XR4DRAMA Knowledge Graph: A Knowledge Graph for Media Planning

Alexandros Vassiliades[a], Spyridon Symeonidis[b], Sotiris Diplaris[c], Georgios Tzanetis, Stefanos Vrochidis[d] and Ioannis Kompatsiaris[e]
Information Technologies Institute, Center for Research and Technology Hellas, Thessaloniki, Greece

Keywords: Knowledge Graph, Media Planning, Points of Interest, POI Management Mechanism.

Abstract: In the previous two decades, knowledge graphs have evolved, inspiring developers to build even more context-related Knowledge Graphs. Because of this development, artificial intelligence applications can now access open domain-specific information in a format that is both semantically rich and machine comprehensible. In this paper, we introduce the XR4DRAMA Knowledge Graph, which can serve as a representation of media planning information. The XR4DRAMA knowledge graph can specifically represent data about the following: (a) Observations and Events (for example, data information from photos and text messages); (b) Spatial and Temporal data, such as coordinates or labels of locations and timestamps; and (c) Tasks and Plans for media planning. In addition, we provide a mechanism that allows Points of Interest to be created or updated based on videos, photos, and text messages sent by users. For improved media coverage of a remote location, Points of Interest serve as markers to journalists.

1 INTRODUCTION

The next step toward making Knowledge Graphs (KGs) the primary knowledge representation format for the Web, looks to be the development of context related KGs, i.e., KGs that can only be used in particular environments (Berners-Lee et al., 2001). In this work, we focus on representing information for media planning, more specifically information about: (a) Observations and Events (for example, information from photos, and information from text messages); (b) Spatial and Temporal data, such as coordinates or labels of locations and timestamps; and (c) Tasks and Plans for media planning. Information representation in Linked Open Data format aids in the data's reuse and linkage with other KGs (Ehrlinger and Wöß, 2016; Villazon-Terrazas et al., 2021). When covering a recording of a documentary in a remote location, unknown to the media production team, a journalist should have access to geospatial data that gives details on the location they will be visiting. This information will enable him or her to accurately and cost-effectively setup the production. For this reason, we offer a tool that allows users to add or update Points-of-Interest (POIs)[1]. We use the term *POI management mechanism* to refer to the process for creating and updating POIs, throughout the paper. A POI, according to the official definition, is a particular place or location point on a map that a user would find useful or interesting. In our situation, POIs also comprise geospatial data that includes details from user-provided videos, photos, and text messages on the condition of a location. As a result, POIs contain data that can assist journalists in actual situations involving media preparation.

The XR4DRAMA KG was created, to serve as the knowledge representation for the XR4DRAMA project[2]. Through the use of many technologies, including eXtended Reality (XR), XR4DRAMA is committed to enhancing situation awareness. Media planning is one of the XR4DRAMA project's primary use cases. In short, XR4DRAMA project stands in three key-points: (a) Facilitating the gathering of all necessary (digital) information for a particular, dif-

[a] https://orcid.org/0000-0003-4569-503X
[b] https://orcid.org/0000-0003-3170-1750
[c] https://orcid.org/0000-0002-9969-6436
[d] https://orcid.org/0000-0002-2505-9178
[e] https://orcid.org/0000-0001-6447-9020

[1] https://xr4drama.eu/2022/07/07/xr4drama-pois-virtual-whiteboards/
[2] https://xr4drama.eu

