VISUAL ANALYTICS FOR SENSE-MAKING IN CRIMINAL INTELLIGENCE ANALYSIS

VALCRI WHITE PAPER SERIES
VALCRI-WP-2017-003
1 January 2017
Edited by B.L. William Wong

Applying Data Science to Criminal Intelligence Analysis

Nadeem Qazi[1], Leishi Zhang[1], Eva Blomqvist[2], Florian Stoffel[3], Patrick Aichroth[4], Christian Weigel[4]

[1]Middlesex University London
The Burroughs, Hendon
London NW4 4BT
UNITED KINGDOM

[2]Linköpings University
SE-581 83 LINKÖPING
SWEDEN

[3]Universitat Konstanz
Universitaetsstrasse 10 78467 Konstanz
GERMANY

[4]Fraunhofer Institute for Digital Media Technology
Ehrenbergstraße 31
98693 Ilmenau
GERMANY

Project Coordinator
Middlesex University London
The Burroughs, Hendon
London NW4 4BT
United Kingdom.

Professor B.L. William Wong
Head, Interaction Design Centre
Faculty of Science and Technology
Email: w.wong@mdx.ac.uk

INTENTIONALLY BLANK

ABSTRACT

A major challenge of criminal intelligence analysis is to process large amount of semi-structured or unstructured data such as textual documents and videos and to extract useful information out of the data to support semantic search, sense-making and decision making. In VALCI, a computational framework is developed that incorporates concept extraction, ontology use and evolution, associative search, and image/video analysis for semantic search and knowledge discovery. In this whitepaper we introduce the key concepts have been applied and their corresponding technologies that have been developed to tackle the challenge.

Keywords

Criminal Intelligence, Semantic Search, Evolving Knowledge Base, Ontology, Associative Search, Concept Extraction, Video Analysis

INTENTIONALLY BLANK

SUMMARY

In this white paper, we present the main data science technologies that we apply for supporting information retrieval and knowledge extraction from large amount of of semi-structured or unstructured crime data. The applied techniques facilitate the use of ontologies for data management and semantic search, the evolution of ontologies, the extraction and analysis of crime-related concepts from data, the analysis of formal concepts in the data for associate search, the assessment of video qualities and the extraction of entities in video data. These technologies are integrated into a computational framework and act as a fundamental building block of the VALCRI System.

ONTOLOGIES

Ontologies are models of some domain of knowledge, most often formalised in a machine readable manner, and they include concepts and relations, as well as additional axioms that are relevant for that domain (Gruber, 1993). In the VALCRI project, we are creating a set of ontologies that correspond to some system requirements specified by the end user. They are expressed in a formal logical language, *W3C Web Ontology Language (OWL)* (Group, 2004), and used as part of the VALCRI system for realising a number of functionalities, including:

1. For expressing data models and allowing integrated access to data, e.g., through querying and search, from multiple diverse data sources, regardless of the underlying storage.
2. For expressing logical constraints and formalising definitions of concepts and relations, so that the system can perform automated reasoning on top of the underlying data using the ontologies. Automated reasoning, sometimes called inferencing, allows not only detecting inconsistencies in data, but also drawing conclusions and deriving new data from existing information.
3. As the basis for computing semantic similarity between elements and sub-structures in the ontologies. Data expressed according to a specific ontology is usually expressed as a graph, and there exist numerous semantic similarity measures that operate on such graphs.
4. As models for informing Information Extraction methods, which can then populate our knowledge base with data extracted from unstructured or semi-structured sources, such as natural language texts.

In summary, ontologies provide at least two beneficial aspects:

- Decoupling of knowledge/information from program code.
- Making assumptions and definitions explicit within the system.

Decoupling of models from program code, means that domain knowledge is not encoded within classes and procedures of the program code, but rather stored separately from the code, i.e., in the ontologies. In contrast to program code, which is (normally) not exchanged at runtime, or even after compile-time, ontologies can be modified and exchanged at any time, even at runtime. This makes an ontology-based system highly configurable. An example could be to encode different policies of the different countries in the ontologies. By switching from "*UK police ontologies*" to "*Dutch police ontologies*", we could reconfigure the system for operating in a different country, without changing a single line of code.

Using ontologies for making assumptions explicit helps the user understanding, and improves maintainability of the system. If what we mean with different concepts and relations is only implicitly encoded within the program code, there is no straight-forward way to, for instance, generate explanations for the user, why something was classified in a certain way, or why a certain data element violated a constraint. When using ontologies, all definitions and constraints are made explicit in the model, which can then be used to explain the inferences made. This does not only apply to the case when inferences are needed but also when querying the data it is useful to avoid encoding important distinctions in the query, but rather making those distinctions explicit in the ontology instead. This allows for simpler queries for retrieving data, and improves maintainability. Less of the knowledge is implicit in queries and program code, but rather explicit in the ontology - the programmer who tries to extend the system does not have to know what someone meant, and what implicit assumptions were made when writing a certain query in the code.

