

Towards a Forensic Event Ontology to Assist Video Surveillance-based Vandalism Detection

Faranak Sobhani¹ and Umberto Straccia²

¹Queen Mary University of London, UK

²ISTI-CNR, Italy

March 22, 2019

Abstract

The detection and representation of events is a critical element in automated surveillance systems. We present here an ontology for representing complex semantic events to assist video surveillance-based vandalism detection. The ontology contains the definition of a rich and articulated event vocabulary that is aimed at aiding forensic analysis to objectively identify and represent complex events. Our ontology has then been applied in the context of London Riots, which took place in 2011. We report also on the experiments conducted to support the classification of complex criminal events from video data.

1 Introduction

In the context of vandalism and terrorist activities, video surveillance forms an integral part of any incident investigation and, thus, there is a critical need for developing an “automated video surveillance system” with the capability of detecting complex events to aid the forensic investigators in solving the criminal cases. As an example, in the aftermath of the London riots in August 2011 police had to scour through more than 200,000 hours of CCTV videos to identify suspects. Around 5,000 offenders were found by trawling through the footage, after a process that took more than five months.

With the aim to develop an open and expandable video analysis framework equipped with tools for analysing, recognising, extracting and classifying events in video, which can be used for searching during investigations with unpredictable characteristics, or exploring normative (or abnormal) behaviours, several efforts for standardising event representation from surveillance footage have been made [9, 10, 11, 22, 23, 28, 30, 37].

While various approaches have relied on offering foundational support for the domain ontology extension, to the best of our knowledge, a systematic ontology for standardising the event vocabulary for forensic analysis and an application of it has not been presented in the literature so far.

In this paper, we present an OWL 2 [25] ontology for the semantic retrieval of complex events to aid video surveillance-based vandalism detection. Specifically, the ontology is a derivative of the DOLCE foundational ontology [7] aimed to represent events that forensic analysts commonly

encounter to aid in the investigation of criminal activities. The systematic categorisation of a large number of events aligned with the philosophical and linguistic theories enables the ontology for interoperability between surveillance systems. We also report on the experiments we conducted with the developed ontology to support the (semi-) automatic classification of complex criminal events from semantically annotated video data.

Our work significantly extends the preliminary works [12, 31]. The work [12] is an embryonal work investigating about the use of an ontology for automated visual surveillance systems, which then has been then further developed in [31]. While our work shares with [31] some basic principles in the development of the ontology, here the level of details is now higher (e.g., the **Endurant** class (see Section 4.2.2 and its sub-classes have not been addressed in [31]) and various ontological errors have been revised. Additionally, and somewhat more important, in our work experiments have been conducted for criminal event classification based on London 2011 riots videos. Furthermore, but less related, is [32] in which the technical challenges facing researchers in developing computer vision techniques to process street-scene videos are addressed. The work focusses on standard image processing methods and does not deal with ontologies in any way.

The remainder of the paper is organised as follows. Related work is addressed in Section 2. For the sake of completeness, Section 3 presents complementary, theoretical material to ease the understanding of the ontology expressions used in our work. Section 4 presents a detailed description of the forensic ontology about complex criminal events. In Section 5 we discuss how to use the ontology to assist video surveillance-based vandalism detection. In Section 6 we conduct some experiments with our ontology based on CCTV footage of London riots from 2011, and finally, Section 7 concludes.

2 Related Work

In [23], the *Event Recognition Language* (ERL) is presented, which can describe hierarchical representation of complex spatiotemporal and logical events. The proposed event structure consists of primitive, single-thread, and multi thread events. Another event representation ontology, called CASE^E, is based on natural language representation and is proposed in [11] and then extended in [10]. Subsequently, in [9, 22] a *Video Event Representation Language* (VERL) was proposed for describing an ontology of events and the companion language called *Video Event Markup Language* (VEML), which is a representation language for describing events in video sequences based on OWL [21]. In [30], event detection is performed using a set of rules using the SWRL language [24].

The *Event Model E* [37] has been developed based on an analysis and abstraction of events in various domains such as research publications, personal media [1], meetings [13], enterprise collaboration [14] and sports [26]. The framework provides a generic structure for the definition of events and is extensible to the requirements ontology of events in the most different concrete applications and domains.

In [28] a formal model of events is presented, called *Event-Model-F*. The model is based on the foundational ontology DOLCE+DnS Ultralite (DUL) and provides comprehensive support to represent time and space, objects and persons, as well as mereological, casual, and correlative relationships between events. In addition, the Event-Model-F provides a flexible means for event composition, modelling event causality and event correlation, and representing different interpretations of the same event. The Event-Model-F is developed following the pattern-oriented approach of DUL, is modularised in different ontologies, and can be easily extended by domain specific ontologies.

While the above-mentioned approaches essentially provide frameworks for the representation of events, none of them address the problem of formalising forensic events in terms of a standard representation language such as OWL 2¹ and, importantly, none have been applied and tested so far in a real use case, which are the topics of the following sections.

3 Description Logics Basics

For the sake of completeness of this work, we recap here succinctly some basic notions related to the *Description Logics* (DLs) family of languages, which is the logical counterpart of the *Ontology Web Language* (OWL 2) [25]. We will in fact use DL expressions to explain the ingredients of the forensic event ontology, which we will use in the following sections. We refer the reader to e.g. [2, 3] for further insights on DLs.

Roughly, DLs allow to describe classes (also called concepts), properties of classes and relationships among classes and properties. Formally, we start with the basic DL called \mathcal{ALC} (*Attributive Language with Complement*) [29]. Elementary descriptions are *atomic concepts*, also called *concept names* (denoted A) and *atomic roles* (denoted r). Complex *concepts* (denoted C) can be built from them inductively with concept constructors. Specifically, concepts in \mathcal{ALC} are formed according to the following syntax rule:

C, D	\rightarrow	A		(atomic concept)
		\top		(universal concept)
		\perp		(bottom concept)
		$C \sqcap D$		(concept conjunction)
		$C \sqcup D$		(concept disjunction)
		$\neg C$		(concept negation)
		$\forall r.C$		(universal restriction)
		$\exists r.C$		(qualified existential restriction) .

Sometimes, the abbreviation $\exists R$ is used in place of the unqualified existential restriction $\exists R.\top$.

We further extend \mathcal{ALC} by allowing both (i) *complex roles* (or, simply *roles*) defined inductively from atomic roles as follows (r, s are roles):

- an atomic role is a role;
- r^- is a role and denotes the inverse of role r ;
- $r \circ s$ is a role and denotes the composition of role r and s ;

and (ii) *existential value restrictions*, which are concepts of the form $\exists r.\{a\}$, where a is an individual.

We recall that, informally, a concept denotes a set of objects, while a role is a binary relation over objects. So, for instance, the concept $\text{Person} \sqcap \exists \text{hasChild.Femal}$ will denote people having a female child, while $\text{Person} \sqcap \exists \text{hasFriend}\{\text{mary}\}$ will denote people having mary as friend.

A *knowledge base* \mathcal{K} consists of a finite set of axioms. An *axiom* may be a

- *General Concept Inclusion axiom* (GCI) $C \sqsubseteq D$, where C and D are concepts (read it as “all instances of C are instances of D ”). An example of GCI is $\text{Male} \sqsubseteq \neg \text{Female}$.

¹We recall that the relationship to our previous work [12, 31, 32] has been addressed in the introductory section

- *concept and role assertion axiom* $C(a)$ and $r(a, b)$, respectively, where a and b are individuals. Examples of assertion axioms are `Person(tim)` (“tim is a person”) and `hasChild(tim, pat)` (“tim has pat as child”).
- *Role Inclusion axiom* (RI) $r \sqsubseteq s$, where r and s are roles (read as “all instances of role r are instances of role s ”). An example of RI is `hasPart ◦ hasPart ⊆ hasPart`, which declares role `hasPart` as transitive.

