

VHAKG: A Multi-modal Knowledge Graph Based on Synchronized Multi-view Videos of Daily Activities

Shusaku Egami

National Institute of Advanced Industrial Science and Technology (AIST)
Koto-ku, Tokyo, Japan
s-egami@aist.go.jp

Takanori Ugai

Fujitsu Ltd.
Kawasaki-shi, Kanagawa, Japan
ugai@fujitsu.com
National Institute of Advanced Industrial Science and Technology (AIST)
Koto-ku, Tokyo, Japan

Swe Nwe Nwe Htun

National Institute of Advanced Industrial Science and Technology (AIST)
Koto-ku, Tokyo, Japan
swenwe.nwehtun@aist.go.jp

Ken Fukuda*

National Institute of Advanced Industrial Science and Technology (AIST)
Koto-ku, Tokyo, Japan
ken.fukuda@aist.go.jp

ABSTRACT

Multi-modal knowledge graphs (MMKGs), which ground various non-symbolic data (e.g., images and videos) into symbols, have attracted attention as resources enabling knowledge processing and machine learning across modalities. However, the construction of MMKGs for videos consisting of multiple events, such as daily activities, is still in the early stages. In this paper, we construct an MMKG based on synchronized multi-view simulated videos of daily activities. Besides representing the content of daily life videos as event-centric knowledge, our MMKG also includes frame-by-frame fine-grained changes, such as bounding boxes within video frames. In addition, we provide support tools for querying our MMKG. As an application example, we demonstrate that our MMKG facilitates benchmarking vision-language models by providing the necessary vision-language datasets for a tailored task.

CCS CONCEPTS

• **Computing methodologies** → **Knowledge representation and reasoning**; • **Information systems** → **Multimedia information systems**; **Semantic web description languages**.

KEYWORDS

Multi-Modal Knowledge Graph, Event-Centric Knowledge Graph, Synthetic Data, Daily Life Video, Visual Question Answering

ACM Reference Format:

Shusaku Egami, Takanori Ugai, Swe Nwe Nwe Htun, and Ken Fukuda. 2024. VHAKG: A Multi-modal Knowledge Graph Based on Synchronized Multi-view Videos of Daily Activities. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management (CIKM '24)*,

*Corresponding author



This work is licensed under a Creative Commons Attribution International 4.0 License.

CIKM '24, October 21–25, 2024, Boise, ID, USA
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0436-9/24/10
<https://doi.org/10.1145/3627673.3679175>

October 21–25, 2024, Boise, ID, USA. ACM, New York, NY, USA, 5 pages.
<https://doi.org/10.1145/3627673.3679175>

1 INTRODUCTION

Multi-modal knowledge graphs (MMKGs) [30], which ground various non-symbolic data to symbols, have attracted attention as a resource that enables knowledge processing across modalities. Typical MMKGs [10, 27] are knowledge graphs (KGs) in which images are grounded to entities in the graph. Grounding video content to entities in a KG requires solving subtasks such as event extraction, object extraction, and relation extraction, and various methods have been proposed. In the case of a long video consisting of multiple events, such as daily activities, the sequential events need to be extracted and listed by the timeline of the events. However, these tasks are difficult to solve using current methods [30]. To the best of our knowledge, there is no available MMKG of videos consisting of such event sequences. Furthermore, the bounding boxes of the objects within the video frames change dynamically since the information in the video changes frame by frame. It is desirable that the MMKGs of videos can effectively represent both visual changes between each frame and contextual changes between each event that is interpreted. Leveraging such MMKGs facilitates the construction of customized pre-training or test datasets for downstream tasks. This capability enables tailored data extraction for specific applications, such as extracting pairs of video frames and corresponding action labels before and after object state changes.

