.

Queen's University Belfast - Research Portal:

[Link to publication record in Queen's University Belfast Research Portal](#)

General rights

Copyright for the publications made accessible via the Queen's University Belfast Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The Research Portal is Queen's institutional repository that provides access to Queen's research output. Every effort has been made to ensure that content in the Research Portal does not infringe any person's rights, or applicable UK laws. If you discover content in the Research Portal that you believe breaches copyright or violates any law, please contact openaccess@qub.ac.uk.

A Speech Based Approach to Surveillance Video Retrieval

Behrang QasemiZadeh

Unit for Natural Language Processing, DERI
National University of Ireland, Galway
Galway, Ireland
behrang.qasemizadeh@deri.org

Jiali Shen, Ian O'Neill, Paul Miller, Philip Hanna,
Darryl Stewart, Hongbin Wang

ECIT
Queen's University of Belfast
Belfast, UK

{j.shen, i.oneill, p.miller, p.hanna, d.stewart, h.wang}
@ecit.qub.ac.uk

Abstract—This paper describes the anatomy of a pilot surveillance system with a speech-based interface for content-based retrieval of video data. The proposed system relies on an ontology-based information sharing architecture and lets components of the system communicate among each other through TCP/IP communication channels. The aim of developing the pilot system was to explore dependencies between image analysis, event detection, video annotation, and speech-based retrieval of the video content in the context of a broader spoken dialogue system.

Surveillance video retrieval; Ontology; Natural language Interface; Video Content Modelling; Event Detection.

I. INTRODUCTION

The ISIS project is a multidisciplinary EPSRC-funded project based at ECIT, Queen's University's institute for Electronics, Communications and Information Technology. The specific aim of the project is to create an intelligent sensor network capable of detecting threats to passenger safety on public transport. However, the individual techniques for detection, classification and information retrieval are in many cases highly generic and relevant to many application contexts. Among the initial experiments conducted for ISIS was an attempt to classify by gender individuals passing through an office doorway (the Doorway Experiment) and then to retrieve footage of a "person", a "male" or a "female" passing through the doorway within a given time period. With the eventual aim of facilitating information retrieval in a busy operations centre, the retrieval enquiry took the form of a spoken natural language query.

A second scenario took as its setting a more complex scene outside an office: here hand-tagged data (based on the British Home Office's i-LIDS data set [1] and representing individuals described by clothing and gender and vehicles described by colour) were used as a basis on which to formulate enquiries concerning the presence, activity and/or location of individuals and vehicles that matched a particular description. While this scenario (the i-LIDS Experiment) proved valuable in terms of developing an underlying data model, its need for hand-tagging and a rigid set of spoken enquiries meant that the simpler Doorway Experiment was a more accurate representation of a replicable technique for automated retrieval and dynamically-phrased information retrieval.

In this paper our description of experimental activity concentrates on our Doorway Experiment, with some reference to the more advanced avenues of enquiry suggested by the forward-looking i-LIDS Experiment.

A. The Doorway Experiment

As presented to the members of the i-LIDS consortium in the autumn of 2008, the Doorway Experiment showed how natural language could be used to progressively refine a user enquiry. In this experiment natural language enquiries adhered to a fixed structure, corresponding to a predefined underlying syntactic grammar, with the expectation that the variables required by the system (gender, start-time, end-time and day) would occur in fixed positions in each user utterance. This template-based approach offered enough flexibility for the user to change variable values at will and so vary the nature of the enquiry over several iterations.

The underlying dataset represented six individuals (three males and three females) passing through an office doorway in the course of a few minutes. Inquiries included:

- Show me the key frames of a person entering the doorway between ten thirty a.m. and eleven a.m. on October the ninth.
- Show me the key frames of a person entering the doorway between ten forty-five a.m. and eleven a.m. on October the ninth.
- Show me the key frames of a male entering the doorway between ten forty-five a.m. and eleven a.m. on October the ninth.
- Show me the key frames of a female entering the doorway between ten forty-five a.m. and eleven a.m. on October the ninth.

II. THE ANATOMY OF THE FULL SYSTEM

Figure 1 shows the anatomy of the prototype system in terms of its building blocks, a set of agents run on different TCP/IP communication channels.

The Data Manager Agent gets the image analysis results from the Image Analysis Agent and it writes the time instants, and the values of annotations, in the appropriate locations in the "image annotation database". The Prolog Inference Engine Agent uses "time stamps" and detects events from the annotations in the "image annotation database". Then, it asks the Data Manager Agent to write the inferred events into the "event annotation database". The

Natural Language Interface Agent gets spoken natural language utterances and parses the input into an intermediate representation encapsulating information about the type of query, e.g. information seeking, in addition to a first order logic representation of natural language utterances. The Knowledge Management Agent then looks at the information provided by the user's utterances and the data repositories to provide the user with a video compilation – or another mode of response, such as a spoken or written answer.