Vassiliades, A., Symeonidis, S., Diplaris, S., Tzanetis, G., Vrochidis, S. and Kompatsiaris, I.
XR4DRAMA Knowledge Graph: A Knowledge Graph for Media Planning.
DOI: 10.5220/0011621600003393
In *Proceedings of the 15th International Conference on Agents and Artificial Intelligence (ICAART 2023) - Volume 3*, pages 124-131
ISBN: 978-989-758-623-1; ISSN: 2184-433X

ficult or even dangerous situation that a media team faces, (b) Utilizing extended reality technologies to simulate an environment "as if on site" in order to accurately predict an event or incidence, and (c) Establishing a shared understanding of an environment and giving users of the project's platform (first responders in the field/control room, citizens, and journalists) the option to update representations of locations as events change, allowing them to comprehend and re-evaluate the effects of particular actions/decisions. As a result, the XR4DRAMA KG may incorporate the findings of several advanced analysis components that process multimodal data and depict the structures they produce (in this project, for the media use case, we integrate visual and textual analysis messages). Additionally, the XR4DRAMA KG provides a innovative method through its POI management mechanism that may generate or update POIs, which contain essential geospatial data that can make it easier for journalists to cover the production recordings.

As part of their day-to-day business journalists and other media houses produce news coverage in various locations. Despite thorough research and preparations, remote production planning very often runs into challenges and difficulties. Many depend on the characteristics of the individual location and the situation on the ground. These challenges can be circumstantial and organisational - like accessibility, noise, the presence of people, the lack of infrastructure (from electricity supply to parking space), the wrong choice of equipment or other filming restrictions. Hence, it is crucial for journalists and media houses to access information about the state of a location, such as the accessibility of the location, among others, in order to plan their media coverage.

The challenge is to combine all of this information into a coherent image to give everyone a precise picture of the location and the situation on the ground in order to prepare themselves for a smooth and safe production. The XR4DRAMA KG is capable of filling the aforementioned gap in distribution of crucial knowledge to journalists, in order for them to be able to plan more efficiently the news coverage.

Our contribution in this paper, is on the one hand the XR4DRAMA KG, which can represent multimodal measures and let journalists report the important elements in a location as effectively as possible. Next, the development of the POI management mechanism for the XR4DRAMA KG, which can be useful in real-life occasions by creating and maintaining POIs that will further facilitate journalists' work by informing them of the location's current condition.

The rest of this paper is organized as follows. Section 2, contains the related work. The XR4DRAMA KG, the POI management mechanism that creates or updates POIs, and the data upon which we built the KG are all presented in Section 3. The POI mechanism system and the evaluation of the KG are both found in section 4. Lastly, we conclude our paper with Section 5.

2 RELATED WORK

In this section we present other KGs that are found in the area of KGs for media planning, fake news detection, media planning through disasters, and other KGs which are related with media, media houses and news in general.

The area of KGs for media planning is not very rich, as only a handful of studies can be classified in this area. For instance, the studies of Opdahl et al. (Opdahl et al., 2016) and Berven et al. (Berven et al., 2018) are two similar studies which present mechanisms for media planning. The basic concept for both these studies is that they offer a news extraction mechanism which based on the semantics of a KG, will extract related posts from social web sites, and other well-known media houses, about an event. This synergy of news extraction, and subsequently representation of knowledge about an event, is set to help a journalists to see what parts of the event have been covered, and what are the restrictions for accessing a location to cover the event. The difference with the XR4DRAMA KG lies mostly in the POI management mechanism, as we offer the most crucial information about an event in a POI, and therefore make it more easily digestible for the journalists, while (Opdahl et al., 2016) and (Berven et al., 2018) return information in the form of text which can be more time consuming for an individual to process.

The area of media planning in disasters based on KGs is also not so rich, as to the best of our knowledge, only a few studies exist in this area. KG in media are mostly used for fake news detection (Pan et al., 2018) and building event-centric news (Rospocher et al., 2016; Tang et al., 2019). Some exceptions are (Wang and Hou, 2018) and (Ni et al., 2019). In the former, the authors propose a method to construct a KG for disaster news based on an address tree. Address Trees, are tree structures which analyze an address having as root the broader region. For example, *home address* $\rightarrow$ *town district* $\rightarrow$ *town*, is a small address tree. In the latter, the authors present a data driven model which generates storylines from huge amount of web information and proposes a KG-based disaster storyline generating framework. For the work of (Wang and Hou, 2018), comparing to

XR4DRAMA KG there is not a mechanism for creating POIs, and the indication for the location is given in string descriptions which can be obscure in some cases, while XR4DRAMA KG represents locations with coordinates. For (Tang et al., 2019), the issue of noise in the data inserted in the KG is addressed, an issue which is not part of the XR4DRAMA KG (see subsection 4.1).