In this paper, we introduce a novel MMKG that integrates fine-grained events and frame-by-frame knowledge of daily activity videos. Specifically, we first generate multi-view synchronized daily activity videos using a virtual space simulator. Then, we construct an MMKG, called VirtualHome-AIST-KG (VHAKG), that represents event-centric and frame-by-frame knowledge, such as Figure 1, based on the generated videos and our designed ontology. Moreover, we compressed VHAKG to remove redundant triples since their data size is enormous and published it on the Web in a permanently available format. Finally, we provide a set of tools to facilitate the use of VHAKG even by users unfamiliar with the graph query language. In addition, as a possible use case of VHAKG, we introduce

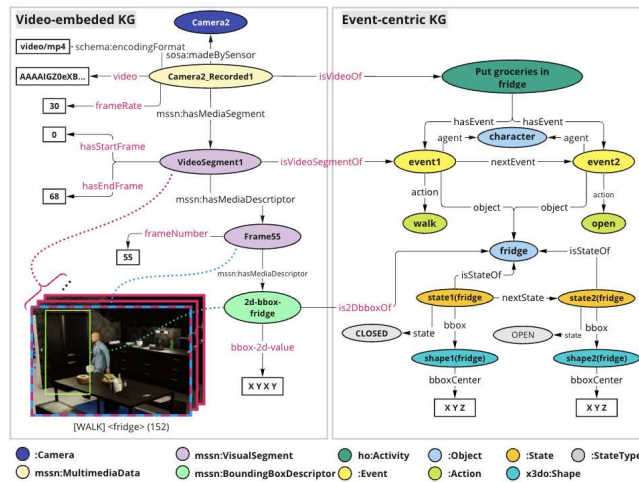


Figure 1: Illustration of VHAKG

the experiments that extract tailored test data sets for visual question answering (VQA) and use them to evaluate the performance of Large Vision-Language Models (LVLMs). The main contributions of this study are summarized as follows:

- **Novelty:** We constructed a novel MMKG that represents event sequences and frame-by-frame knowledge.
- **Availability:** We compressed our MMKG, which has a huge data size, and published it on the web in a permanently accessible format. Our MMKG is not subject to ethical review since it is artificial data.
- **Utility:** We provided support tools to use our MMKG and also introduced an example of creating a test set for evaluating LVLMs.
- **Predicted Impact:** Development of simulation-to-reality (sim2real) approaches [8, 12, 22] for daily life support and the publication of MMKGs of videos, similar to our study, are expected.

2 RELATED WORK

Zhu et al. [30] comprehensively surveyed and summarized previous works on MMKGs. Unfortunately, many MMKGs are not publicly available or are inaccessible. We focus on publicly available MMKGs whose entities (i.e., nodes) link directly to image or video files. *IMGPedia* [10] is an MMKG that grounds Wikimedia Commons images into *DBpedia* [2] entities. *MMpedia* [27] is an MMKG that matches entities corresponding to images retrieved from search engines. These MMKGs are still available because they are intended to share data using semantic web technologies. *VisionKG* [29] is an MMKG containing bounding boxes of objects extracted from various image datasets such as *MS-COCO* [16], *CI-FAR* [13], and *PASCAL VOC* [9]. The raw dataset is not publicly available at this time, but a useful interface is available. Our MMKG differs from these studies because it describes timelines of temporally fine-grained events in the videos and frame-by-frame bounding boxes.

Although Zhu et al. [30] mentioned that the extraction of sequential events from a long video containing multiple events has

not yet been addressed, the KGs of such events are constructed using different approaches from automatic event extraction. Vizcarra et al. [26] constructed a KG based on manual annotation data to videos. However, the KG is not publicly available because they focused on knowledge representation methods. In our previous work [8], we developed *VirtualHome2KG*, a framework for constructing KGs from fine-grained event data generated by *VirtualHome* [21] simulator. The *VirtualHome* simulator renders human activities and outputs environmental information based on the input program data. The program data consists of multiple action steps, as shown below.

```

Watch movie
Sit down on a couch in front of the TV. Use remote
to turn on the TV.
[WALK] <couch> (275)
[SIT] <couch> (275)
[GRAB] <remote_control> (1000)
...
[WATCH] <television> (297)
    
```

VirtualHome2KG structures the environmental information output by *VirtualHome* based on an ontology to construct event-centric KGs (EKGs [11]), where nodes are events and entities, while edges are event-event relations, event-entity relations, and entity-entity relations. Unlike a temporal knowledge graph (TKG) [4], which consists of triples with timestamps, EKG is represented as an edge-labeled directed graph. However, their entities are not directly linked to the video files because these EKGs focus on representing only the content of the videos. In addition, the frame-by-frame knowledge, such as objects' 2D bounding boxes in each image, is not represented. In contrast, this paper introduces a novel MMKG, which embeds synchronized video data captured by multiple cameras into entities and represents event sequences and frame-by-frame knowledge in videos.

