

---

# Fantastic Data and How to Query Them\*

---

**Trung-Kien Tran**

Bosch Center for Artificial Intelligence,  
Renningen, Germany

**Anh Le-Tuan, Manh Nguyen-Duc, Jicheng Yuan and Danh Le-Phuoc**

Open Distributed Systems,  
Technical University of Berlin

## Abstract

It is commonly acknowledged that the availability of the huge amount of (training) data is one of the most important factors for many recent advances in Artificial Intelligence (AI). However, datasets are often designed for specific tasks in narrow AI sub areas and there is no unified way to manage and access them. This not only creates unnecessary overheads when training or deploying Machine Learning models but also limits the understanding of the data, which is very important for data-centric AI. In this paper, we present our vision about a unified framework for different datasets so that they can be integrated and queried easily, e.g., using standard query languages. We demonstrate this in our ongoing work to create a framework for datasets in Computer Vision and show its advantages in different scenarios. Our demonstration is available at <https://vision.semkg.org>.

## 1 Vision

**Motivation** The field of Artificial Intelligence (AI) advanced significantly in recent years and most of the breakthroughs are data-intensive AI or its combinations with other techniques. Therefore, it is reasonable to argue that to make AI better, we need to have better ways to create (training) data and make it available to be consumed by AI models. In different AI areas, e.g., Computer Vision (CV) and Natural Language Processing (NLP) etc., and even within the same area, datasets are organised differently, e.g. stored in different formats and using different labels. This hinders not only the advancements in individual AI area but also the path towards Artificial General Intelligence (AGI). Let us make an analogy with the developments of human intelligence (HI) to demonstrate the importance of sharing data and knowledge. To advance in sciences and technologies, over thousands of years, we have accessed all kinds of knowledge and data, and this is because we have (written) languages. Based on previous discoveries, new ones were discovered and recorded in shareable forms. Similarly, it is commonly acknowledged that the current modern AI era was triggered by the availability of huge amount of data. For example, ImageNet [5] helps to showcase the success of Deep Learning in CV, many datasets created from Wikipedia and other websites provide input for NLP models; although language models like GPT [31] or BERT [6] do not need labelled data, it is still required in many downstream tasks. These datasets are so fragmented that even there exist algorithms/architectures for AGI models, there is no data to train them. This calls for the need of a unified framework to share and access datasets.

**General Ideas** We propose a unified framework for datasets in data-centric AI. Within this framework, not only different datasets in one AI area but also those in different AI areas are integrated

---

\*NeurIPS Data-Centric AI Workshop

and linked together. Existing resources such as ConceptNet [35], Wikidata [39], have a similar goal in a sense that they integrate data from different sources and make them shareable, but they rather focus on some specific application domains, and particularly, none of them is connected with datasets in other areas such as CV. On the other hand, in CV areas, so-called *scene graphs* were introduced to model the relationship between detected objects in images [12]. However, they lack the cross-domain interoperability and cannot be queried via standard query languages. Therefore, we argue that these existing resources should be unified and new-coming resources should easily be integrated. We believe that it is very beneficial, for example, it can help to avoid distribution shift<sup>2</sup>, create more robust models for training and testing [14]. Also, the shift from model tweaking to deep understanding of data furthermore requires that datasets need to be better organised [33]. Additionally, we believe that the term “data” should be extended to contain not only training data but also abstract knowledge, e.g., commonsense or causal relations [34].

In the next sections, we present our ongoing work to demonstrate a unified framework for datasets in CV in which images annotations are described in a knowledge graph (KG) and labels are linked to the well-known knowledge base Wikidata so that we can interlink the annotation labels across label spaces in different datasets under shared semantics. We then describe our future work and concrete steps to realise our vision.