The following studies (Rospocher et al., 2016; Tang et al., 2019), present methods and tools to automatically build KGs from news articles. As news articles describe changes in the world through the events they report, an approach is presented to create event-centric KGs using state-of-the-art natural language processing and semantic web techniques. Such event-centric KGs capture long-term developments and histories on hundreds of thousands of entities and are complementary to the static encyclopedic information in traditional knowledge graphs. Even though these two studies might not solve exactly the same problem with XR4DRAMA KG, the crucial information can be accessed through sophisticated SPARQL queries, which might not be user-friendly even with an UI. On the other hand, XR4DRAMA through its POI management mechanism serves the crucial information about an event, with a POI, which is more easily understandable by an individual.

One can have a more detailed view at KGs for media by reading the survey (Opdahl et al., 2022).

3 XR4DRAMA KNOWLEDGE GRAPH

The project's platform's back-end includes the XR4DRAMA KG. Because of this, a detailed investigation of the multimodal mapping mechanism that accepts messages from the visual and textual analysis components and transmits their content into the XR4DRAMA KG will be skipped. But one can find a blueprint of these messages here[3]. Moreover, the source code of the multimodal mapping mechanism can be found here[4]. The idea behind the pipeline is to map the data into the KG after the multimodal mapping mechanism has received the message from a component. The POI management mechanism of the XR4DRAMA KG will then construct a new POI or update an existing one based on the information in the message and the information from the KG, when the message is received from the textual or visual analysis component. In the second case, the premise is that a generated POI's status has changed, for instance, the area has become crowded, necessitating an update of the POI's metadata. The pipeline is depicted in Figure 1, where each circled number denotes a different step's sequence.

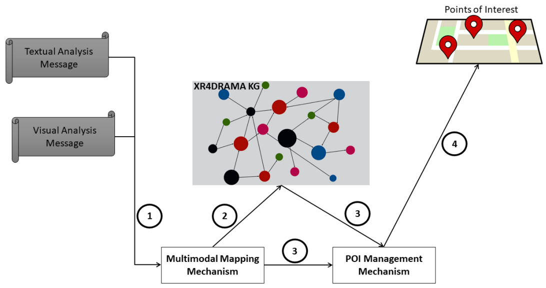

Figure 1: Pipeline of the XR4DRAMA KG.

3.1 Nature of Data

The use of a semantic KG was necessary to meet the project's needs due to the system's multimodality, diversity, and need for homogeneity and fusion. There is an underlying data storage facility for this purpose, thus the XR4DRAMA KG is not in charge of archiving and storing raw data files. Instead, the XR4DRAMA KG stores raw data metadata, analysis findings, and other material with semantic value that might be mapped and combined with other candidates to construct the knowledge base.

The main types of information that needed to be recorded in the XR4DRAMA KG were: general data about virtual reality experiments, visual analysis results from images and videos, and textual analysis results derived from online retrieved content.

The visual analysis component is off-the-shelf tool, which is result of some of our previous work (Batziou et al., 2023). In more detail, the visual analysis is a computer vision mechanism that can identify objects in an image or video, and also classify the image into a specific category which is called verge of the image. The Verge classifier was used to assign the photos (or video frames) to one or more classes depending on context (Andreadis et al., 2020), while the model conducting semantic segmentation (Qiu et al., 2021) on images was trained to extract semantic labels and percentages per pixel on images.