From a semantics point of view, an *interpretation* \mathcal{I} is a pair $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ consisting of a non-empty set $\Delta^{\mathcal{I}}$ (called the *domain*) and of an *interpretation function* $\cdot^{\mathcal{I}}$ that assigns to each atomic concept a subset of $\Delta^{\mathcal{I}}$, to each role a subset of $\Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$ and to each individual a an element in $\Delta^{\mathcal{I}}$. Eventually, the mapping $\cdot^{\mathcal{I}}$ is extended to complex roles as follows:

$$\begin{aligned} (r^-)^{\mathcal{I}} &= \{ \langle x, y \rangle \mid \langle y, x \rangle \in r^{\mathcal{I}} \} \\ (r \circ s)^{\mathcal{I}} &= \{ \langle x, y \rangle \mid \exists z \text{ s.t. } \langle x, z \rangle \in r^{\mathcal{I}} \text{ and } \langle z, y \rangle \in s^{\mathcal{I}} \} . \end{aligned}$$

The mapping $\cdot^{\mathcal{I}}$ is extended to complex concepts as follows:

$$\begin{aligned} \top^{\mathcal{I}} &= \Delta^{\mathcal{I}} \\ \perp^{\mathcal{I}} &= \emptyset \\ (C \sqcap D)^{\mathcal{I}} &= C^{\mathcal{I}} \cap D^{\mathcal{I}} \\ (C \sqcup D)^{\mathcal{I}} &= C^{\mathcal{I}} \cup D^{\mathcal{I}} \\ (\neg C)^{\mathcal{I}} &= \Delta^{\mathcal{I}} \setminus C^{\mathcal{I}} \\ (\forall r.C)^{\mathcal{I}} &= \{ x \in \Delta^{\mathcal{I}} \mid r^{\mathcal{I}}(x) \subseteq C^{\mathcal{I}} \} \\ (\exists r.C)^{\mathcal{I}} &= \{ x \in \Delta^{\mathcal{I}} \mid r^{\mathcal{I}}(x) \cap C^{\mathcal{I}} \neq \emptyset \} \\ (\exists r.\{a\})^{\mathcal{I}} &= \{ x \in \Delta^{\mathcal{I}} \mid r^{\mathcal{I}}(x) \cap \{a^{\mathcal{I}}\} \neq \emptyset \} , \end{aligned}$$

where $r^{\mathcal{I}}(x) = \{ y \mid \langle x, y \rangle \in r^{\mathcal{I}} \}$.

The *satisfiability* of an axiom E in an interpretation $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$, denoted $\mathcal{I} \models E$, is defined as follows: $\mathcal{I} \models C \sqsubseteq D$ iff $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$, $\mathcal{I} \models r \sqsubseteq s$ iff $r^{\mathcal{I}} \subseteq s^{\mathcal{I}}$, $\mathcal{I} \models C(a)$ iff $a^{\mathcal{I}} \in C^{\mathcal{I}}$, and $\mathcal{I} \models r(a, b)$ iff $\langle a^{\mathcal{I}}, b^{\mathcal{I}} \rangle \in r^{\mathcal{I}}$. Given a knowledge base \mathcal{K} , we say that \mathcal{I} *satisfies* \mathcal{K} iff \mathcal{I} satisfies each element in \mathcal{K} . If $\mathcal{I} \models \mathcal{K}$ we say that \mathcal{I} is a *model* of \mathcal{K} . An axiom E is a *logical consequence* of a knowledge base \mathcal{K} , denoted $\mathcal{K} \models E$, iff every model of \mathcal{K} satisfies E . Determining whether $\mathcal{K} \models C(a)$ is called the *instance checking problem*, while determining whether $\mathcal{K} \models C \sqsubseteq D$ is called the *subsumption problem*.

4 A Forensic Event Ontology

In the following, we present an OWL 2 ontology to support to some extent the semantic retrieval of complex events to aid automatic or semi-automatic video surveillance-based vandalism detection. The idea is to develop an ontology that not only conveys a shared vocabulary, but some inferences based on it may assist a human being to support the video analysis by hinting to videos that may be more relevant than others in the detection of criminal events.

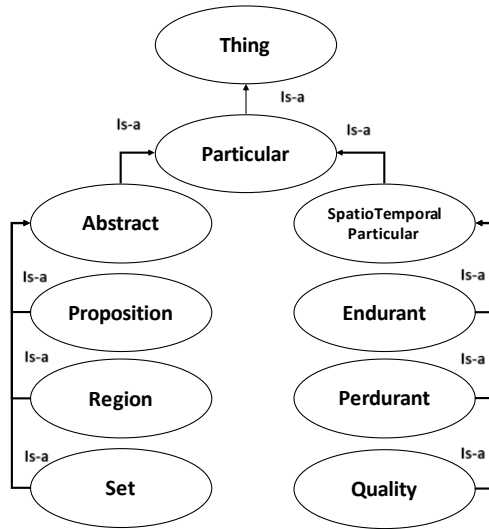


Figure 1: The foundational DOLCE ontology.

4.1 The Role of a Foundation Ontology

To facilitate the elimination of the terminological ambiguity and the understanding and interoperability among people and machines [19], it is common practice to consider a so-called *foundational ontology*. Let us note that several efforts have been taken by researchers in defining the foundational ontologies, such as BFO,² SUMO,³ UFO⁴ and DOLCE,⁵ to name a few. AS DOLCE ontology offers a cognitive bias with the ontological categories underlying natural language and human common sense, the same is selected for our proposed extension. We recall that the DOLCE foundational ontology (see Figure 1) encompasses **Endurant** and **Perdurant** entities. **Endurant** entities are ever-present at any time as opposed to **Perdurant** entities that extended in time by accumulating different temporal parts. A more thorough explanation on the DOLCE events conceptualisation can be found e.g. in [7].

4.2 A Forensic Complex Event Ontology

Our complex event classes extend DOLCE’s **Perdurant** class. To assign the action classes into respective categories, we follow a four-way classification of action-verbs: namely, into **State**, **Process**, **Achievement** and **Accomplishment** using event properties such as **telic**, **stage** and **cumulative** (see [27, 35, 36]). The distinction between these concepts are derived from the event properties as illustrated in Table 1, which we summarise below.

- **State [-telic,-stage]** This action category represents a long, non-dynamic event in which

²<http://ifomis.uni-saarland.de/bfo/>

³<http://www.adampease.org/OP/>

⁴<https://oxygen.informatik.tu-cottbus.de/drupal7/ufo/>

⁵<http://www.loa.istc.cnr.it/old/Papers/DOLCE2.1-FOL.pdf>

Table 1: Classification of Event Types.

State	-telic	-stage	cumulative
Process	-telic	+stage	-
Achievement	+telic	-stage	not cumulative
Accomplishment	+telic	+stage	not cumulative

every instance is the same: there cannot be any distinction made between the stages. States are cumulative and homogenous in nature.

- **Process** [-telic, +stage] The action category, like **State**, is atelic, but unlike **State**, the action undertaken are dynamic. The actions appear progressively and thus can be split into a set of stages for analysis.
- **Accomplishment** [+telic, +stage] Accomplishments are telic and cumulative activities, and thus behave differently from both **State** and **Process**. The performed action can be analysed in stages and in this way, they are similar to **Process**. Intuitively, an accomplishment is an activity which moves toward a finishing point as it has variously been called in the literature. Accomplishment is also cumulative activity.
- **Achievement** [-telic, -stage] Achievements are similar to **Accomplishment** in their telicity. They are also not cumulative with respect to contiguous events. Achievements do not go on or progress, because they are near instantaneous, and are over as soon as they have begun.