An example use of ontology in VALCRI is for crime classification. Police analysts do not only work with the legal categories of crimes, but rather have their own ways of grouping crimes, partly depending on the task they are working on. While such grouping can very well be done by clustering or other machine learning algorithms, such methods are not usually able to put a label on what the group actually means nor explain based on what criteria that the grouping is being made. Ontologies on the contrary provide a rule-based approach to this problem, where membership criteria and meaning of the groups can be made explicit in the ontology, and general-purpose inference engines can be used to perform the classification of new crime instances.

Another interesting feature of ontologies is its ability to express a hierarchical view on the data, i.e., concepts can be arranged in hierarchies. An example could be that a specific crime report describes a theft from a motor vehicle as involving the offender stealing the car satnav. In another crime report there might be a description of a car radio being stolen. While superficially, and textually, these descriptions may not be very similar, if the ontology includes the information that both satnav and radio are actually part of the console, then it suddenly emerges that these crimes may actually be related and belong to some series of "console thefts" in the area.

ONTOLOGY ENGINEERING IN VALCRI

In order to achieve the effects mentioned in the previous section, however, the ontologies of the system must be designed, stored and used in an appropriate way. First of all it is important to structure the ontologies in a modular way, so that certain parts can be replaced as needed. In VALCRI we have identified three levels of ontologies.

At the top level we have ontologies, and patterns (explained further below), that are considered as more or less fixed, i.e. both static over a longer period of time, and generic enough to hold in most contexts where the system is envisioned to be applied. This level consists of both generic Ontology Design Patterns (Blomqvist & Sandkuhl, 2005) (Gangemi, Ontology Design Patterns for Semantic Web Content, 2005) (Gangemi & Presutti, Ontology Design Patterns , 2009) as well as a domain ontology of the criminal intelligence domain. Examples of concepts on this level include things such as crimes, investigations, suspects and offenders, as well as crime patterns such as the notion of participations, i.e. that an actor can participate in some event, such as a person in a crime.

On the second level we have context-specific specialisations of the previous level. This is for instance anything that is specific to a certain setting where the system would be applied, such as legal concepts and consequences of the local law, and specific data elements that are recorded in the local data stores in that context. Examples of such concepts include legal classifications of crime types in a certain legal jurisdiction, as well as the specific attributes (data elements) collected in existing databases of the police organisation using VALCRI.

Finally, there is a **third level of ontologies** which are local even within a deployment context, i.e. can be personal to a certain user, or a group of users. Here the user can make additions to the two previous levels, e.g. add new classes of crimes that he or she is specifically targeting. While the latter level is only created while the system is being used, the two upper levels have to be created before the system is deployed. Our intention is that the top level only has to be created once; hence, our work in this project should be entirely reusable, while the middle level has to be developed for each new (deployment) context of the system. Therefore we also work with support for ontology engineering in this project, in order to support the development of new such "middle level" ontologies in the system configuration process before deploying it, as well as adding specialisations on the third, local, level.

Key facilitators for this ontology architecture are the notions of *ontology network* and *Ontology Design Patterns*. An **ontology network** was defined in the context of the *NeOn* project (Suárez-Figueroa, Gómez-Pérez, Motta, & Gangemi, 2012) and is a set of aligned (interlinked) ontologies that together can solve some set of requirements. However, the intention is also that the modules of the network, i.e. the individual ontologies, can be used standalone without the rest of the network if required. *Ontology Design Patterns* (**ODPs**) (Blomqvist & Sandkuhl, 2005) (Gangemi, Ontology Design Patterns for Semantic Web Content, 2005) (Gangemi & Presutti, Ontology Design Patterns , 2009) are on one hand comparable to design patterns in other areas, such as software, in that they represent generic solutions to common problems. On the other hand ODPs can also be used as small components, e.g. like pieces of a puzzle that act as readily available building blocks for composing the ontologies. In order to produce ontologies using ODPs we are applying a tailored version of the *eXtreme Design* (**XD**) ontology engineering methodology (Blomqvist , Hammar, & Presutti, 2016), together with the *WebProtégé* ontology engineering environment (extended with a plugin for using ODPs (Hammar, 2015). The aim is to develop this methodology, and associated tools (including the ontology evolution support described in the next section), to the level where technical staff configuring and deploying the VALCRI system (i.e. non-ontology experts) can create the and maintain the ontologies needed, based on specialising the "fixed" top level ontologies and patterns.

AN EVOLVING KNOWLEDGE BASE

In a system like VALCRI, data will be constantly changing. Some data changes slowly, such as geographical data about places and categories of crimes defined in legal regulations, while other data changes rapidly, such as near real-time reports of cars passing some CCTV camera and being recognised through automatic number plate recognition (ANPR) software, or financial transactions. To accommodate both these scenarios, VALCRI both needs to contain methods for ingesting new data entries from legacy databases, such as when a street name changes, or when a new crime is reported, and manage the consequences of such changes, as well as a stream processing component for filtering and processing rapid data streams in near real-time. While the former is mainly a matter of data transformation, mapping the data to the VALCRI ontologies and versioning, the latter is an even more challenging task, mainly due to scalability requirements. In VALCRI we are applying novel techniques for RDF stream processing (RSP) and Semantic Complex Event Processing (Semantic CEP) for performing such online filtering and processing. A demo scenario was developed to showcase some of the potential of the RSP component of the system (Keskisärkkä & Blomqvist, 2015).