3 DATASETS

We describe a method for generating data on daily activities, structuring them based on an ontology, and constructing an MMKG that is practically distributable on the web through data compression.

3.1 Data Generation

We generated a large number of simulated videos and EKGs of a wide variety of daily activities using *VirtualHome2KG*. In this step, we first added the following three new functions to the *VirtualHome* simulator and extended it to generate more diverse daily activity datasets frame-by-frame: (1) implementing renderable motions of various primitive actions, (2) automatically annotating 2D bounding boxes of objects in each video frame, and (3) adding synchronous camera mode with adjusted viewing angle and position. The original *VirtualHome* can factually render only about 10 primitive actions, and the generated daily activity videos are not diverse enough. Thus, we implemented 38 motions corresponding to various primitive actions, e.g., "eat," "pour," and "wipe."

In addition, we implemented a function to output the 2D bounding boxes of objects in the video frame every 5 frames. The 2D bounding box is automatically detected by ray casting from cameras inside the environment. Therefore, we can collect the ground truth used in computer vision tasks without any annotation work on the bounding box, which used to require much manual work.




Figure 2: Number of actions (red means new actions that became executable by this study)

Furthermore, we installed new fixed cameras on the diagonal of each room and increased the viewing angle of the cameras from 40 degrees to 70 degrees to capture the entire room situation. We released the new simulator with these improvements on the web as VirtualHome-AIST¹.

We manually created over 700 various daily activity scenarios (i.e., program data described in Section 2) with reference to the existing dataset [21] and simulated them using VirtualHome-AIST. Each activity was rendered simultaneously using 5 camera modes and generated over 3500 synchronized multi-view videos. Furthermore, we integrated VirtualHome-AIST into VirtualHome2KG and constructed EKGs of over 700 daily activities. These EKGs are integrated into video-embedded KGs described in Section 3.2.

Figure 2 shows the number of actions included in the generated video. An activity consists of an average of 10.2 events, and each event has one action.

3.2 MMKG Construction

Our MMKG is defined as a graph $\mathcal{G} = \{\mathcal{E}, \mathcal{R}, \mathcal{L}, \mathcal{T}\}$, where $\mathcal{E}, \mathcal{R}, \mathcal{L}$ are set of entities, relations, and literal values, and $\mathcal{T} \subseteq \mathcal{E} \times \mathcal{R} \times (\mathcal{E} \cup \mathcal{L})$ are sets of triples. A set of literal values $\mathcal{L} = \{\mathcal{L}_{\mathcal{K}}, \mathcal{L}_{\mathcal{M}}\}$ denotes that $\mathcal{L}_{\mathcal{K}}$ is the set of the KG’s literal values and $\mathcal{L}_{\mathcal{M}}$ is the set of multi-modal data.

We first designed the schema of VHAKG with reference to the existing ontologies. We interpreted the videos captured by multi-camera cameras installed in VirtualHome-AIST as sensor data. Thus, we reused MSSN-Onto [1], which is an extension of the Semantic Sensor Network Ontology [23] for multimedia content, and extended the VirtualHome2KG’s ontology. Figure 1 shows an example of the modeling of VHAKG to be constructed in this study. The right side of the figure shows the EKG constructed in Section 3.1. The left side of the figure is the KG described in this section. In the EKG, “Activity” corresponds to the whole video, and “Event” corresponds to “VideoSegment.” The 2D bounding box links to the corresponding “Object” of the EKG. The video data is embedded as a literal value encoded in base64. Consequently, video files are free of broken links and reference errors, ensuring their permanent availability and sharing.

¹https://github.com/aistairc/VirtualHome_aist

Table 1: Number of triples and data size of constructed KG. The round brackets mean the compression ratio. ImageKG is image-embedded KG, and VideoKG is video-embedded KG.