4.2.1 Forensic Perdurant Entities

Perdurant entities, extend in time by accumulating different temporal parts and some of their proper temporal parts may be not present. To this end, Perdurant entities are divided into the classes **Event** and **Stative**, classified according to their temporal characteristics.

Axiom sets (1)-(5) below provide a subset of our formal extension of the Perdurant vocabulary. The forensic extension of the ontology structure is shown in Figure 2.

$$\begin{aligned}
 \text{Perdurant} &\sqsubseteq \text{SpatioTemporalParticular} \\
 \text{Perdurant} &\sqsubseteq \exists \text{participant. Endurant} \\
 \text{Fighting} &\sqsubseteq \exists \text{participant. GroupOfPeople} \\
 \text{Perdurant} &\sqsubseteq \neg \text{Endurant} \\
 \text{Kicking} &\sqsubseteq \neg \text{Vehicle}
 \end{aligned} \tag{1}$$

$$\begin{aligned}
 \text{State} &\sqsubseteq \text{Stative} \\
 \text{MetaLevelEvent} &\sqsubseteq \text{State} \\
 \text{Accusing} &\sqsubseteq \text{MetaLevelEvent} \\
 \text{Beiliving} &\sqsubseteq \text{MetaLevelEvent} \\
 \text{PsychologicalAggression} &\sqsubseteq \text{State} \\
 \text{Blaming} &\sqsubseteq \text{PsychologicalAggression} \\
 \text{Bullying} &\sqsubseteq \text{PsychologicalAggression}
 \end{aligned} \tag{2}$$

$$\begin{array}{lcl}
\text{Process} & \sqsubseteq & \text{Stative} \\
\text{Action} & \sqsubseteq & \text{Process} \\
\text{Gesture} & \sqsubseteq & \text{Process} \\
\text{PhysicalAggression} & \sqsubseteq & \text{Process} \\
\text{ActivePhysicalAggression} & \sqsubseteq & \text{PhysicalAggression}
\end{array} \tag{3}$$

$$\begin{array}{lcl}
\text{Accomplishment} & \sqsubseteq & \text{Event} \\
\text{CriminalEvent} & \sqsubseteq & \text{Accomplishment} \\
\text{EventCategory} & \sqsubseteq & \text{Accomplishment} \\
\text{Crimecategory} & \sqsubseteq & \text{Stative}
\end{array} \tag{4}$$

$$\begin{array}{lcl}
\text{Achievement} & \sqsubseteq & \text{Event} \\
\text{Saying} & \sqsubseteq & \text{Achievement} \\
\text{Seeing} & \sqsubseteq & \text{Achievement} .
\end{array} \tag{5}$$

An excerpt of the forensic ontology is shown in Figure 3 and strictly adheres to the above terminological determination of action categories as mentioned before and extends the classes with suitable event concepts.

The concept **State** offers representation for **MetaLevelEvent** which encompasses abstract human events such as **Accusing**, **Believing** and **Liking** among others. As previously stated, the concept **State** represents a collection of events which are exhibited by a human that is time-consuming, non-dynamic, cumulative and homogenous. The other sub-class of **State** is **PsychologicalAggression** which characterises the human actions such as **Blaming**, **Decrying**, **Harassing** and so forth. The concept **Process** includes several human action categories that represent dynamic events which can be split into several intermediate stages for analysis. For the purposes of clarity, the concept **Process** offers three sub-concepts namely **Action**, **Gesture** and **PhysicalAggression**. The **Action** class incorporates different event such as **Dancing**, **Greeting**, **Hugging** among other concepts defined. The concept **Gesture** formalises the different interest points related to human gestures. In order to eliminate the ambiguity traditionally present in human gestures across cross-cultural impact, the action performed during the gesture is captured and represented in the ontology and, thus, enabling the removal of subjectivity from the concept definition. The final sub-class of the **Process** class includes the concept **PhysicalAggression** and formalises human conflicting actions.

By and large, the human action categorised into **State** and **Process** represent the microscopic movements of humans.

From the automatic surveillance viewpoint, these microscopic events may be extracted from media items. In contrast, the event representation formalised by means of the concepts **Achievement** and **Accomplishment** offer a rich combination of human events that allow for the construction of complex events with or without the combination of microscopic features. For instance, the concept hierarchy for **Vandalism** is illustrated in Figure 4, while the concept hierarchy for **CyberCrime** is shown in Figure 5 instead.

4.2.2 Forensic Endurant Entities

DOLCE is based on fundamental distinction among **Endurant** and **Perdurant** entities. The difference between **Endurant** and **Perdurant** entities is related to their behaviour in time. **Endurant** are wholly present at any time they are present. Philosophers believe that **endurant** are entities that are in

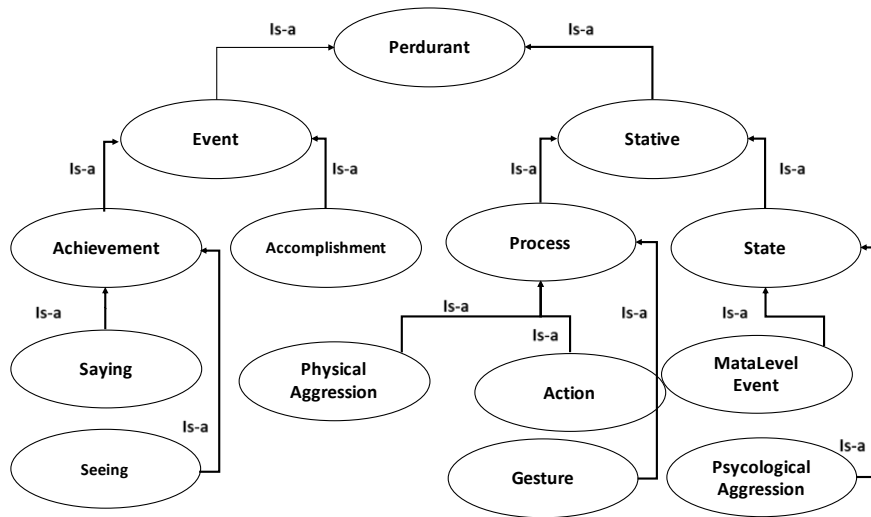


Figure 3: The Perdurant class hierarchy for forensic events descriptions.

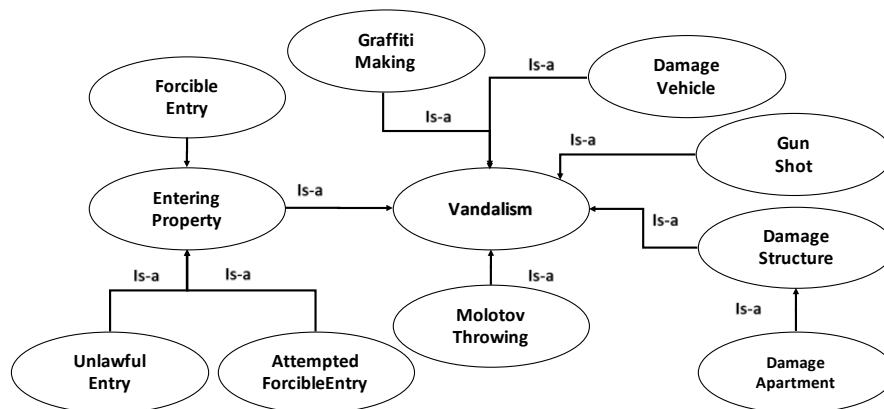


Figure 4: The concept hierarchy of Vandalism, direct subclass of CrimeAgainstProperty. The latter is a subclass of class Accomplishment.

time while lacking, however, temporal parts [19]. Therefore, the proposed vocabulary structure of all possible forensic entities also extends on Endurants entities.