However, it is not only the data that will change, also the models, ontologies, will have to change over time. For instance, new crime categories emerge over time and others become obsolete, new methods for conducting various crime types emerge and others disappear. For example, during recent years car thefts are increasingly based on "hacking" the electronics of the car, e.g. the electronic key for instance, which was a method not applicable prior to the modern cars that have a "keyless" unlock and start procedure. Although these crimes may still be sorted under the same legal classification, e.g. "theft of motor vehicle", for investigation purposes it may be beneficial to have a sub-category of the legal classification, representing this particular type of crime. By describing this group of crimes in a formal concept definition in an ontology, one can use the ontology for finding additional crimes reported that can be classified under this category, and analyse their attributes in order to create profiles of this type of crime. Depending on the work processes and restrictions set for the system, one may either add such new categories on the middle level, i.e. the context-specific ontologies, or allow them only on the individual (or group) level, e.g. as tools for a certain individual or group using the system. The intention is that concepts can be added, modified or deleted from the ontologies on the lower levels through an ontology editing interface, which can be viewed as part of an "administrator interface" to the VALCRI system. Changes to the ontologies could then be analysed and visualised in various ways (Lambrix et al., 2016).

Another feature of the VALCRI system is to be able to automatically suggest such editing actions, e.g. to propose new concepts, modifications of existing ones, or even deletion of obsolete ones. This is envisioned as a semi-automatic ontology evolution process, where a system administrator validates suggestions. Suggestions can be based on analysing a set of "signals", including cues from the information extraction and data ingestion processes, and user actions. For instance, terms and relations in textual data that cannot be mapped to any existing structure in the ontologies can be cues to something being missing in the ontologies, while parts of the ontologies that are rarely used by the system, and which do not have any recent instances added in the data may be candidates for modification or deletion.

CONCEPT EXTRACTION AND SIMILARITY ANALYSIS

Based on end-user feedback and analysis of typical applications in intelligence data analysis domain, it became clear that intelligence analysts are interested in information referencing a common topic and not every piece of potential useful information that can be found in a text document, for example witness statements or the modus operandi of a burglary. This observation is the motivation of what is called **concept extraction** in VALRCI. It is important to note that the term *concept* is not referring to a concept of an ontology. Instead, a concept refers to a general term that describes a set of words that build a semantically coherent group, which are used as a topic specific dictionary. While state of the art information extraction techniques can give hints on relevant terms or topics (Group, 2004), concept extraction is task- and application dependent, as the input in form of the concept dictionaries can be defined and selected by the intelligence data analyst. This adds flexibility and possibilities for task and data-based adaptions to the information extraction process of VALCRI.

The process of concept extraction has three phases. At first, the text data has to be processed that it can be used by automated data analysis methods. This can be done by state of the art pre-processing and text processing pipelines, as provided by Stanford CoreNLP or similar toolkits. Afterwards, the generated tokens should be lemmatised or stemmed, so that for the following processing problems caused by morphology or inflection of words is mitigated effectively. In phase two, a full text search of the words from the active concept lists in the text documents to process is executed.

Correspondingly to phase one, the concept words are required to be lemmatised or stemmed. The last step contains the creation of word contexts. For example, if a concept dictionary only contains nouns, a natural way of creating the context of a noun is to extract nouns plus proceeding adjectives, e.g. the match CAR, preceded by the adjective SILVER will be extended to the match of SILVER CAR. Similarly, different rule-sets for context creation can be created. Concept extraction generates sequences of words (n-grams) that are guaranteed to contain at least one concept word, plus additional context created by the last processing step.

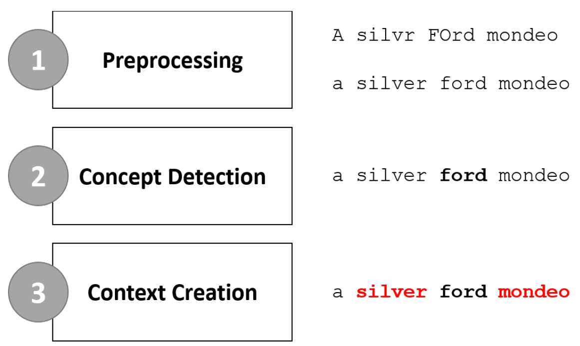

Figure 1: The three phases of Concept Extraction. Black denotes a concept term, red indicates context terms.