Since it was decided not to deep copy structures from a SOLR instance[5], which do not serve any requirements, only a small number of the generated as-

[3] https : / / xr4drama.eu / wp - content / uploads / 2021 / 12 / d3.5_xr4drama_semanticrepresentationfusiondss_20211201_v1.2.pdf

[4] https://github.com/valexande/xr4drama-icaart-paper

[5] https://solr.apache.org/

sets were chosen to be included in the XR4DRAMA KG. These assets include the text itself, the sentences that make up the text, and the named entity relationships that may be found there.

3.2 A Knowledge Graph for Media Planning

In this subsection, we give a high-level overview of the KG schema's structure (i.e., the XR4DRAMA ontology) and the guiding principles of each class. You can find the KG and the programs created to populate it here[4]. A high-level overview of the main XR4DRAMA ontology classes is shown in Figure 2.

- **InformationOfInterest:** The fundamental entities of interest to aid in decision-making.
- **Location:** The location of an event is represented by this class. It may be displayed with the location's name or with its coordinates (e.g., Vicenza).
- **Metadata:** All the supplementary data that is provided with the analysis results, and can be used in the decision-making process is referred to as metadata.
- **MultimediaObject:** Indicates the type of the transmission that is used to transfer information, it can be either Audio, Textual, or Video.
- **Observation:** This class, which is utilized by each individual component, describes the method of assessing the date.
- **Result:** This class, represents information about the outcome of an observation.
- **Procedure:** This class, represents information about the procedures that should be taken during an observation (i.e., tasks that should be performed).
- **Project:** Each observation is described in this class, along with some relevant data.
- **RiskReport:** This class describes the total outcome of all risk levels derived from various components. We added this class as some locations can be risky to access due to various reasons, even in a media planning scenario.
- **User:** Users are journalists, and each journalist is given a unique ID.

We also analyze the purpose of the various *object type* properties, i.e., properties that connect instances from one class with instances from another class.

- **hasMultimedia:** This property identifies if the observation was provided via a textual, video, or image post.
- **hasResult:** This property shows how the event in the observation turned out.
- **usedProcedure:** This property identifies the method used to extract the information from the observation, such as whether the information was taken from the visual or textual analysis component.
- **hasInformationOfInterest:** This property identifies the most important data in an observation, such as the type of scene, for example *airfield*, and the recognized objects, such as *building, automobile*, etc. Domain experts suggested what was deemed important information.
- **hasMetadata:** Based on the type of the observation, this property indicates the observation's metadata. This property indicates the metadata of the observation, based on the type of the observation. If the observation is a result of the: (a) visual analysis component (see Table 1), and (b) textual analysis component (see Table 3).
- **hasLocation:** This property shows where the important information in an observation is located. Observe that the location is specified using latitude and longitude.
- **consistsOf:** This property lists the observations that make up a project. A project is a collection of observations with neighboring coordinates.
- **hasProjectLocation:** This property indicates the location of the project.
- **includes:** This property indicates the name of the user that created the project.
- **hasUserLocation:** This property indicates the location of the user.
- **hasRiskReport:** This property indicates the risk level of the project, i.e., if it is an emergency or not.

3.3 Point of Interest Creation and Update

In order to make it easier for journalists and media organizations to complete a remote production mission, POIs aim to create some points in a region (i.e., pins on a map) that convey vital geographical information. Any user can add or modify POIs using a phone application (which is not currently public). The user can either provide a picture or video that the visual analysis component analyzes, and some of the important data in the image or video is then passed to a new or existing POI (the pipeline for a textual message is similar). Notice that in Figure 1, we also show that