Axiom set (6) describes a subset formalization of the Endurant vocabulary and an excerpt of the

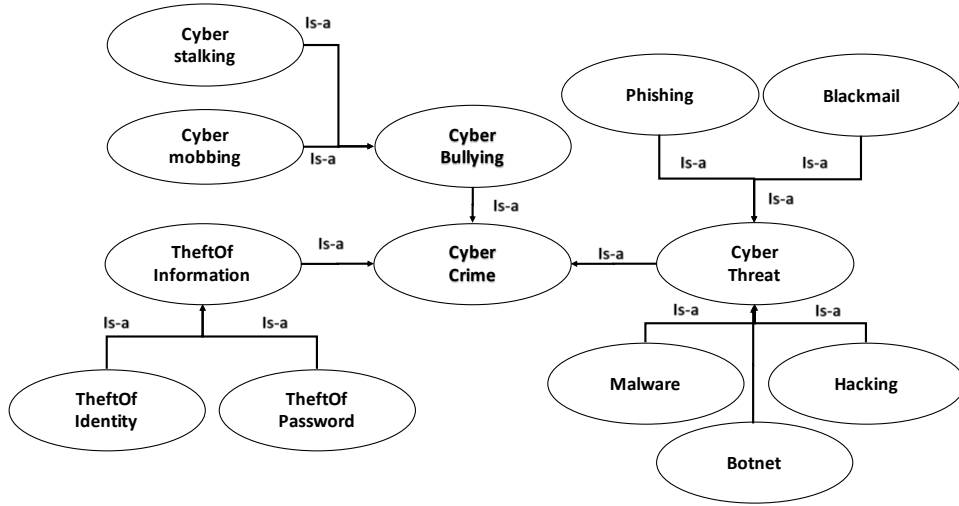


Figure 5: The concept hierarchy of CyberCrime.

forensic extension of the ontology structure shown in Figure 6.

$$\begin{aligned}
 \text{Endurant} &\sqsubseteq \text{SpatioTemporalParticular} \\
 \text{Endurant} &\sqsubseteq \exists \text{participantIn.Perdurant} \\
 \text{participantIn} &= \text{participant}^- \\
 \text{NonPhysicalEndurant} &\sqsubseteq \text{Endurant} \\
 \text{PhysicalEndurant} &\sqsubseteq \text{Endurant} \\
 \text{ArbitrarySum} &\sqsubseteq \text{Endurant} .
 \end{aligned} \tag{6}$$

5 Assisting Video Surveillance-based Vandalism Detection

We next show how the so far developed ontology is expected to be used to assist video surveillance-based vandalism detection.

5.1 Annotating Media Objects, viz. Surveillance Videos

Given surveillance videos and any media in general, we need a method to annotate them by using the terminology provided by our ontology. This gives rise to a set of facts that, together with the inferred facts, may support a more effective automatic or, more likely, semi-automatic retrieval of relevant information, such as e.g. vandalic acts. Specifically, the inferred information may suggest a user look at some e.g. video sequences or video still images, rather than to others first.

The general model we are inspired on is based on [20]. Conceptually, according to [20], a media object o (e.g., an image region, a video sequence, a piece of text, etc.) is annotated with one (or more) entities t of the ontology (see e.g. Figure 7).

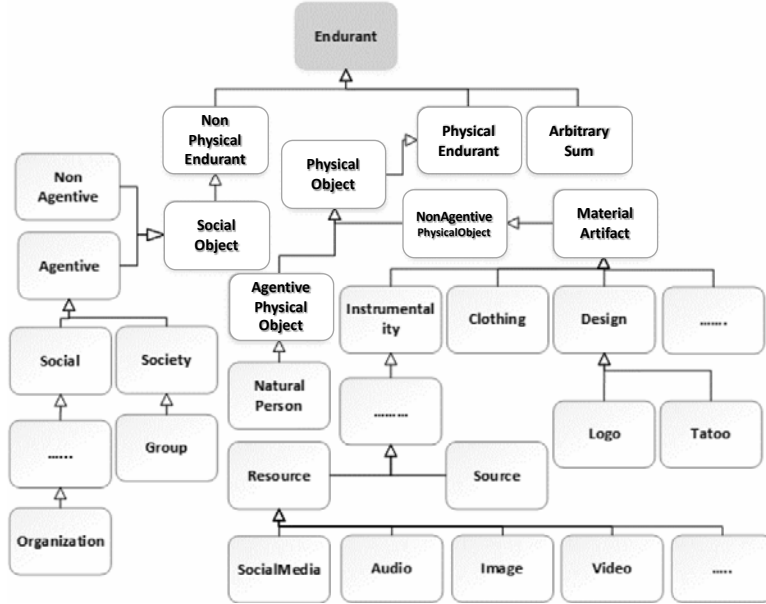


Figure 6: Excerpt of the Endurant concept hierarchy in the forensic ontology.

For instance, stating that an image object o is about a `DamageVehicle` can be represented conceptually via the DL expression

$$(\exists \text{isAbout}.\text{DamageVehicle})(o) .$$

As specified in [20], such an annotation may come manually from a user or, if, available, from an image classifier. In the latter case, it may annotate the image automatically, or, semi-automatically by suggesting to a human *annotator*, which are the most relevant entities of the ontology that may be used for a specific media object o . Note, however, that, the above methodology just illustrates the concept. In our case, for the sake of ease the annotation, we may not enforce the use of the object property `isAbout` (see Example 2 later on). Generally, we will annotate a `Resource` with `Perdurants` and `Endurants`: thus, if an image is annotated with e.g. a perdurant that is a damaged vehicle, then this means that the image *is about* a damaged vehicle.

We recall that `Resources` (and `Sources`) are modelled as follows (see Figure 8):

$$\begin{aligned} \text{Source} &\sqsubseteq \text{Endurant} \sqcap \exists \text{has}.\text{Resource} \\ &\quad \sqcap \exists \text{hasCameraId}.\text{string} \\ &\quad \sqcap \exists \text{hasLatitude}.\text{string} \\ &\quad \sqcap \exists \text{hasLongitude}.\text{string} \\ &\quad \sqcap \exists \text{hasLocationName}.\text{string} \\ \text{Resource} &\sqsubseteq \text{Endurant} \sqcap \exists \text{has}.\text{Perdurant} \\ \text{has} &= \text{isFrom}^- \\ \text{has} \circ \text{has} &\sqsubseteq \text{has} . \end{aligned}$$

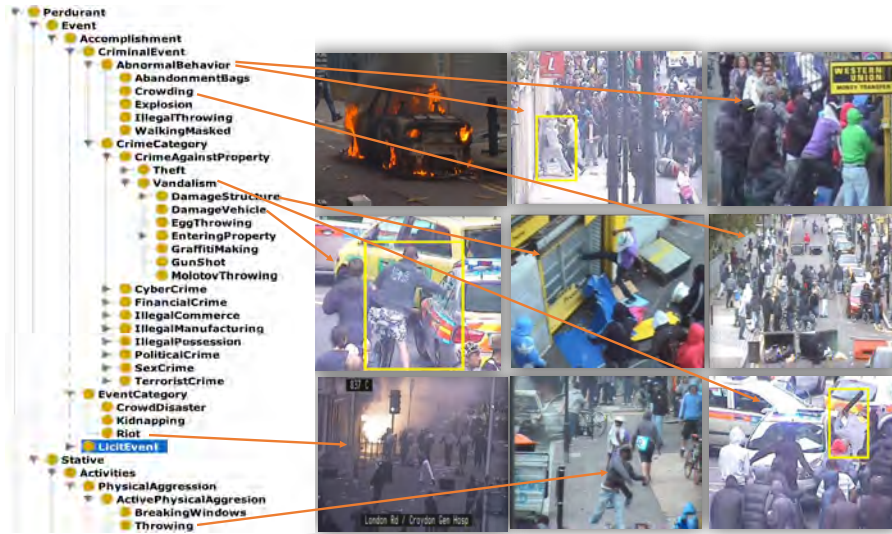


Figure 7: Examples of still image annotations from the London Riots 2011 of events as per Table 2.