While the technology as outlined here may seem quite static, it gains application relevance by the choice of concepts to search for, the actual composition of concept dictionaries, and the context matches (concept terms and their context) that are extracted in the last step as described above. Each phase and its component can be parameterised to the needs of the analyst, pre-conditions or application requirements, as well as task-specific requirements. Additionally, these parameters stay on a non-technical level, so that analysts without specific training in information or linguistic technologies can adapt them to their needs.

Each component is building up on the outputs of preceding analysis steps. This allows the easy modification and extension of the pre-processing, concept word detection, and context creation if the data that is exchanged stays the same. This is especially true for methodological extensions, such as a fuzzy text search instead of exact searching of concept terms, machine learning to adapt or extend the concept dictionaries, or any human interaction with the corresponding automated data analysis processes. Not at last, the process can be parallelised easily, as there are no interdependencies of the concept extraction processes above document level, which makes this process suitable for interactive systems used by data analysts.

Many automatic tasks in the analysis of intelligence data, such as grouping, searching or ordering, require a numerical definition of **similarity**. While similarity of data that is of the same data type, such as text or numbers, is already well defined, the similarity analysis of crime reports requires a multi-type capable definition of similarity. This is due to the fact, that a single data record in intelligence domain is typically defined by a number of fields that are not of the same data type, e.g. text, numeric time stamps, categories, or simply numeric values. Therefore, the definition of a similarity needs to be able to unify different similarity scores, as state of the art techniques that identify groups/clusters or search for similar data records require the similarity of a given data record to another to be a single number. In consequence, the definition of similarity of two data records needs to expanded to the unified similarity of each of the same data fields of each data record.

For unification, the Gower Coefficient (Gower, 1971) can be used, which gives a general definition of a similarity that can be composed out of a number of different similarity scores. In addition, the Gower Coefficient allows that different components of the similarity are assigned different weights, which can be utilized to assign different importance to different fields of the data records. This translates well to the area of criminal intelligence analysis, where because of the analyst's experience, application or task requirements different fields of the data are of different importance. Using this technology, groups of crimes based on the time they happened can be created, that at the same time incorporate the similarity of the location assigned to a data record, by a customizable and potentially different importance (weight). Based on experiences with real-world data, the following parts of the final similarity score should be available: binary similarity of flags, similarity of dates, similarity of time, as well as similarity of basic data types such as text data (e.g. by cosine distances of larger documents or lexical similarity for shorter ones or keywords) and numeric values. These building blocks of a similarity measure allow a high degree of flexibility for adaption to the data or common practices of the analyst, e.g. by computing bins for data fields that correspond the time of day before the similarity is computed. Additionally, this approach is flexible enough to be extended in various ways, e.g. to compute the geographic distance of city districts for a data record with the corresponding field, instead of computing lexical similarity.

ASSOCIATIVE SEARCH

Criminal intelligence analysis is defined by INTERPOL (INTERPOL, 2016) as the *"… study [of] criminals, crime suspects, incidents, issues and trends ... [to] identify relationships or connections between different crimes in different places.*" Analysts face a number of significant difficulties in the information analysis process, including making sense of data from multiple sources that need to be collated and organized into meaningful ways that can lead to understanding and insight. Data is also of varying quality and reliability may be out of sequence, lacking context, and may be missing, ambiguous and uncertain. Analysts may also be working on several cases at the same time and may be difficult to distinguish the relevance or similarities among these cases. In our investigations of how analysts think and reason about problems in criminal intelligence analysis (Wong & Kodagoda, 2016), we have observed a practice we refer to as "*associative questioning*". This is the practice where an analyst asks a variety of questions to learn more about the diverse nature of the context in which the crimes were committed. During a crime matching process of assigning crimes or criminal to the previous solved or unsolved crimes, analysts often investigate clues to solve the crimes and hence seek to extract the hidden associations among connected objects during a crime investigation. The analyst might be interested in answering queries such as HOW a crime is related to another crime; WHO are the victims of this crime; WHAT are the other crimes that have been committed using the same Modus operandi; and many more similar questions. This kind of search is still in its infancy stage in data science domain and needs particular attention.

The data science fraternity in recent decade, have produce a great deal of scientific researches and studies on crime data analysis especially for entity extraction, crime matching, crime pattern detection (Nath, 2006) crime trends prediction (Chandra & Gupta, 2008), money laundering detection (K.Cao & Do, 2012), criminal career analysis (Bruin & Cocx, 2006) classification, cluster outlier and social network analysis to assist criminal analysis process. However not very much literature is available for link based search in criminal analysis.

In VALCRI we aim to give a data science based solution to this type of associative questioning. In this section we first define the notion of search functionality in VALCRI introducing a new term called as associative search. The intuition behind this relatively new type of search is to search across the different datasets to retrieve related meaningful *concepts* instead of simple keywords. For example in responding to query "*Tell me about Marks and spencer*" the system should retrieve information about the company, its director, associated business distribution channels and other meaningfully associated concepts rather than information that contains the keyword “Mark and spencer”.