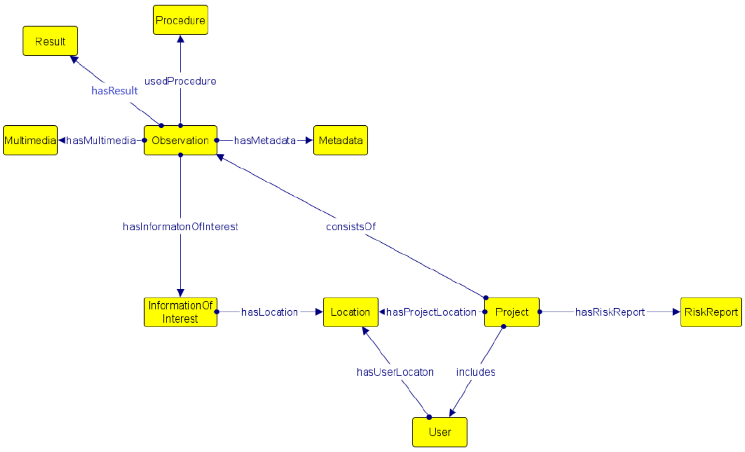

Figure 2: XR4DRAMA KG high level illustration.

the POIs receive information from the KG. Here we analyze only the information from the messages, as the information from the KG refers to some IDs that relate to the projects and the observations (see subsection 3.2).

It is simple to comprehend why a POI would need to be formed: if an event had taken place and there were none already present in the region. The updating of POIs, on the other hand, takes place when there are already POIs in the region and part of the information in them needs to be updated since the event's state has changed. For instance, the area has become crowded. The data from a visual message that is sent to a POI during creation (see Table 1) or updating (see Table 2) is shown below.

Table 1: Information passed from a visual message to a POI when created.

Label	**Value**	**Example**
category	string	Education
subcategory	string	Universities
current_user	string	journalist_1
objectsDetected	list of strings	[cabinet,chair]
sceneRecognition	string	theater
type	string	Point
coordinates	list of floats	[11.55,45.54]

Table 2: Information passed from a visual message to a POI when updated.

Label	**Value**	**Example**
objectsDetected	list of strings	[cabinet,chair]
sceneRecognition	string	theater

One can notice that when a POI is created the information passed from the visual analysis messages are: (i) what objects have been recognized, and (ii) the label of the scene (i.e., the scene is the verge classification of the image see subsection 3.1). Since some of the aforementioned data may be dynamic, the POI will still be constructed even if some are absent. The *category* and *subcategory* characterize the area which was recognized, in our running example an Education area, and more particularly, a University was recognized. In addition, the *current_user* is the name of the user who sent the message. Last but not least, the *coordinates* are also given in the format: [longitude, latitude]. The current user, the category, the subcategory, and the location are required pieces of information. However, only a limited amount of information can be altered if a POI already exists. The data which can be updated are: (i) and (ii).

We also analyze the information from a textual message that is passed to a POI when is created (see Table 3) or updated (see Table 4).

Table 3: Information passed from a textual message to a POI when created.

Label	**Value**	**Example**
category	string	Education
subcategory	string	Universities
current_user	string	journalist_1
sourceText	string	the theater has become crowded
objectsDetected	list of strings	[cabinet, man]
label	string	theater
type	string	Point
coordinates	list of floats	[11.55,45.54]

Similarly when a POI is created the information passed from the textual analysis messages are: (i)

Table 4: Information passed from a textual message to a POI when updated.

Label	Value	Example
sourceText	string	the theater has become crowded
objectsDetected	list of strings	[cabinet, man]
label	string	theater

which are the detected objects, (ii) an auxiliary label that characterizes the location, and (iii) the source text of the textual message. The aforementioned data can be dynamic, meaning that even if some are missing the POI will still be created. The necessary data is the current_user, the category, the subcategory, and the coordinates. On the other hand, if a POI already exists only some information can be updated. The data which can be updated are: (i-iii).

4 EVALUATION

The XR4DRAMA KG was evaluated with two separate methods. On the one hand, we analyzed the XR4DRAMA KG's consistency and completeness using two separate evaluation techniques. By developing a set of competency questions that the KG must be able to answer with the knowledge it contains (subsection 4.1), we evaluated the comprehensiveness of the XR4DRAMA KG. Then, we evaluated the XR4DRAMA KG's consistency by seeing if it adhered to a particular set of SHACL constraints (subsection 4.1). On the other hand, by calculating the precision-recall-F1 scores utilized in information extraction systems (subsection 4.2), the POI management mechanism was evaluated.