	# of triples	Size [GB]
ImageKG	134,945,485 (-)	62.0 (-)
VideoKG	131,786,665 (97.7%)	17.3 (27.9%)
VideoKG (compressed)	37,646,681 (27.9%)	12.5 (20.0%)

VirtualHome-AIST outputs a frame-by-frame image, a file recording the start and end frame numbers of each action step, and 2D bounding boxes every 5 frames. We integrate these data and transform them into a KG in RDF format based on the designed schema using Python scripts with RDFLib [25].

3.3 MMKG Compression

3.3.1 Video compression. We first created a KG with all images embedded every five frames; however, we found that the data size became too large to share practically. Thus, we used MPEG [14] compression to reduce the data size significantly. Each video frame entity has a frame number instead of having a base64 value, and the video entity has a base64 value for the entire compressed video. This allows the use of well-known tools such as FFmpeg [24] to recover and extract arbitrary frame images from VHAKG.

3.3.2 Removing redundant triples. Constructing a KG, as shown in Figure 1, will create redundant triples about the 2D bounding boxes. We reduced the number of entities and triples by referring to the previous entities if the current 2D bounding boxes have not changed since the previous frame.

3.3.3 Results. Table 1 shows the number of triples and data size of VHAKG. In image-embedded KG, each image was converted to JPEG with a quality of 90 using Pillow [5]. As a result, our approach significantly compressed the number of triples and data size, and made it possible to share them securely in a research data repository². VHAKG, consisting of the compressed video-embedded KG and EKG, is available at Zenodo³.

4 APPLICATIONS

4.1 Support tools

We have developed and released support tools (GUI and command-line)⁴ to make VHAKG accessible to users who are unfamiliar with SPARQL [6]. The GUI displays the specified videos and images and is provided as Docker Compose [7]. The back-end system automatically loads VHAKG by launching GraphDB [17] as a triple store on the local machine. The GUI executes template-based SPARQL queries that search for videos matching the conditions specified in the UI and restore and display the videos from their base64 values. The seek bar moves to the position of the specified frame. In addition, the command-line tool can extract videos and images and save them with annotation labels on a local machine.

²For example, Zenodo’s size limit is currently 50 GB.

³<https://doi.org/10.5281/zenodo.11438499>

⁴<https://github.com/aistairc/vhakg-tools>

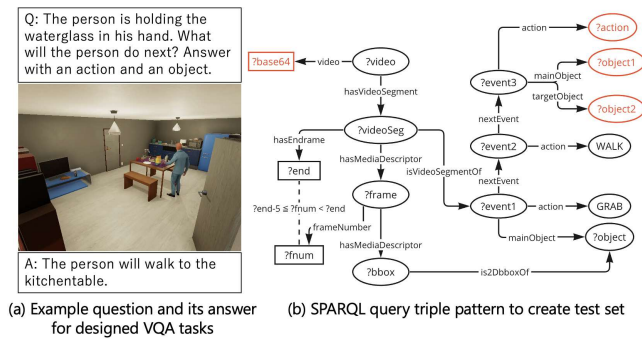


Figure 3: Example question and query pattern

4.2 Example Benchmarking of LVLMS

SPARQL querying enables users to extract videos, video frames, entities, and their labels from VHAKG. Thereby, users can create customized image and video annotation datasets for specific use cases. As a demonstration of VHAKG, we designed a new VQA task, created a test dataset, and conducted an example experiment.

4.2.1 Task design. We designed a task to understand a character’s daily activities from an input image and predict and explain the character’s next action. An example of input data is a single image and a question, as shown in Figure 3(a). This image is extracted from the daily activity video included in VHAKG. The models are required to understand the meaning of this input image and explain what action the character will take next. Therefore, this task is more difficult than existing vision-language tasks, such as caption generation.

4.2.2 Dataset preparation. The evaluation dataset should be able to predict the next action from the image features to some extent. Thus, we extracted the image immediately after the character grabbed something as the question data and extracted the next action and the target objects as the correct answer labels. Figure 3(b) shows the triple pattern of SPARQL queries to obtain these test sets. Red entities and literals are extracted. This triple pattern queries the event’s action following the event that the character grabs something and walks to somewhere, the object, the video, and the frame number. We then created short sentences based on the extracted actions and objects, as shown in Figure 3(a), based on a simple template. We extracted 100 pairs of data as the test set and another 5 pairs of data as samples for few-shot learning. Such data extraction is possible because VHAKG comprehensively represents video data, event-centric knowledge, and frame-by-frame knowledge.