Note that in the last RI axiom, \circ is role composition and, thus, $\text{has} \circ \text{has} \sqsubseteq \text{has}$ dictates that the property has is transitive, while with $\text{has} = \text{isFrom}^-$ we say that isFrom is the inverse of has . Therefore, isFrom is transitive as well.

The following example illustrates the mechanism of image of annotation together with a meaningful inference.

Example 1. Consider the following DL axioms resulting from annotating images of a video (video6) registered by a camera (cameraC004):

```

participateIn(personA, throwing5)
Throwing(throwing5)
NaturalPerson(personA)
Throwing  $\sqsubseteq$  ActivePhysicalAggression
ActivePhysicalAggression  $\sqsubseteq$  PhysicalAggression
PhysicalAggression  $\sqsubseteq$  Process
isFrom(throwing5, endurant6)
Resource(endurant6)
hasVideoid(endurant6, video6)
Source(endurant7)
hasCameraId(endurant7, cameraC004)
has(endurant7, endurant6) .

```

Now, as isFrom is transitive, we may infer:

$\text{isFrom}(\text{throwing5}, \text{endurant7}) .$




Figure 8: Modelling transitive property axiom in our defined ontology.



Figure 9: Example of DamageVehicle and DamageStructure scenes in CCTV.

Then, it is not difficult to see that we finally infer

$$\exists \text{participateIn} . (\text{PhysicalAggression} \sqcap \exists \text{isFrom} . (\text{Source} \sqcap \exists \text{hasCameraID} . \{\text{cameraC004}\})) (\text{PersonA}) ,$$

which can be read as:

“A person (PersonA) participated in a physical aggression that has been registered by camera C004”.

□

5.2 Modelling GCIs for Vandalism Event Detection

As we are focusing on forensic domain and dealing with variety of concepts aiming at aiding forensic analysis, to objectively identify and represent complex events, we next show that a (manually build) *General Concept Inclusion* (GCI) axiom may help to classify high-level events in terms of a composition of some lower level events. The following are such GCI examples:

DamageVehicle:

$$\begin{aligned} & \text{Perdurant} \sqcap \\ & \exists \text{participant} . (\text{Vehicle} \sqcap \\ & \quad \exists \text{participantIn} . (\text{BreakingDoor} \sqcup \text{BreakingWindows})) \\ & \quad \sqsubseteq \text{DamageVehicle} . \end{aligned}$$



Figure 10: Example of Vandalism scenes in CCTV videos.

“If an event involves a vehicle that is subject of a breaking door or breaking windows then the event is about a damaged vehicle” (see Figure 9).

DamageStructure:

$$\begin{aligned} & \text{Perdurant} \sqcap \\ & \quad \exists \text{participant.}(\text{Structure} \sqcap \\ & \quad \quad \exists \text{participantlin.Kicking}) \sqsubseteq \text{DamageStructure} . \end{aligned}$$

“If an event involves a structure that is subject of kicking, then the event is about a damaged structure” (see Figure 9).

The following example illustrates the use of such GCIs.

Example 2. *Suppose we have an image classifier that is able to provide us with the following facts. Specifically, assume it is able to identify vehicles and breaking windows:*

$$\begin{aligned} & \text{participant}(\text{Perdurant2}, \text{Endurant1}), \text{Vehicle}(\text{Endurant1}) \\ & \text{BreakingWindows}(\text{Perdurant2}) . \end{aligned}$$

From these facts and the GCI about DamageVehicle, we may infer that the image is about a damaged vehicle, i.e. we may infer

$$\text{DamageVehicle}(\text{Perdurant2}) .$$

□

The following set of GCIs illustrates instead how one may have multiple GCIs to classify a single event, such as those for Vandalism (see, e.g. Figure 10).⁶

$$\begin{aligned} & \text{Perdurant} \sqcap \exists \text{part.}(\text{Crowding} \sqcap \text{DamageStructure}) \sqsubseteq \text{Vandalism} \\ & \text{Perdurant} \sqcap \exists \text{part.}(\text{Crowding} \sqcap \text{DamageVehicle}) \sqsubseteq \text{Vandalism} \\ & \text{Perdurant} \sqcap \exists \text{part.}(\text{Explosion} \sqcap \text{Throwing}) \sqsubseteq \text{Vandalism} . \end{aligned}$$

Note that in the example above, we assume that events (perdurant) may be complex in the sense that they may compose by multiple sub-events (parts). So, e.g. in the last GCI, we roughly state

“If a (complex) event involves both throwing and an explosion (two sub-events) then the event is about vandalism”.

⁶Recall that all these GCIs provide sufficient conditions to be an instance of Vandalism, but no necessary condition.



Figure 11: Examples of events that happen in the same location (`locatedSameAs`) from CCTV.

Following our previous examples, we next are going to formulate another kind of background knowledge. Our main focus in this example is on recognizing high-level events, which occur in the same location (same street in our modelling). In order to model this scenario, we may use the *Semantic Web Rule Language* (SWRL) to model the `locatedSameAs` role and then use it in GCIs. The SWRL rule is:

“Two perdurants that occur in the same street occur in the same place.

$$\text{Perdurant}(?p1), \text{Perdurant}(?p2), \text{hasLocationName}(?p1, ?l1), \\ \text{hasLocationName}(?p2, ?l2), \text{SameAs}(?l1, ?l2) \rightarrow \text{locatedSameAs}(?p1, ?p2) .$$

The following axioms illustrate how to use the previously defined relation (few examples captured from our data set by these rules are illustrated in Figure 11).

$$\text{Perdurant} \sqcap \\ \exists \text{part.}(\text{Crowding} \sqcap \exists \text{locatedSameAs.} \text{Explosion}) \sqsubseteq \text{Vandalism}$$

$$\text{Perdurant} \sqcap \\ \exists \text{part.}(\text{Crowding} \sqcap \exists \text{locatedSameAs.} \text{DamageStructure}) \sqsubseteq \text{Vandalism}$$

$$\text{Perdurant} \sqcap \\ \exists \text{part.}(\text{Crowding} \sqcap \exists \text{locatedSameAs.} \text{Throwing}) \sqsubseteq \text{Vandalism}$$

$$\text{Perdurant} \sqcap \\ \exists \text{part.}(\text{DamageStructure} \sqcap \exists \text{locatedSameAs.} \text{Throwing}) \sqsubseteq \text{Vandalism} .$$

6 Experiments

We conducted two experiments with our ontology, which we are going to describe in the following.⁷

In the first case, we evaluated the classification effectiveness of manually built GCIs to identify crime events, while in the second case we drop the manual-built GCIs and, try to learn such GCIs instead automatically from examples and compare their effectiveness with respect the manually built ones.