4.1 Completness and Consistency of the Knowledge Graph

Competency Questions (CQs) compiled during the creation of the official ontology requirements specification document (ORSD) were used to assess the completeness of the XR4DRAMA KG (Suárez-Figueroa et al., 2009). For this reason, we asked a group of specialists to create a series of questions that they would like the XR4DRAMA KG to answer before we built it. The experts were authority workers from *Autorita' di bacino distrettuale delle alpi orientali*[6] and journalists from *Deutsche Welle*[7]. A total of 32 CQs were gathered, and we have included a sample of 10 of them in Figure 3. The full list of CQs may be accessed here[4].

1) What is the risk level of the observation?
2) Which is the emergency in the observation?
3) What is the detection/creation time of the observation?
4) Which is the area in the observation?
5) What is the probability of the area in observation?
6) Which is the Stress level of the between time intervals?
7) Which is the objects found in video?
8) Which is the most/least risky observation?
9) What is the multimedia type used in observation?
10) How many people are in danger between time intervals [$t1$]-[$t2$]?

Figure 3: Batch of Competency Questions.

The completeness of the XR4DRAMA KG was found adequate, as each CQ when translated into a SPARQL counterpart returned the desired information. For instance, the fourth CQ from Figure 3 when translated into a SPARQL counterpart (see Example 1), for the observation *VisualMetadata_2c60537511c240c9add7fb2eb4e7459e_0* returned *amphitheater*. If the observation is visual, the name will be created from the text *VisualMetadata_* (if not, *TextualMetadata_*) and a unique simmoid value.

Example 1. *SELECT DISTINCT ?area WHERE* {
xr:VisualMetadata_2c60537511c240c9add7 fb2eb4e7459e_0
xr:hasInformationOfInterest ?info .
?info xr:hasLocation ?location .
?location xr:hasArea ?area. }

In addition to the CQs, we carried out a validation process to examine the syntactic and structural quality of the KB's metadata and to verify their consistency. Custom SHACL consistency checking rules and native ontology consistency checking, such OWL DL reasoning, were used to adhere to the closed-world criterion. One can find constraint violations, such as cardinality inconsistencies, incomplete, or missing information, by employing the first method. By employing the latter, the terminological semantics, or TBox, are taken into account as validation, much like in the case of class disjointness. Out of 56 SHACL rules, 21 of which referenced to object type properties and 35 to data type properties, the consistency of the XR4DRAMA KG was deemed sufficient because none of them returned any rule invalidation. We also looked for instances that belong to the intersection of classes because we did not want that to happen, but none were found.

[6]http://www.alpiorientali.it/
[7]https://www.dw.com/en/news/s-30701

4.2 POI Management Mechanism Evaluation

The standard precision, recall, and F1-score used for information extraction systems (Equations 1, 2 and 3) were applied to the evaluation of the POI management mechanism in order to create or update POIs from visual and textual messages. The POI management mechanism can be regarded as an information extraction mechanism, as a query is received (in our case a message from the textual or visual analysis components), and some information is extracted from the KG (a POI is created or updated).

$$precision = \frac{|\{RelevantInstance\} \cap \{RetrievedInstance\}|}{|\{RetrievedInstance\}|} \quad (1)$$

$$recall = \frac{|\{RelevantInstance\} \cap \{RetrievedInstance\}|}{|\{RelevantInstance\}|} \quad (2)$$

$$F1 = 2 * \frac{recall * precision}{recall + precision} \quad (3)$$

Retrieved Instances are considered all the visual (or textual) messages for which the POI management mechanism, *did not* return an error when we casted a message in order to create or update a POI.

Relevant Instances are considered all the the visual (or textual) messages for which the POI management mechanism, managed to create or update a POI, when we casted a message with them.