Associative search in VALCRI is defined as the search along the networks of associations between objects such as people, places, other organizations, products, events, services. It is different from both keyword search and semantic based search techniques as described in previous sections. *Keyword based search* does not consider the meaning of the given query. It searches for exact words from related documents. It does not extract entity-specific information, association, or relationships from knowledge-base. *Semantic based search*, on other hand looks for the content that matches the meaning of the question, however it still does not actually leverage the power of associations of concepts in the search domain. Therefore there is a need for a search methodology particularly in criminal analysis to assist an analysist in making sense from the available data. This search mechanism inspired by the (Spivack, 2008) approach aims at developing functions that mimic our brain activities while searching for relevant information, through a web of associated documents, narrowing down the set of concepts and entities that are connected to all those starting points and result in various forms of associations to the input query.

The important issue in making association is to understand the user intention behind the query. Query Expansion, Relevance feedback and Pseudo-feedback are recommended methodologies quoted in literature to understand the user intention behind a query. A new algorithm for Query Expansion based on the Distance Constraint Activation model of human memory is recently been proposed by (Zhao & Luo, 2014). Another important issue is the establishment of the associations among the related concepts over the given data set. Researchers have demonstrated the use of ontology, probabilistic based models such as *Latent Dirichlet Allocation*, fuzzy-Based Methods and many others to find the relevancy among the semantic concept in a document. In addition to these the latest trend is towards the use of elastic search to find the relevancy among the text documents. The elastic search uses concept of vector space model and thus is based on standard similarity algorithm which is known as term frequency (TF) or inverse document frequency (IDF) therefore it does not necessarily depict the associative relations while searching. The search giant *Google* has implemented a knowledge graph to find the association between all the real world objects. The purpose of Knowledge Graph is to find the right information, get the best summary available and go deeper and broader into the content all in order to return more relevant information to the user and understand better what content they are looking. The social web site LinkedIn while applying this idea has patented relevance algorithm to calculate the relevance score, which in turn is used to establish the association among the people.

In VALCRI we take approaches based on exploratory data analysis and data mining techniques i.e. Formal Concept Analysis, partition clustering, frequent item set mining and A-Priori algorithm to determine the association rules between the linked concepts in the data.

Formal Concept Analysis represents the subject domain through a formal context made of objects and attributes of the concepts from the subject domain (Wille, 2009). A concept is constituted by its extension, comprising all objects belonging to this concept and its intension, comprising all attributes (properties, meanings) that apply to all objects of the extension. The set of objects and attributes, together with their relation to each other form a formal context, which can be represented by a cross table called as formal concept table. One of the major outputs of this cross table is a concept lattice representing formal concepts through nodes having attributes placed over the nodes and objects under the nodes. In order to handle information overload of associative search we have used this FCA approach and implemented it through five general though associated concepts i.e. WHAT (what has happened), WHO (who has committed the crime identifying offender or group of offenders), WHEN (When an offence was happened), WHERE (the geo-spatial information about the offence) HOW (The modus-operandi used in the committing a crime). Each of these concepts i.e. (WHO, WHEN, WHERE, WHAT and HOW) represent a question and are linked with each other through a set of properties or attributes. The objective is to find the connected information between offences (WHAT) AND offender (WHO) both temporally and spatially (WHEN & WHERE) generating association rules among all these concepts for a given set of Concepts representing a crime pattern. These questions are made through SQL queries form VALCRI dataset and later visualized in three separate tempo-spatial formal contexts i.e. Geographic and Temporal profile of the offender, Offender Network and Crime hotspot. We also have demonstrated the use of frequent item set mining and A-Priori algorithm to determine the association rules between the linked concepts in the data. This worked has been presented at IEEE SMC 2016 conference (Qazi, Wong, Kodagoda, & Adderley, 2016) we plan to incorporate in VALCRI in future.

In our second approach, we have employed the idea of partition clustering coupled with multi-dimensional scaling to visualize association to assist the crime matching process. This is implemented in the VALCRI following a pipeline consisting of *data filtering*, *dynamic clustering*, *Multi-dimensional scaling (MDS)*, *association extraction* and *visualisation*. This visualisation scheme creates a user defined two dimensional space consisting of key process indicators (KPI) such as total number of crimes, number of solved and unsolved crimes, offenders, similarity between the crimes reports, location of the offence represented through post-code, modus operandi of the offender such as entry position fixture type used, search location, entry type etc. For this purpose a dynamic feature vector based on the user specified concepts from the unstructured text is created and later fed in *K-mean*s clustering algorithm to generate clusters of solved, unsolved and associated offenders. The associations are extracted based on the graph theory and similarity of crime pattern. Three types of associations including associations between crimes as solved and unsolved, association between crimes and offenders generating a one degree criminal network identifying direct and indirect association, along with associations between offenders and offenders have been measured graphical theory and crime pattern similarity. This visualising framework is implemented in VALCRI as self-executing service which starts automatically after receiving a crime pattern consisting of modus operandi and other crime related information and generates the visualisation of the association among connected crime objects.