## 2 A Case Study of Vision Knowledge Graph

**The Vision Knowledge Graph** To realise the ideas outlined above, we started to build a unified knowledge graph, namely VisionKG [18], for CV datasets (e.g. COCO [20], KITTI [8] or Visual Genome [15]). VisionKG is a Resource Description Framework (RDF) based knowledge graph [10] that contains RDF statements [26] describing the metadata of the images and the semantic of their annotations. RDF is a standardised data model recommended by the W3C particularly for semantic data integration and as a formal representation for shared human-machine understanding. Therefore, RDF can be used to represent many semantic structures of popular label taxonomies such as Wordnet [7], ConceptNet [35], and Freebase [3] which are used in many CV datasets, e.g., Imagenet [16], OpenImage [17] and VisualGenome [15].

Figure 1 illustrates the process of creating VisionKG. We first collect CV datasets and extract their annotation labels ①. To create a unified data model for the annotation labels and visual features, we follow the Linked Data principles [2] and use the RDF data model. The data entities (e.g. images, boxes, labels) are named using Uniform Resource Identifiers (URI). RDF data model allows the data to be expressed using *triples* of the form  $\langle \text{subject}, \text{predicate}, \text{object} \rangle$ . For example, to describe “an image contains a bounding box for a person” in COCO dataset, we first assign unique URIs, e.g., *vision.semkg.org/img01* and *vision.semkg.org/box01*, for the image and the bounding box, respectively to create the following triples for such image:  $\langle \text{img01}, \text{hasBox}, \text{box01} \rangle, \langle \text{box01}, \text{hasObject}, \text{obj01} \rangle, \langle \text{obj01}, \text{rdf:type}, \text{Person} \rangle$ . The predicates *hasBox*, *hasObject*, and *rdf:type* are predefined, in which *rdf:type* is used to express an object/image belongs to a specific class/type, e.g., *Person*, in the knowledge base; for simplicity, we skip the prefix in these triples. Furthermore, we add metadata and semantic annotations for the images, e.g., where the images come from or what are the relations of the boxes in an images (Figure 1 ②). The earlier provides the semantic relationships to facilitate semantic reasoning capability in below.

**Unifying and Expanding Labels** Since labels, attributes, and relationships in COCO, KITTI, and Visual Genome datasets are either just text or mapped to WordNet [7], we link them to the corresponding predicates and classes in Wikidata.<sup>3</sup> Wikidata is an open knowledge graph commonly used in other application domains, and therefore integrating datasets via Wikidata make them available also to other domains. Another advantage is that we can utilise the existing class hierarchy as shown in Figure 2 to add more labels to existing datasets using a semantic *reasoner* to expand/materialise the labels. For example, a box that is labelled as *a pedestrian* is also annotated as *a person* (Figure 2 ①) because in the hierarchy of the knowledge base, *pedestrian* is a subclass of *person* (Figure 2 ②). Hence, our knowledge graph can interlink the annotation labels across label spaces under shared semantic understanding. Along with the semantic relationships, thanks to the URI representation of image instances, different types of annotations for different learning tasks can be linked together

<sup>2</sup><http://ai.stanford.edu/blog/wilds/>

<sup>3</sup><https://www.wikidata.org/>




Figure 1: The overview of VisionKG

via the URI of the images. For example, COCO [20], COCO-Stuff [4], and VisualGenome [15] share many common images for different annotations in object detection, semantic segmentation, and visual relationship detection.