The intuition behind the retrieved and relevant instances, is that retrieved pairs from the moment that the POI management mechanism did not return any error they are capable of retrieving information (through a POI), while relevant are instance which managed to create or update a POI and therefore contain information relevant to a project.

The number of retrieved textual and visual messages are indicated by the variables $Retrieved_t$ and $Retrieved_v$, respectively. The numbers of relevant textual and visual messages are $Relevant_t$ and $Relevant_v$, respectively. Next, $recall_t$, $recall_v$, are the recall scores for the textual and visual messages, $F1_t$, $F1_v$ are the F1 scores for the textual and visual messages, respectively, and $precision_t$, $precision_v$ are the precision scores for the textual and visual messages.

One can find the dataset used to test our POI management mechanism here[4]. It consists of a set of 1501 text messages and 800 visual messages. Be aware that the values of each label in each message were chosen at random from a gold standard dataset assembled by domain experts, in order to tackle potential biases. It is interesting that all messages—textual or visual—were considered to have been successfully retrieved, which means that our POI management mechanism never returned an error for any given message, whether it was textual or visual. The resulting values are $Retrieved_t = 1501$ and $Retrieved_v = 800$. The same does not hold for the relevant messages, either textual or visual, as there were 1376 $Relevant_t$ messages for the textual analysis component and 697 $Relevant_v$ messages for the visual analysis component.

Table 5 contains the precision, recall, and F1 scores for both textual and visual messages based on the aforementioned data. The results are rounded to four decimals.

Table 5: Precision, Recall and F1-scores for textual and visual messages.

	Precision	Recall	F1
Textual Messages	0.917	1.0	0.956
Visual Messages	0.871	1.0	0.931

5 DISCUSSION AND CONCLUSION

The XR4DRAMA KG was founded with the goal of assisting journalists in managing and disseminating information regarding the condition of a location, so that they may provide the best coverage possible for events that took place there. In addition, the XR4DRAMA KG provides a innovative mechanism, the POI management mechanism that can create or update POIs, which contain vital geographical data required by journalists in order to have a clear view of the area and plan appropriately for the coverage of an event. The XR4DRAMA KG was constructed in order to work as a KG that would assist journalists in media planning.

Regarding the evaluation our goal was to examine the POI management mechanism using the precision-recall-F1 metrics for information extraction systems, as well as the completeness and consistency of the XR4DRAMA KG. CQs, which were gathered by experts, were used to assess the KG's completeness (subsection 4.1). More specifically, we converted each CQ into a SPARQL equivalent, and we anticipated that each one would return results. This is evidence that our KG may deliver significant information in a broader media planning scenario. A series of 56 SHACL restrict rules, of which 21 related to object type properties and 35 to data type properties, were used to test the consistency of the KG (subsection 4.1); none of them resulted in the rule being invalidated. Additionally, we looked to see whether there were any instances that belonged to the intersection of the classes, but none were found. This demonstrates the coherence of our KG, proving that it is free of noise and contradicting information.

Our POI management mechanism received strong F1 scores for both the visual (93.1%) and textual (95.6%) message, demonstrating that it can be utilized independently to create/update POIs in a broad media planning scenario. Additionally, we can comment that this occurred when updating POIs, which means that new information could not be added to POIs that already existed in the area, and were missing relevant instances for both textual and visual messages. The updated messages were straying outside of the bounding boxes of all existing POIs because each POI has a box around it. Be aware that the POIs' bounding boxes are part of a broader bounding box that encompasses the area that requires news coverage. It seems reasonable to take into account a bounding box for the POIs and the area that requires news coverage; otherwise, we risk adding POIs to the area that are situated in a completely other area of the map.

The high Recall scores—100% for both the visual and textual messages—can also be mentioned. This essentially indicates that there were no textual or visual messages that indicated an error. If we examine the two scenarios in which an error may be returned, the reason for not doing so is pretty clear: (i) The message's coordinates do not fall within a bounding box that designates the location where an event has occurred, or (ii) The user will identify a non-matching category-subcategory tuple. It is difficult for the user to choose the incorrect selection in both cases since the user sends messages using a mobile application (which is now private) that displays the permissible category-subcategory tuples, and the bounding box with a blue hue over an area.