Figure 1. The Anatomy of the full system. Each block shows an agent run on a TCP/IP communication channel. The hatched lines show references, while the solid lines show the data flow among agents.

To make sure the system functions correctly in terms of semantics, a set of ontologies supports the system's agents. The ontologies ensure that the system uses the same vocabulary for the same semantic concepts in different agents. There are four classes of ontology in the system: a time ontology, properties ontologies, an object ontology, and an events ontology. The definitions in the Time Ontology are based on DAML-Time [2] and its recent version [3], and together with the time stamps, these serve the temporal needs of the system – e.g. providing temporal tags as well as an inference engine when temporal query analysis is required. A Property Ontology provides vocabulary for describing an attribute of an object i.e. what can be sensed by sensors – e.g. if colour is considered as an attribute for vehicle objects, then the system's agents will use only the vocabulary introduced by the "Colour Property" to annotate or retrieve the video. The Property Ontology provides other services, such as comparing two vocabularies or mapping them

together. The Object Ontology provides a classification of objects and it may be varied from one application domain to the next. The most important role of the Object Ontology is to provide the system with vocabularies of shared concepts in a domain of interest – e.g. in the doorway scenario, the doorway, people, and stationary objects like computers might be the objects of interest. An Event Ontology like the one defined in [4] presents definitions for events, and the temporal or causal relationships between them. In the following, we briefly describe each of the system's agents.

A. Image Analysis Agent

The Image Analysis Agent can identify people and their properties, such as gender. At the current stage of development within ISIS gender profiling is based on the human face. The process is outlined in Fig. 2. A doorway camera is positioned to capture the faces of those who have just entered the target zone. Viola's technique [5] is used for face detection. The system is trained using AdaBoost [5] to select the most discriminative rectangle features, an approach that has been used in many applications.

Then Principal Component Analysis (PCA) is carried out on the normalized faces (24*24 pixels). In principal component analysis (PCA), eigenvalues of the matrix composed by training faces are calculated by ranking the eigenvectors with respect to their eigenvalues and selecting the first M principal components. Theory suggests that these components contain most of the information needed to separate the training faces. The first 5 eigenvectors are shown in Figure 3. A conversion matrix is also generated in this step, which is applied to transfer the input testing face into the new space.

Figure 2. Flowchart of human gender profiling.

Figure 3. The first four eigenfaces.

In this system, the first 50 components are input to a

trained Support Vector Machine (SVM) as the gender classifier. The SVM system has been trained by nearly 4000 faces with same PCA processing.

B. Data Manager Agent

The Data Manager Agent is responsible for any transaction on the data repositories. This includes usual data administrative tasks in addition to checking data consistency, and compatibility against the system's ontologies. For example, the Data Manager checks whether annotations provided by the Image Analyzer Agent are valid according to the property and object ontologies. In addition, the Data Manager Agent records all the transaction times from the Image Analyzer Agent and keeps track of the temporal order of annotations. When retrieving information, the Data Manager Agent assigns time tags to image annotations and provides a high level representation of actual stored information. The following shows an example of an annotation provided by the Data Manager Agent:

```
<object type="vehicle" id="v1"time="8">
  <properties>
    <colour>blue</colour>
    <location>Gp12528</location>
  </properties>
</object>
```

Figure 4 shows the steps the Data Manager Agent takes when it gets a new input from the Image Analyzer Agent.

Figure 4. The figure shows the steps the Data Manager Agent takes when it receives new annotations from the Image Analyzer.

C. Prolog Inference Engine

The Prolog Inference Engine, or Event Detection Agent, detects events by monitoring annotations of objects in time instants. The Event Detection Agent recognizes events in a similar way to [6]. Then it asks the Data Manager Agent to assert inferred facts into the preserved data repository. An example of annotation for an event can be as follows:

```
<event type="move" id="jzglkncoos">
  <event_objects>
    <agent>v1</agent>
  </event_objects>
  <event_details>
    <from_location
      inherited_from="location">
      oos
    </from_location>
  </event_details>
  <time stime="0" etime="2"/>
</event>
```

Figure 5 shows the steps the Data Manager takes when asserting events inferred by the Event Detection Agent.