4.2.3 Experiments. In our preliminary experiment, we extracted events after the “grab” event and found that 76 out of 100 events were “walk,” which was highly biased. In contrast, we reduced the bias in this experiment by extracting events after the “walk” following “grab,” as shown in Figure 4. In this way, it is possible to create a test set with reduced bias by querying VHAKG.

Note that the purpose of this section is to show that VHAKG can produce benchmark datasets for VQA tasks and that this study does not aim to develop or compare LVLMS. We show an example of an experiment to evaluate whether LVLMS can answer questions




Figure 4: Distribution of actions in test datasets

Table 2: Evaluation results of VQA task

Model	Method	BLEU	ROUGE-1	METEOR
GPT-4V	Zero-shot	0.249	0.354	0.271
	5-shot	0.470	0.544	0.484
GPT-4o	Zero-shot	0.0808	0.232	0.111
	5-shot	0.481	0.565	0.497

shown in Figure 3(a) by In-context Learning. We selected GPT-4V (gpt-4-1106-vision-preview [18]) and GPT-4o (gpt-4o-2024-05-13 [19]) as LVLMS. The evaluation was performed in zero-shot learning and 5-shot learning. All output token sizes were set to 50. We employed BLEU [20], ROUGE-1 [15], and METEOR [3] as evaluation metrics to calculate the similarity between the short sentences output and the correct answers. Table 2 shows the evaluation results.

As a result, GPT-4 and GPT-4o could hardly predict the correct answer in the zero-shot learning of this task. Therefore, these results imply that the new benchmark dataset created from VHAKG was an unknown set of images and labels that had not yet been trained on GPT-4V and GPT-4o. In contrast, the accuracy improved in few-shot learning. Therefore, these results imply that the extracted sample data are good examples for this task, and GPT-4V and GPT-4o can learn this context to some extent.

5 CONCLUSION

In this study, we provided VHAKG, which is a novel MMKG based on multi-view videos of daily activities consisting of multiple events. Moreover, we presented VirtualHome-AIST, support tools of VHAKG, and an example experiment. VHAKG integrated different modalities by embedding videos as literal values within the KG. By compressing videos and removing redundant triples, we reduced the data size and made it permanently available through a reliable research repository. Moreover, VHAKG can contribute to the creation of benchmark datasets and pre-training datasets for LVLMS since VHAKG represents both event sequences and frame-by-frame visual changes. The remaining issues in this study are increasing the variety of videos and linking to other KGs. In the future, we plan to generate multi-agent and egocentric videos and link VHAKG to real visual dataset KGs [28, 29] for sim2real tasks.

ACKNOWLEDGMENTS

This paper is based on results obtained from a project, JPNP20006, commissioned by the New Energy and Industrial Technology Development Organization (NEDO), and JSPS KAKENHI Grant Numbers JP22K18008 and JP23H03688.