The visualizing Framework attempts to tackle the visualization of multi-dimensional yet associated crime variables in two layers of visualization. In the first layer we presents a two dimensional Cluster space visualizing multi-dimensional temporal and spatially associated crime variables such as crime similarity, solved and un-solved crimes, nominals, victim, suspects and offenders illustrating the association between solved and unsolved crime . The second level of the association between the solved crimes and associated offender revealing hot spots and Criminal profiles, temporal, geo-spatial and modus operandi based one degree network of the offenders in the form of hierarchal tree visualisation in a separate Card.

VIDEO ANALYSIS

For the video analysis tools of VALCRI we identified three major topics for automatic analysis that mainly aim at reducing the amount of time an (operational) analysis needs to invest when reviewing videos for case related content. These three tools are *video quality estimation*, *video motion estimation* and *video object detection*.

Video quality estimation gives a low level fine grained measure of different quality properties of the video such as black frames, freezes, over- or under exposure. The idea behind automatically measuring such data is that the analysist may decide to focus on HQ videos or parts of the video first in order to find what he is looking for.

The same idea is behind the motion estimation. We use a machine learning approach in order to classify global camera motion. In VALCRI we are aiming to train this model for specific camera motion that usually occurs, when something happens in a video. This might be caused due to the automatic motion sensing of the camera or the manual intervention of a camera operator. We assume that these patterns significantly differ from the usual motion of (surveillance) cameras. If this assumption is true we are able to identify and mark short time instances where the camera has moved in a non-standard way and thus lead to important points in the timeline of the video.

While the first two analysis and classification tasks are on a comparatively semantically low level, object detection classifies and annotates semantically higher concepts such as whole object categories. In the past 20 years a large and growing amount of literature describing techniques for automatic object detection has been published. Object detection aims at determining the locations and sizes of general objects such as Faces, and Persons. Automatic approaches for object detection often build the basis of subsequent tasks such as object tracking, object classification or individual identification.

Although object recognition, i.e. automatic classification of detected objects into predefined categories (e.g. car, house, person), is often referred to as a separate field of research, object localization and object classification are in fact very much related. One could argue that for object recognition in a given image one has to detect all objects (a restricted class of objects depending on the dataset), localize them with a bounding box and categorize that bounding box with a label. For object detection, however, one only has to differentiate between two classes (usually the object of interest, e.g. persons, versus the background), and estimate bounding boxes around each object of interest. However, the task of object classification then becomes trivial: One just has to train multiple object detectors - one for each object category - and classification can then just be done by evaluating which detector fires at a certain position. In modern deep learning-based algorithms, however, multiple object classes are trained at once and thus object detection and classification can be performed by one single unified framework.

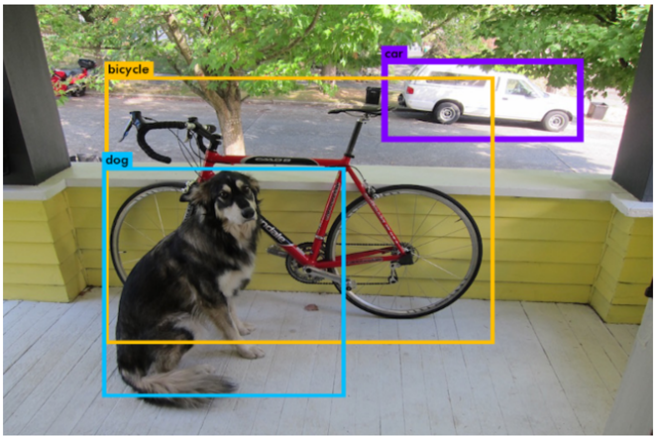

Figure 2: Object detection results of different object categories by YOLO

Figure 1 shows some examples for object detection by means of the deep-learning framework YOLO (Redmon, Divvala, Girshick, & Farhadi, 2016).

We aim at applying a state-of-the-art approach for object detection and classification using deep learning techniques. Recently developed deep learning techniques have revolutionized the field of computer vision and particularly object detection and classification. Especially Convolutional Neural Networks (CNN) have gained much attention in the computer vision community and have thus been used for a variety of different image and video processing tasks such as object detection and recognition as well as image classification among others. CNNs are biologically-inspired variants of Multi-Layer-Perceptrons (MLPs). Thus, a CNN is a type of feed-forward artificial neural network in which the connectivity pattern between its neurons is inspired by the organization of the visual cortex of mammals, whose individual neurons are arranged in such a way that they respond to overlapping regions tiling the visual field.