⁷The ontologies used in the experiments and experimental results can be found at <http://www.umbertostraccia.it/cs/ftp/ForensicOntology.zip>.

Table 2: Criminal event classes considered.

Vandalism (13, 57)	Riot (4, 21)	AbnormalBehavior (2, 80)
Crowding (1, 64)	DamageStructure (3, 9)	DamageVehicle (3, 16)
Throwing (1, 30)		

Table 3: Ontology Metrics.

Axioms	9889	SubclassOf axioms count	532
Logical axiom count	7176	EquivalentClasses axioms count	5
Class count	483	DisjointClasses axioms count	11
Object property count	148	GCI count	38
Data property count	51	Hidden GCI Count	5
Individual count	1800		
DL expressivity	SHIQ(D)		

SubObjectPropertyOf axioms count	93	SubDataProperty axioms count	11
InverseObjectProperties axioms count	20	DataPropertyDomain axioms count	1
TransitiveObjectProperty axioms count	5	DataPropertyRange axioms count	5
SymmetricObjectProperty axioms count	2		
ObjectPropertyDomain axioms count	19		
ObjectPropertyRange axioms count	18		

ClassAssertion axioms count	1793	AnnotationAssertion axioms count	195
ObjectPropertyAssertion axioms count	2964		
DataPropertyAssertion axioms count	1706		

6.1 Classification via Manually Built GCIs

Roughly, we have considered a number of crime videos, annotated them manually and then checked whether the manually built GCIs, as described in Section 5.2, were able to determine crime events correctly.

Setup. Specifically, we considered our ontology and around 3.07 TB of video data about the London riot 2011,⁸ of which 929 (GB) is in a non-proprietary format. We considered 140 videos (however, the videos cannot be made publicly available). Within these videos, all the available CCTV cameras (35 CCTV) along with their features such as latitude, longitude, start time, end time and street name, have been annotated manually according to our methodology described in Section 5 and included into our ontology. We have also calculated all the geographic distances between each camera. The resulting ontology contains 1800 created individuals of which, e.g. 106 are of type `Event`.

Then, we considered criminal events occurring in the videos (specifically, we focused on vandalic events). For each class of events, we manually built one or more GCIs, as illustrated in Section 5.2. The list of crime events considered is reported in Table 2. In it, the first number in parenthesis

⁸These are part of the EU funded project LASIE “Large Scale Information Exploitation of Forensic Data”, <http://www.lasie-project.eu>.

reports the number of GCIs we built for each of them, while the second number indicates the number of event instances (individuals) we created during the manual video annotation process. So, for instance, for the event `DamageStructure` we have built 3 classification GCIs and we have created 9 instances of `DamageStructure` during the manual video annotation process. For further clarification, the 3 GCIs for `DamageStructure` are

$$\begin{aligned} & \text{Perdurant } \sqcap \\ & \quad \exists \text{participant.}(\text{Structure } \sqcap \\ & \quad \quad \exists \text{participantIn.Kicking}) \sqsubseteq \text{DamageStructure} \\ & \text{Perdurant } \sqcap \\ & \quad \exists \text{participant.}(\text{Structure } \sqcap \\ & \quad \quad \exists \text{participantIn.Beating}) \sqsubseteq \text{DamageStructure} \\ & \text{Perdurant } \sqcap \\ & \quad \exists \text{participant.}(\text{Structure } \sqcap \\ & \quad \quad \exists \text{participantIn.BreakingWindows}) \sqsubseteq \text{DamageStructure} , \end{aligned}$$

while, e.g., an instance of `DamageStructure` is the individual `Kicking1`, whose related information excerpt is:

`Kicking(Kicking1), isFrom(Kicking1, 2bdf), Resource(2bdf), isFrom(2bdf, C004),
has(2bdf, pr11), part(pr11, Kicking1), part(pr11, BreackingWindows3),
BreackingWindows(BreackingWindows3), ...`

As a matter of general information, the global metric statistics of the so built ontology is reported in Table 3.

Evaluation. Let \mathcal{O} be the built ontology from which we drop axioms stating explicitly that an individual is an instance of a crime event listed in Table 2. Please note that without the GCIs none of the crime events instances in \mathcal{O} can be inferred to be instances of the crime events in Table 2.⁹

Now, on \mathcal{O} we run an OWL 2 reasoner that determines the instances of all crime event classes in the ontology.

To determine the classification effectiveness of the GCIs, we compute the so-called micro/macro averages of precision, recall and F1-score w.r.t. inferred data, which for clarity and completeness we report below. Specifically, we measure the outcomes of binary classification using a 2×2 *contingency table* or *confusion matrix*, as follows: for each class C in Table 2, consider the following contingency table

		True Condition	
		in C	is not in C
Prediction	in C	TP_C	FP_C
	is not in C	FN_C	TN_C

where,

- the *true positives* TP_C is calculated as the number of positive predictions for the class C provided that the actual value is positive;

⁹Roughly, crime events are subclasses of the `Event` class, while crime event instances are instances of the class `Stative` (see Figure 5).

- the *false positives* FP_C is calculated as the number of positive predictions for the class C provided that the actual value is negative;
- the *false negatives* FN_C is calculated as the number of negative predictions for the class C provided that the actual value is positive;
- the *true negatives* TN_C is calculated as the number of negative predictions for the class C provided that the actual value is negative.

Formally, for each crime event class C , let $true(C)$ be the set of manually determined instances e of it, *i.e.* $e \in true(C)$ denotes that fact that e has to be an instance of C . Then¹⁰

$$\begin{aligned}
TP_C &= \{e \mid \mathcal{O} \models C(e) \text{ and } e \in true(C)\} \\
FP_C &= \{e \mid \mathcal{O} \models C(e) \text{ but } e \notin true(C)\} \\
FN_C &= \{e \mid \mathcal{O} \not\models C(e) \text{ but } e \in true(C)\} \\
TN_C &= \{e \mid \mathcal{O} \not\models C(e) \text{ and } e \notin true(C)\} .
\end{aligned}$$

Thereafter, with $|C| = |TP_C| + |FP_C|$, $|true(C)| = |TP_C| + |FN_C|$, we do the following:

1. for each class C in Table 2, we determine

$$\begin{aligned}
Precision_C &= \frac{|TP_C|}{|C|} \\
Recall_C &= \frac{|TP_C|}{|true(C)|} \\
F1_C &= 2 \cdot \frac{Precision_C \cdot Recall_C}{Precision_C + Recall_C}
\end{aligned}$$

2. for each of these measures we also compute the micro- and macro-average (where N is the number of classes in Table 2):

$$\begin{aligned}
Precision_{micro} &= \frac{\sum_C |TP_C|}{\sum_C |C|} \\
Recall_{micro} &= \frac{\sum_C |TP_C|}{\sum_C |true(C)|} \\
F1_{micro} &= 2 \cdot \frac{Precision_{micro} \cdot Recall_{micro}}{Precision_{micro} + Recall_{micro}} \\
Precision_{macro} &= \frac{\sum_C Precision_C}{N} \\
Recall_{macro} &= \frac{\sum_C Recall_C}{N} \\
F1_{macro} &= 2 \cdot \frac{Precision_{macro} \cdot Recall_{macro}}{Precision_{macro} + Recall_{macro}}
\end{aligned}$$

The evaluation result of the first test is shown in Table 4.

¹⁰Recall that $\mathcal{O} \models C(e)$ dictates that e has been automatically classified/inferred as being instance of class C .

Table 4: Results for the experiment on classification of manually build GCIs .