In terms of future work, we intend to provide a method that will make the POIs more beneficial while making decisions. Additionally, we will provide POIs with a list of tasks that must be taken in order to complete a remote production mission more accurately.

ACKNOWLEDGEMENTS

This work has been funded by XR4DRAMA Horizon 2020 project, grant agreement number 952133.

REFERENCES

Andreadis, S., Moumtzidou, A., Apostolidis, K., Gkountakos, K., Galanopoulos, D., Michail, E., Gialampoukidis, I., Vrochidis, S., Mezaris, V., and Kompatsiaris, I. (2020). Verge in vbs 2020. In *International Conference on Multimedia Modeling*, pages 778–783. Springer.

Batziou, E., Ioannidis, K., Patras, I., Vrochidis, S., and Kompatsiaris, I. (2023). Low-light image enhancement based on u-net and haar. In *In Proceedings of the 29th International Conference on Multimedia Modeling (MMM 2023)*. Springer.

Berners-Lee, T., Hendler, J., and Lassila, O. (2001). The semantic web. *Scientific american*, 284(5):34–43.

Berven, A., Christensen, O. A., Moldeklev, S., Opdahl, A. L., and Villanger, K. J. (2018). News hunter: building and mining knowledge graphs for newsroom systems. *NOKOBIT*, 26:1–11.

Ehrlinger, L. and Wöß, W. (2016). Towards a definition of knowledge graphs. *SEMANTiCS (Posters, Demos, SuCCESS)*, 48(1-4):2.

Ni, J., Liu, X., Zhou, Q., and Cao, L. (2019). A knowledge graph based disaster storyline generation framework. In *2019 Chinese Control And Decision Conference (CCDC)*, pages 4432–4437. IEEE.

Opdahl, A. L., Al-Moslmi, T., Dang-Nguyen, D.-T., Gallofré Ocaña, M., Tessem, B., and Veres, C. (2022). Semantic knowledge graphs for the news: A review. *ACM Computing Surveys (CSUR)*.

Opdahl, A. L., Berven, A., Alipour, K., Christensen, O. A., and Villanger, K. J. (2016). Knowledge graphs for newsroom systems. *NOKOBIT*, 24:1–4.

Pan, J. Z., Pavlova, S., Li, C., Li, N., Li, Y., and Liu, J. (2018). Content based fake news detection using knowledge graphs. In *International semantic web conference*, pages 669–683. Springer.

Qiu, W., Li, W., Liu, X., and Huang, X. (2021). Subjective street scene perceptions for shanghai with large-scale application of computer vision and machine learning. Technical report, EasyChair.

Rospocher, M., van Erp, M., Vossen, P., Fokkens, A., Aldabe, I., Rigau, G., Soroa, A., Ploeger, T., and Bogaard, T. (2016). Building event-centric knowledge graphs from news. *Journal of Web Semantics*, 37:132–151.

Suárez-Figueroa, M. C., Gómez-Pérez, A., and Villazón-Terrazas, B. (2009). How to write and use the ontology requirements specification document. In *OTM Confederated International Conferences" On the Move to Meaningful Internet Systems"*, pages 966–982. Springer.

Tang, J., Feng, Y., and Zhao, D. (2019). Learning to update knowledge graphs by reading news. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2632–2641.

Villazon-Terrazas, B., Ortiz-Rodriguez, F., Tiwari, S. M., and Shandilya, S. K. (2021). *Knowledge graphs and semantic web*. Springer.

Wang, Y. and Hou, X. (2018). A method for constructing knowledge graph of disaster news based on address tree. In *2018 5th International Conference on Systems and Informatics (ICSAI)*, pages 305–310. IEEE.