One of the first attempts to utilize CNNs for object detection called *Regions with CNN features (R-CNN)* was published by Girshick *et al.* in 2014 (Girschick, Donahue, Darrell, & Malik, 2014). The CNN localization problem is solved by operating within the ”recognition using regions” paradigm, which means that the input images is divided into *N x N* overlapping blocks and for each region the system estimates the probability that it belongs to a certain class.The proposed object detection system consists of three modules. The system (1) takes an input image, (2) extracts around 2000 bottom-up region proposals over several scales, (3) computes features for each proposal using a large CNN, and then (4) classifies each region using class-specific linear SVMs. The authors proposed a simple and scalable detection algorithm that improved mean average precision (mAP) by more than 30% relative to the previous best result on the PASCAL Visual Object Classes (VOC) Challenge, a widely known competition for object recognition algorithms. Although the proposed R-CNN algorithm achieved excellent detection accuracies, notable drawbacks can be found. Most notably, R-CNN is slow both in training and test time because it performs a CNN forward pass for each object proposal, without sharing computation. Thus, a number of extensions have been proposed in the recent past to overcome these limitations. *Fast RCNN* (Girshick, 2015) for instance employs several innovations to improve training and testing speed while at the same time increasing detection accuracy. While R-CNN takes about 47 seconds per image on a GPU, Fast R-CNN is about 213 times faster. Speed-Up is achieved by replacing the SVM with fully connected layers. Thus the architecture can be trained end-to-end. Furthermore, fully connected layers are accelerated by compressing them with a truncated version of the Singular Value Decomposition (SVD). Fast R-CNN achieves near real-time rates using very deep networks *when ignoring the time spent on region proposals* which actually turned out to be the bottleneck at test-time.

To overcome this issue, *Faster R-CNN* was proposed in 2015 by Ren *et al.* in (Ren, He, Girshick, & Sun, 2015). The authors introduce a Region Proposal Network (RPN) that shares full-image convolutional features with the detection network, thus enabling nearly cost-free region proposals. An RPN is a fully convolutional network that predicts object bounds and object scores at each position. The RPN is trained end-to-end to generate high-quality region proposals, which are used by Fast R-CNN for detection. Furthermore, RPN and Fast R-CNN are fused into a single network by sharing their convolutional features. The proposed detection system has a frame rate of 5 frames per second (including all steps) on a GPU, while achieving state-of-the-art results for object detection.

In the same year, a unified, real-time capable object detection framework named *YOLO (You Only Look Once)* was presented in (Redmon, Divvala, Girshick, & Farhadi, 2016). Prior work on object detection repurposes classifiers to perform detection. Instead, within YOLO object detection is framed as a regression problem to spatially separate bounding boxes and associated class probabilities. A single neural network predicts bounding boxes and class probabilities directly from full images in one evaluation. Since the whole detection pipeline is a single network, it can be optimized end-to-end directly on detection performance. Compared to previous deep-learning frameworks for object detection such as R-CNN and its variants, YOLO is extremely fast. The base YOLO model processes images in real-time at 45 frames per second on a GPU. A smaller version of the network, Fast YOLO, processes an astounding 155 frames per second while still achieving double the mean average precision of other real-time detectors. Compared to other state-of-the-art detection systems, YOLO makes more localization errors but is less likely to predict false positives on background. (Redmon, Divvala, Girshick, & Farhadi, 2016). Due to the impressive speed of YOLO and the state-of-the-art results achieved on extreme challenging datasets we will apply, re-train and adapt YOLO for the task at hand within this project.

Especially in video, where the occurrence of such objects of interest is usually low the automatic annotations again may help to find important parts in videos in less time than required for completely going over the videos.

All these automatic classification results will be designed with care. Since automatic approaches are prone to errors we will take care that their results will always be clearly identifiable as results of an automatic tool. We will also ensure, that all video data and analyst is allowed to watch is always completely viewable. None of the automatic annotations will obscure or hide videos or parts of videos but rather use visualization techniques such as heat maps in order to communicate their results.

REFERENCES

Blomqvist , E., Hammar, K., & Presutti, V. (2016). Engineering Ontologies with Patterns – The eXtreme Design Methodology. *Ontology Engineering win Ontology Design Pattern - Foundations and APplications.* IOS press.

Blomqvist, E., & Sandkuhl, K. (2005). Patterns in Ontology Engineering: Classification of Ontology Patterns. *Proc. of ICEIS’05, Miami Beach, Florida.*

Bruin, D., & Cocx, T. (2006). Data Mining Approaches to Criminal Career Analysis. *Sixth Int. Conf. Data Min*, (S. 171–177,).

Chandra, B., & Gupta, M. (2008). A multivariate time series clustering approach for crime trends prediction. *IEEE Int. Conf. Syst. Man Cybern*, (S. 892–896).

Gangemi, A. (2005). Ontology Design Patterns for Semantic Web Content. *The Semantic Web - ISWC 2005, LNCS Vol.3729.* Springer.

Gangemi, A., & Presutti, V. (2009). Ontology Design Patterns . In *In Handbook on Ontologies, 2nd Ed. Int Handbooks on Information Systems.* Springer.

Girschick, R., Donahue, J., Darrell, T., & Malik, J. (2014). Rich feature hierarchies for accurate object detection and semantic segmentation. *Computer Vision and Pattern Recognition (CVPR).*

Girshick, R. (2015). Fast R-CNN. *International Conference on Computer Vision (ICCV).*

Gower, J. C. (1971). A General Coefficient of Similarity and Some of Its Properties. *Biometrics*, 857-871.