Event	TP	FP	FN	TN	$ C $	$ trueC $	$Precision_C$	$Recall_C$	$F1_C$
Vandalism	42	0	15	168	42	57	1.00	0.74	0.85
DamageVehicle	11	0	5	209	11	16	1.00	0.69	0.81
DamageStructure	9	0	0	216	9	9	0.89	0.89	0.89
Crowding	60	1	4	160	61	64	0.98	0.94	0.96
Throwing	30	0	0	195	30	30	1.00	1.00	1.00
Riot	5	0	16	204	5	21	1.00	0.24	0.38
AbnormalBehaviour	70	22	10	123	92	80	0.76	0.88	0.81
	$Precision_{micro}$	$Recall_{micro}$	$F1_{micro}$	$Precision_{macro}$	$Recall_{macro}$	$F1_{macro}$			
	0.91	0.82	0.86	0.96	0.78	0.86			

6.2 Classification via Automatically Learned GCIs

In the second experiment, we apply a concept learning approach to replace the manually built GCIs describing the crime events listed in Table 2. To this end, the DL-Learner¹¹ system was used to find descriptions of the criminal events in Table 2, based on existing instances of these classes. We recall roughly that DL-learner is, among others, a framework for supervised Machine Learning in OWL 2, which is capable of making class expression suggestions for a specified class C by relying on the instances of C .

Setup. Let now \mathcal{O} be the ontology as in Section 6.1, but from which we also drop additionally the manually created GCIs for the crime event listed in Table 2. On it we used the CELOE algorithm [6, 15] with its default settings to generate suggestion definitions (inclusion axioms) for each target class C by using the manually identified crime event instances.

Specifically, we used a K -fold cross style validation method [8], which divides the available crime event instances into a K disjoint subsets. That is, we split each target class C into K disjoint subsets C_1, \dots, C_K of equal size. In our experiment, K is the number of instances of C and, thus, each C_i has size one. For each C_i , the *training set* is $(\bigcup_{i=1}^K C_i) \setminus C_i$ and is denoted as $Trainset_i$. Then, for each C_i we run CELOE on the training set $Trainset_i$ and generated at most 10 class expressions of the form $D_j \sqsubseteq C_i$, out of which we have chosen the best solution (denoted $D_{C_i} \sqsubseteq C_i$ or GCI_i). If the best solution is not unique, we select the first listed one.

The best-selected GCIs found by CELOE for each of the target classes in Table 2 are:

PhysicalAggression \sqcap \exists immediateRelation.Structure \sqsubseteq DamageStructure

\exists immediateRelation.Vehicle \sqsubseteq DamageVehicle

\exists immediateRelation.Vandalism \sqsubseteq AbnormalBehavior

\exists immediateRelation.Arm \sqsubseteq Throwing

\exists immediateRelation.Group \sqsubseteq Crowding .

With the help of a reasoner, we then infer all instances in \mathcal{O} , that are not in $Trainset_i$, that are instances of the selected D_{C_i} and consider them as our *result set* (denoted $Resultset_i$).

Evaluation. To determine the classification effectiveness of the learned GCIs, *i.e.* of GCI_i , precision

¹¹<http://dl-learner.org/>

Table 5: Results for the experiment on classification using DL-Learner CELOE algorithm.

Event	$Precision_C$	$Recall_C$	$F1_C$
DamageVehicle	0.69	0.98	0.81
Damage Structure	1.00	1.00	1.00
Crowding	0.96	1.00	0.98
Throwing	0.86	0.99	0.92
AbnormalBehavior	0.69	0.99	0.81

$Precision_{micro}$	$Recall_{micro}$	$F1_{micro}$	$Precision_{macro}$	$Recall_{macro}$	$F1_{macro}$
0.753	0.964	0.845	0.599	0.709	0.649

(Pr_i) and recall (Re_i) across the folds is computed:

- the *false positives* FP_i is calculated as the difference between the instances in $Resultset_i$ and those in $true(C)$;
- the *true positives* TP_i is calculated as the difference between $Resultset_i$ and the false positive;
- the *false negatives* FN_i is calculated as the difference between C_i and $Resultset_i$.

Using FP_i, TP_i and FN_i , we have that $Pr_i = \frac{|TP_i|}{|TP_i|+|FP_i|}$, while $Re_i = \frac{|TP_i|}{|TP_i|+|FN_i|}$. We then determine both average precision ($Precision_C$) and recall ($Recall_C$) across the folds, *i.e.*

$$Precision_C = \frac{1}{K} \cdot \sum_{i=1}^k Pr_i$$

$$Recall_C = \frac{1}{K} \cdot \sum_{i=1}^k Re_i .$$

Using these measures we then compute the micro- and macro-averages of precision, recall and F1 as in Section 6.1.

The evaluation results of the second test are shown in Table 5.

Discussion. The results are generally promising. In the manually built GCI case, precision and F1 are reasonably good, though in one case (Riot) the recall and, thus, F1 is not satisfactory. For the learned GCI case, the individual measures are generally comparable to the manual ones.

Given that the learned GCIs are completely different than the manually built ones, it is surprising that both sets perform more or less the same. However, please note that DL-Learner was neither able to learn a GCI for Vandalism nor for Riot. This fact is reflected in the generally worse micro/macro precision, recall and F1 measures.

Eventually, we also merged the manually built GCIs and the learned ones together and tested them as in Section 6.1. The results in Table 6 show, however, that globally their effectiveness is as for the manual case (and does not improve).

Table 6: Results of merging manual and learned GCIs.

Event	$Precision_C$	$Recall_C$	$F1_C$
Vandalism	1.00	0.74	0.85
DamageVehicle	1.00	0.69	0.81
Damage Structure	0.89	0.89	0.89
Crowding	0.98	0.94	0.96
Throwing	1.00	1.00	1.00
Riot	1.00	0.24	0.38
AbnormalBehavior	0.76	0.89	0.82

$Precision_{micro}$	$Recall_{micro}$	$F1_{micro}$	$Precision_{macro}$	$Recall_{macro}$	$F1_{macro}$
0.90	0.82	0.86	0.95	0.77	0.85

7 Conclusions

In this work, we have proposed an extensive ontology for representing complex criminal events. The proposed ontology focuses on events that are often required by forensic analysts. In this context, the *Perdurant*, as defined in the DOLCE ontology as an occurrence in time, and the *Endurant*, defined in the DOLCE ontology as contentious in time, have both been extended to represent all forensics entities together with meaningful entities for video surveillance-based vandalism detection. The aim of the built ontology is to support the interoperability of the automated surveillance system.

To classify high-level events in terms of the composition of lower level events we focused on both manually built and automatically learned GCIs and have compared the evaluation results of both experiments. The results are generally promising and the effectiveness of machine derived definitions for high-level crime events is encouraging though needs further development.

In the future, we intend to deal with vague or imprecise knowledge and we would like to work on the problem of automatically learn fuzzy concept description [4, 5, 16, 17, 18, 33, 34] as most of the involved entities are fuzzy.

Acknowledgements

This work is partially funded by the European Union’s Seventh Framework Programme, grant agreement number 607480 (LASIE IP project).

References

- [1] Preetha Appan and Hari Sundaram. Networked multimedia event exploration. In *Proceedings of the 12th Annual ACM International Conference on Multimedia*, pages 40–47. ACM, 2004.