REFERENCES

- [1] Chinnapong Angsachotmetee, Richard Chbeir, and Yudith Cardinale. 2020. MSSN-Onto: An ontology-based approach for flexible event processing in Multimedia Sensor Networks. *Future Generation Computer Systems* 108 (July 2020), 1140–1158. <https://doi.org/10.1016/j.future.2018.01.044>
- [2] Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary Ives. 2007. DBpedia: A Nucleus for a Web of Open Data. In *The Semantic Web*, Karl Aberer, Key-Sun Choi, Natasha Noy, Dean Allemang, Kyung-Il Lee, Lyndon Nixon, Jennifer Golbeck, Peter Mika, Diana Maynard, Riichiro Mizoguchi, Guus Schreiber, and Philippe Cudré-Mauroux (Eds.). Springer, Berlin, Heidelberg, 722–735. https://doi.org/10.1007/978-3-540-76298-0_52
- [3] Satanjeev Banerjee and Alon Lavie. 2005. METEOR: An Automatic Metric for MT Evaluation with Improved Correlation with Human Judgments. In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, Jade Goldstein, Alon Lavie, Chin-Yew Lin, and Clare Voss (Eds.). Association for Computational Linguistics, Ann Arbor, Michigan, 65–72. <https://aclanthology.org/W05-0909>
- [4] Borui Cai, Yong Xiang, Longxiang Gao, He Zhang, Yunfeng Li, and Jianxin Li. 2023. Temporal Knowledge Graph Completion: A Survey. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*. International Joint Conferences on Artificial Intelligence Organization, Macau, SAR China, 6545–6553. <https://doi.org/10.24963/ijcai.2023/734>
- [5] Jeffrey A. Clark. 2024. Pillow. <https://pillow.readthedocs.io/en/stable/index.html> Accessed: 2024-05-27.
- [6] World Wide Web Consortium. 2013. SPARQL 1.1 overview. (2013).
- [7] Docker Inc. 2024. Docker Compose overview. <https://docs.docker.com/compose/> Accessed: 2024-05-27.
- [8] Shusaku Egami, Takanori Ugai, Mikiko Oono, Koji Kitamura, and Ken Fukuda. 2023. Synthesizing Event-Centric Knowledge Graphs of Daily Activities Using Virtual Space. *IEEE Access* 11 (March 2023), 23857–23873. <https://doi.org/10.1109/ACCESS.2023.3253807>
- [9] Mark Everingham, Luc Van Gool, Christopher K. I. Williams, John Winn, and Andrew Zisserman. 2010. The Pascal Visual Object Classes (VOC) Challenge. *International Journal of Computer Vision* 88, 2 (June 2010), 303–338. <https://doi.org/10.1007/s11263-009-0275-4>
- [10] Sebastián Ferrada, Benjamin Bustos, and Aidan Hogan. 2017. IMGpedia: A Linked Dataset with Content-Based Analysis of Wikimedia Images. In *The Semantic Web – ISWC 2017*, Claudia d’Amato, Miriam Fernandez, Valentina Tamma, Freddy Lecue, Philippe Cudré-Mauroux, Juan Sequeda, Christoph Lange, and Jeff Heflin (Eds.). Springer International Publishing, Cham, 84–93. https://doi.org/10.1007/978-3-319-68204-4_8
- [11] Saiping Guan, Xueqi Cheng, Long Bai, Fujun Zhang, Zixuan Li, Yutao Zeng, Xiaolong Jin, and Jiafeng Guo. 2022. What is Event Knowledge Graph: A Survey. *IEEE Transactions on Knowledge and Data Engineering* (2022), 1–20. <https://doi.org/10.1109/TKDE.2022.3180362>
- [12] Abhishek Kadian, Joanne Truong, Aaron Gokaslan, Alexander Clegg, Erik Wilmans, Stefan Lee, Manolis Savva, Sonia Chernova, and Dhruv Batra. 2020. Sim2Real Predictivity: Does Evaluation in Simulation Predict Real-World Performance? *IEEE Robotics and Automation Letters* 5, 4 (2020), 6670–6677. <https://doi.org/10.1109/LRA.2020.3013848>
- [13] Alex Krizhevsky, Geoffrey Hinton, et al. 2009. Learning multiple layers of features from tiny images. (2009).
- [14] Didier Le Gall. 1991. MPEG: A video compression standard for multimedia applications. *Commun. ACM* 34, 4 (1991), 46–58.
- [15] Chin-Yew Lin. 2004. ROUGE: A Package for Automatic Evaluation of Summaries. In *Text Summarization Branches Out*. Association for Computational Linguistics, Barcelona, Spain, 74–81. <https://aclanthology.org/W04-1013>