Figure 5. Asserting an inferred event into data repository

D. Natural Language Interface

The main components of the Natural Language Interface are a speech recognizer and a semantic parser. The speech recognizer used in this case was Sphinx [6]. The semantic parser is formed around ontology-related vocabularies. When the Semantic Parser gets an input, it looks for a proper mapping from the natural language utterances to the semantics that are defined in the ontologies of the system. This is done by a set of Definite Clause Grammar rules implemented in Prolog. The following is an example parse of a natural language utterance:

```

<Event prolog_rule_id="eventlexical1"
event_type="get_into" id="49702660200">
  <lexme pos="verb" class="occurance"
tense="none" aspect="progressive">
    entering
  </lexme>
</Event>

```

In the given example, the semantic parser has mapped an utterance like “entering” to a XML entity of the type “event”. Furthermore, as the above XML shows, the utterance “entering” is classified as an event of semantic type “get_into”. The event of type “get_into” has a formal definition in the system’s event ontology.

E. Knowledge Manager

The Knowledge Manager agent gets a semantic parse tree as an input; and it maps the parse information on to the facts previously asserted by the Data Manager Agent. The Knowledge Manager may use additional inference rules – e.g. to detect higher level types of events or to map linguistic variables on to their canonical forms. The process is as follows:

- 1) Find the pattern of the input parse tree,
- 2) Extract facts from the parse tree and assert them into the memory,
- 3) Translate the facts into their canonical forms,
- 4) Locate video annotations with the criteria specified by extracted facts, and
- 5) Perform a video compilation based on the located video annotations.

At the third step, translation of facts into their canonical forms, the Knowledge Manager may use ontologies at different levels of granularity to perform its task. For example, it may use definitions for higher level types of events, or it may use definitions at a finer level of granularity to translate facts and locate video annotations.

III. EXPERIMENTAL RESULTS

A. Visual Detection (Offline)

The Gender Profiling module was tested with 1254 face pictures (641 male and 615 female). This is to test the performance of the gender profiling module. In the experiment, 532 males and 500 females were correctly profiled, more details is listed in the Table 1.

B. Visual Detection (Online)

The integrated system was tested in a lab equipped with a doorway camera. This is to test the performance and robustness of the whole system. During the 3 hours running, 22 instances of people entering (5 females among them) the doorway, and 100% of them were detected and correctly profiled.

TABLE I. GENDER PROFILING PERFORMANCE

Gender	Correct/All	Accuracy
Female	500/615	81.30%
Male	532/641	83.00%
General	1032/1256	81.97%

IV. CONCLUSION, LOOKING AHEAD

The paper describes the architecture of a pilot surveillance system with a natural language interface. We have designed, and developed an agent-based system that employs ontologies for information sharing among the agents. Future enhancements to the experimental system, for which technical solutions have already been identified, include enabling the system speak a commentary on the information retrieved – the commentary will correspond to the selection criteria uttered by the user – as well as displaying the user’s selection criteria as confirmatory on-screen text. Our longer-term goal is to develop the experimental natural language interface into a comprehensive interactive natural language dialogue system. A dialogue-based information retrieval system will:

- prompt operators for important search criteria that they have omitted;
- suggest more discriminating constraints that the operator may wish to use for narrower searches;
- make and confirm inferences concerning vaguely formulated operator enquiries;
- and confirm turn-by-turn modifications to search constraints that the operator may introduce in the course of an evolving dialogue.

ACKNOWLEDGMENT

The demonstration system was developed as part of the ISIS project under Grant No. EP/E028640/1.

REFERENCES

- [1] Home Office Scientific Development Branch Imagery library for intelligent detection systems (i-LIDS), <http://scienceandresearch.homeoffice.gov.uk/hosdb/cctv-imaging-technology/video-based-detection-systems/i-lids/>
- [2] J. R. Hobbs, G. Ferguson, J. Allen, P. Hayes, I. Niles, and A. Pease, “A DAML Ontology of Time”, 2002.
- [3] J. R. Hobbs, and F. Pan, “Time Ontology in OWL”, Ontology Engineering Patterns Task Force of the Semantic Web Best Practices and Deployment Working Group, World Wide Web Consortium (W3C) notes, 2006.
- [4] A. R. Francois, R. Nevatia, J. R. Hobbs, and R. C. Bolles, “VERL: An Ontology Framework for Representing and Annotating Video Events”, IEEE MultiMedia 12, 2005.
- [5] P. Viola and M. Jones, “Rapid Object Detection Using a Boosted Cascade of Simple Features”, in Proc. of IEEE CVPR 2001, Vol. 1, pp 1-511 – 1-518, 2001.
- [6] V. D. Shet, D. Harwood, L. S. Davis, “VidMAP: Video Monitoring of Activity with Prolog”, AVSS, 2005.
- [7] Sphinx-4, A speech recognizer written entirely in the Java, <http://cmusphinx.sourceforge.net/sphinx4/>