- [2] Franz Baader, Diego Calvanese, Deborah McGuinness, Daniele Nardi, and Peter F. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press, 2003.
- [3] Franz Baader, Ian Horrocks, and Ulrike Sattler. Description logics. In Steffen Staab and Rudi Studer, editors, *Handbook on Ontologies*, International Handbooks on Information Systems, pages 21–43. Springer Verlag, 2009.
- [4] Fernando Bobillo and Umberto Straccia. Fuzzy ontology representation using OWL 2. *International Journal of Approximate Reasoning*, 52:1073–1094, 2011.
- [5] Fernando Bobillo and Umberto Straccia. The fuzzy ontology reasoner *fuzzyDL*. *Knowledge-Based Systems*, 95:12 – 34, 2016.
- [6] Lorenz Bühmann, Jens Lehmann, and Patrick Westphal. DL-Learner framework for inductive learning on the semantic web. *Web Semantics: Science, Services and Agents on the World Wide Web*, 39:15–24, 2016.
- [7] Roberto Casati and Achille Varzi. Events. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Winter 2015 edition, 2015. <http://plato.stanford.edu/archives/win2015/entries/events/>.
- [8] George Forman and Martin Scholz. Apples-to-apples in cross-validation studies: pitfalls in classifier performance measurement. *ACM SIGKDD Explorations Newsletter*, 12(1):49–57, 2010.
- [9] Alexandre R. J. François, Ram Nevatia, Jerry R. Hobbs, and Robert C. Bolles. VERL: an ontology framework for representing and annotating video events. *IEEE MultiMedia*, 12(4):76–86, 2005.
- [10] Asaad Hakeem, Khurram Shafique, and Mubarak Shah. An object-based video coding framework for video sequences obtained from static cameras. In *Proceedings of the 13th Annual ACM International Conference on Multimedia*, pages 608–617. ACM, 2005.
- [11] Asaad Hakeem, Yaser Sheikh, and Mubarak Shah. CASE^E: A hierarchical event representation for the analysis of videos. In *Proceedings of the 19th National Conference on Artificial Intelligence*, pages 263–268. AAAI Press, 2004.
- [12] Craig Henderson, Saverio G. Blasi, Faranak Sobhani, , and Ebroül Izquierdo. On the impurity of street-scene video footage. In *6th International Conference on Imaging for Crime Prevention and Detection (ICDP-15)*. The Institution of Engineering and Technology (IET), 2015.
- [13] Ramesh Jain, Pilho Kim, and Zhao Li. Experiential meeting system. In *Proceedings of the 2003 ACM SIGMM Workshop on Experiential Telepresence*, pages 1–12. ACM, 2003.
- [14] Pilho Kim, Mark Podlaseck, and Gopal Pingali. Personal chronicling tools for enhancing information archival and collaboration in enterprises. In *Proceedings of the the 1st ACM Workshop on Continuous Archival and Retrieval of Personal Experiences*, pages 56–65. ACM, 2004.

- [15] Jens Lehmann, Sören Auer, Lorenz Bühmann, and Sebastian Tramp. Class expression learning for ontology engineering. *Web Semantics: Science, Services and Agents on the World Wide Web*, 9(1):71–81, 2011.
- [16] Francesca A. Lisi and Umberto Straccia. A logic-based computational method for the automated induction of fuzzy ontology axioms. *Fundamenta Informaticae*, 124(4):503–519, 2013.
- [17] Francesca Alessandra Lisi and Umberto Straccia. A foil-like method for learning under incompleteness and vagueness. In *23rd International Conference on Inductive Logic Programming*, volume 8812 of *Lecture Notes in Artificial Intelligence*, pages 123–139, Berlin, 2014. Springer Verlag. Revised Selected Papers.
- [18] Thomas Lukasiewicz and Umberto Straccia. Managing uncertainty and vagueness in description logics for the semantic web. *Journal of Web Semantics*, 6:291–308, 2008.
- [19] Claudio Masolo, Stefano Borgo, Aldo Gangemi, Nicola Guarino, and Alessandro Oltramari. Wonderweb Deliverable D18, Ontology Library (final). *ICT project*, 33052, 2003.
- [20] Carlo Meghini, Fabrizio Sebastiani, and Umberto Straccia. A model of multimedia information retrieval. *Journal of the ACM*, 48(5):909–970, 2001.
- [21] Boris Motik, Bernardo Cuenca Grau, Ian Horrocks, Zhe Wu, Achille Fokoue, Carsten Lutz, et al. OWL 2 Web Ontology Language Profiles. *W3C recommendation*, 27:61, 2009.
- [22] Ram Nevatia, Jerry R. Hobbs, and Bob Bolles. An ontology for video event representation. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, page 119. IEEE Computer Society, 2004.
- [23] Ram Nevatia, Tao Zhao, and Somboon Hongeng. Hierarchical language-based representation of events in video streams. In *IEEE Conference on Computer Vision and Pattern Recognition*, page 39. IEEE Computer Society, 2003.
- [24] SWRL: A Semantic Web Rule Language Combining OWL and RuleML. <https://www.w3.org/Submission/SWRL/>. W3C, 2004.
- [25] OWL 2 Web Ontology Language Document Overview. <http://www.w3.org/TR/2009/REC-owl2-overview-20091027/>. W3C, 2009.
- [26] Gopal Sarma Pingali, Agata Opalach, Yves D. Jean, and Ingrid B. Carlbom. Instantly indexed multimedia databases of real world events. *IEEE Transactions on Multimedia*, 4(2):269–282, 2002.
- [27] Susan Rothstein. Chapter 1: Verb classes and aspectual classification. In *Structuring Events: A Study in the Semantics of Lexical Aspect*, pages 1–35. Wiley Online Library, 2004.
- [28] Ansgar Scherp, Thomas Franz, Carsten Saathoff, and Steffen Staab. F–A model of events based on the foundational ontology DOLCE+DnS ultralight. In *Proceedings of the Fifth International Conference on Knowledge Capture*, pages 137–144. ACM, 2009.
- [29] Manfred Schmidt-Schauß and Gert Smolka. Attributive concept descriptions with complements. *Artificial Intelligence*, 48:1–26, 1991.

- [30] Lauro Snidaro, Massimo Belluz, and Gian Luca Foresti. Representing and recognizing complex events in surveillance applications. In *Fourth IEEE International Conference on Advanced Video and Signal Based Surveillance*, pages 493–498. IEEE Computer Society, 2007.
- [31] Faranak Sobhani, Krishna Chandramouli, Qianni Zhang, and Ebroul Izquierdo. Formal representation of events in a surveillance domain ontology. In *2016 IEEE International Conference on Image Processing*, pages 913–917. IEEE Computer Society, 2016.
- [32] Faranak Sobhani, Nur Farhan Kahar, and Qianni Zhang. An ontology framework for automated visual surveillance system. In *13th International Workshop on Content-Based Multimedia Indexing*, pages 1–7. IEEE Computer Society, 2015.
- [33] Umberto Straccia. *Foundations of Fuzzy Logic and Semantic Web Languages*. CRC Studies in Informatics Series. Chapman & Hall, 2013.
- [34] Umberto Straccia and Matteo Mucci. pFOIL-DL: Learning (fuzzy) \mathcal{EL} concept descriptions from crisp OWL data using a probabilistic ensemble estimation. In *Proceedings of the 30th Annual ACM Symposium on Applied Computing (SAC-15)*, pages 345–352, Salamanca, Spain, 2015. ACM.
- [35] Zeno Vendler. Verbs and times. *The Philosophical Review*, 62(2):143–160, 1957.
- [36] Zeno Vendler, editor. *Linguistics in Philosophy*. G - Reference, Information and Interdisciplinary Subjects Series. Cornell University Press, 1967.
- [37] Utz Westermann and Ramesh Jain. Toward a common event model for multimedia applications. *IEEE Multimedia*, 14(1), 2007.