- [16] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. 2014. Microsoft COCO: Common Objects in Context. In *Computer Vision – ECCV 2014*, David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars (Eds.). Springer International Publishing, Cham, 740–755. https://doi.org/10.1007/978-3-319-10602-1_48
- [17] Ontotext USA, Inc. 2024. Ontotext GraphDB. <https://www.ontotext.com/products/graphdb/> Accessed: 2024-05-27.
- [18] OpenAI. 2023. New models and developer products announced at DevDay. <https://openai.com/index/new-models-and-developer-products-announced-at-devday/> Accessed: 2024-05-27.
- [19] OpenAI. 2024. Hello GPT-4o. <https://openai.com/index/hello-gpt-4o/> Accessed: 2024-06-03.
- [20] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a Method for Automatic Evaluation of Machine Translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, Pierre Isabelle, Eugene Charniak, and Dekang Lin (Eds.). Association for Computational Linguistics, Philadelphia, Pennsylvania, USA, 311–318. <https://doi.org/10.3115/1073083.1073135>
- [21] Xavier Puig, Kevin Ra, Marko Boben, Jiaman Li, Tingwu Wang, Sanja Fidler, and Antonio Torralba. 2018. VirtualHome: Simulating Household Activities Via Programs. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, Salt Lake City, UT, 8494–8502. <https://doi.org/10.1109/CVPR.2018.00886>
- [22] Yue Qiu, Yoshiki Nagasaki, Kensho Hara, Hirokatsu Kataoka, Ryota Suzuki, Kenji Iwata, and Yutaka Satoh. 2023. VirtualHome Action Genome: A Simulated Spatio-Temporal Scene Graph Dataset with Consistent Relationship Labels. In *2023 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*. IEEE, Waikoloa, HI, USA, 3340–3349. <https://doi.org/10.1109/WACV56688.2023.00335>
- [23] Kerry Taylor, Armin Haller, Maxime Lefrançois, Simon Cox, Raul Garcia-Castro, Danh Le-Phuoc, Joshua Lieberman, and Claus Stadler. 2019. The Semantic Sensor Network Ontology, Revamped. *Proceedings of the Journal Track co-located with the 18th International Semantic Web Conference (ISWC 2019)* (2019), 1–9.
- [24] FFmpeg Team. 2000. FFmpeg. <https://ffmpeg.org/> Accessed: 2024-05-27.
- [25] RDFLib Team. 2023. rdflib 7.0.0 – rdflib 7.0.0 documentation. <https://rdflib.readthedocs.io/en/stable/> Accessed: 2024-05-27.
- [26] Julio Vizcarra, Satoshi Nishimura, and Ken Fukuda. 2021. Ontology-based human behavior indexing with multimodal video data. *2021 IEEE 15th International Conference on Semantic Computing (ICSC)* (March 2021), 262–267. <https://doi.org/10.1109/ICSC50631.2021.00052> ISSN: 2325-6516.
- [27] Yanan Wu, Xiaowei Wu, Junwen Li, Yue Zhang, Haofen Wang, Wen Du, Zhidong He, Jingping Liu, and Tong Ruan. 2023. MMpedia: A Large-Scale Multi-modal Knowledge Graph. In *The Semantic Web – ISWC 2023*, Terry R. Payne, Valentina Presutti, Guilin Qi, María Poveda-Villalón, Giorgos Stoilos, Laura Hollink, Zoi Kaoudi, Gong Cheng, and Juanzi Li (Eds.). Springer Nature Switzerland, Cham, 18–37. https://doi.org/10.1007/978-3-031-47243-5_2
- [28] Yasunori Yamamoto, Shusaku Egami, Yuya Yoshikawa, and Ken Fukuda. 2023. Towards Semantic Data Management of Visual Computing Datasets: Increasing Usability of MetaVD. In *Proceedings of the ISWC 2023 Posters, Demos and Industry Tracks: From Novel Ideas to Industrial Practice co-located with 22nd International Semantic Web Conference (ISWC 2023)*.
- [29] Jicheng Yuan, Anh Le-Tuan, Manh Nguyen-Duc, Trung-Kien Tran, Manfred Hauswirth, and Danh Le-Phuoc. 2024. VisionKG: Unleashing the Power of Visual Datasets via Knowledge Graph. In *The Semantic Web*, Albert Meroño Peñuela, Anastasia Dimou, Raphaël Troncy, Olaf Hartig, Maribel Acosta, Mehwish Alam, Heiko Paulheim, and Pasquale Lisena (Eds.). Springer Nature Switzerland, Cham, 75–93. https://doi.org/10.1007/978-3-031-60635-9_5
- [30] Xiangru Zhu, Zhixu Li, Xiaodan Wang, Xueyao Jiang, Penglei Sun, Xuwu Wang, Yanghua Xiao, and Nicholas Jing Yuan. 2024. Multi-Modal Knowledge Graph Construction and Application: A Survey. *IEEE TRANSACTIONS ON KNOWLEDGE AND DATA ENGINEERING* 36, 2 (2024).