Group, O. W. (2004). *OWL Web Ontology Language Reference.* https://www.w3.org/TR/owl-ref/, last retrieved 25/11/2016.

Gruber, T. (1993). A translation approach to portable ontology specifications. *Knowledge Acquisition, 5(2)*, 199-220.

Hammar, K. (2015). Ontology Design Patterns in WebProtégé. . *In: Proceedings of the ISWC 2015 Posters & Demonstrations Track co-located with the 14th International Semantic Web Conference (ISWC-2015).* Betlehem, USA : CEUR-WS.

INTERPOL. (2016). *Criminal Intelligence Analysis.* http://www.interpol.int/INTERPOL-expertise/Criminal-Intelligence-Analysis.

K.Cao, D., & Do, P. (2012). Applying Data Mining in Money Laundering Detection. *ACIIDS'12 Proceedings of the 4th Asian conference on Intelligent Information and Database Systems*, (S. 207–216).

Nath, S. V. (2006). Crime Pattern Detection Using Data Mining. *Florida Atlantic University Oracle Corporation*, vol. 1, no. 954, pp. 1–4.

Qazi, N., Wong, W., Kodagoda, N., & Adderley, R. (2016). Associative Searchthrough Formal Concept Analysis in Criminal Intelligence Analysisu. *IEEE International Conference on Systems, Man, and Cybernetics SMC 2016.*

Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (2016). You Only Look Once: Unified, Real-Time Object Detection. *IEEE Conference On Computer Vision And Pattern Recognition (CVPR)* (pp. -). Las Vegas: IEEE.

Ren, S., He, K., Girshick, R., & Sun, J. (2015). Faster R-CNN: Towards Real-Time Object Detection. *Advances in Neural Information Processing Systems (NIPS)*.

Spivack, N. (2008). *Associative Search and the Semantic Web: The Next Step Beyond Natural Language Search.* Reterieved from http://www.novaspivack.com/technology/associative-search-and-the-semantic-web-the-next-step-beyond-natural-language-search.

Suárez-Figueroa, M., Gómez-Pérez, A., Motta, E., & Gangemi, A. (2012). *Ontology Engineering in a Networked World.* Springer.

Wille, R. (1982). Restructuring lattice theory: An approach based on hierarchies of concepts. *ICFCA'09, 7th International Conference on Formal Concept Analysis*, (S. 314). Berlin, Heidelberg.

Wille, R. (1982). Restructuring lattice theory: An approach based on hierarchies of concepts. *In: Rival, I. (ed.) Ordered Sets* , (S. 445-470).

Wille, R. (2009). Restructuring lattice theory: An approach based on hierarchies of concepts. *ICFCA'09, 7th International Conference on Formal Concept Analysis*, (S. 314). Berlin, Heidelberg.

Wille, R. (2009). Restructuring lattice theory: An approach based on hierarchies of concepts. *ICFCA'09, 7th International Conference on Formal Concept Analysis, (S. 314).* Berlin, Heidelberg.

Wong, W., & Kodagoda, N. (2016). How analysts think: Anchoring, Laddering and Associations . *Proceedings of the Human Factors and Ergonomics Society 60th Annual Meeting.* Washington, D.C., USA : SAGE Publications.

Zhao, X., & Luo, X. (2014). Query Expansion Based on the Distance Constraint Activation of Human Memory. *Ninth International Conference on Semantics, Knowledge and Grids*, (S. 143-150). eijing.

The research leading to the results reported here has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) through Project VALCRI, European Commission Grant Agreement Number FP7-IP-608142, awarded to Middlesex University and partners.

	VALCRI Partners	Country
1	Middlesex University London Professor B.L. William Wong, Project Coordinator Professor Ifan Shepherd, Deputy Project Coordinator	United Kingdom
2	Space Applications Services NV Mr Rani Pinchuck	Belgium
3	Universitat Konstanz Professor Daniel Keim	Germany
4	Linkopings Universitet Professor Henrik Eriksson	Sweden
5	City University of London Professor Jason Dykes	United Kingdom
6	Katholieke Universiteit Leuven Professor Frank Verbruggen	Belgium
7	A E Solutions (BI) Limited Dr Rick Adderley	United Kingdom
8	Technische Universitaet Graz Professor Dietrich Albert	Austria
9	Fraunhofer-Gesellschaft Zur Foerderung Der Angewandten Forschung E.V. Mr. Patrick Aichroft	Germany
10	Technische Universitaet Wien Assoc. Prof. Margit Pohl	Austria
11	ObjectSecurity Ltd Mr Rudolf Schriener	United Kingdom
12	Unabhaengiges Landeszentrum fuer Datenschutz Dr Marit Hansen	Germany
13	i-Intelligence Mr Chris Pallaris	Switzerland
14	Exipple Studio SL Mr German Leon	Spain
15	Lokale Politie Antwerpen	Belgium
16	Belgian Federal Police	Belgium
17	West Midlands Police	United Kingdom