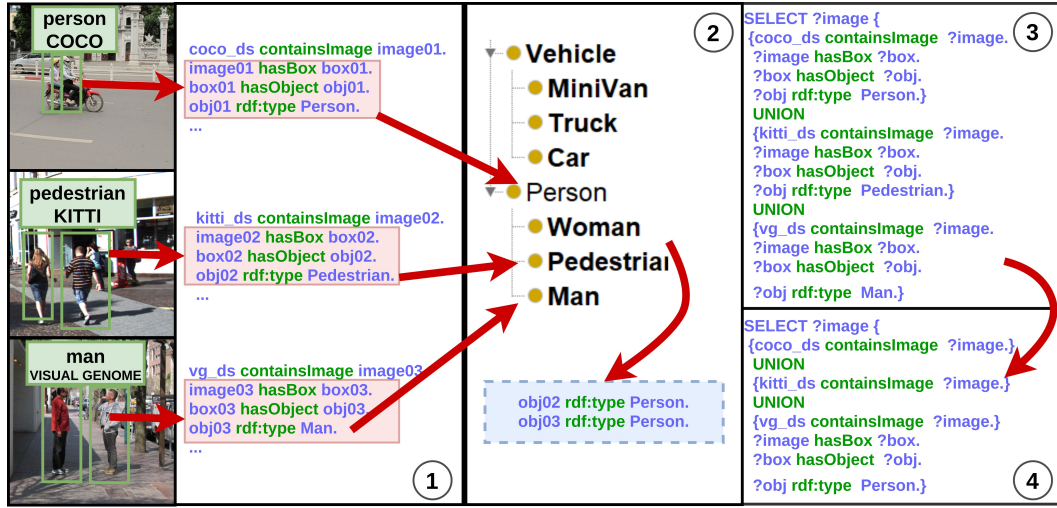


Figure 2: Mapping labels in COCO, KITTI, and Visual Genome to classes in the knowledge base (1). Expanding the labels according to the class hierarchy (2). And examples of two equivalent queries to obtain images that contain Person from COCO, KITTI, and Visual Genome datasets (3,4).

**Use Cases** VisionKG enables a more effective way to organise training data and offers more robust ways to analyse and evaluate trained Deep Neural Networks (DNNs). Specifically, datasets and analysis can be done using rich semantic query languages such as SPARQL.<sup>4</sup> The SPARQL query language provides users the ability to describe queries using RDF statements which are similar to those in the SQL language. The first use case is that VisionKG can be used to obtain mixed-datasets in an elegant way. For instance, one can query for images of *person* from COCO, KITTI, and Visual Genome using a simple query (Figure 2 ④) instead of the more complex query (Figure 2 ③) that covers all possible cases: *pedestrian* in KITTI or *man* in Visual Genome. This is possible because we already aligned labels of existing datasets with the common taxonomy and expanded them according to the taxonomy hierarchy. Similarly, one can conveniently create mixed-test data. In advanced settings, users can have more fine-grained criteria for retrieving images such as the query for “*images that contain a person holding a cat*” shown in Figure 1 ①. Other use cases include analytic queries to have deeper understanding about the dataset and performances of DNNs on some particular sub-set of the data. Figure 1 ② illustrates a complex query to “*search for the trained models to detect Car on*

<sup>4</sup><https://www.w3.org/TR/rdf-sparql-query/>

*images of Car in crowded traffic scenes or in mountain areas*". More example SPARQL queries can be found in our online VisionKG system.<sup>5</sup>

### 3 Towards a Unified Framework for Querying Data on The Fly

Our case study shows how the data can be organised in a unified way for both training and test phases. This approach advocates for an AI strategy that shifts from emphasis on proprietary datasets, to the sharing of data across entities for knowledge creation, namely building a knowledge graph guiding learning algorithms, called **Learning Knowledge Graph**  $\mathcal{L}^{KG}$ . Such a strategy leads us to look towards building a unified framework to facilitate the ability to query data in a declarative fashion for data loading, training, validation and test phases of Machine Learning pipelines. These pipelines are not restricted to neural networks but also learning approaches such as probabilistic logic programming [25] or a fusion of various learning algorithms, e.g. [29]. Such a framework with  $\mathcal{L}^{KG}$  will enable a learning algorithm to dynamically retrieve data programmatically. Hence, it does not have to assume training/validation/test sets to be fixed. This might lead to more robust models for a set of learning architectures. In particular,  $\mathcal{L}^{KG}$  can give such a training algorithm the access to a much bigger set of training samples than those in a traditional training pipeline. For example, we are planning to integrate into  $\mathcal{L}^{KG}$  a larger set of classification samples than those in the currently popular datasets such as 14 millions images of ImageNet [5] or even 300 millions images of JFT-300M [37] with probably open label spaces, e.g. [42, 14, 11]. Note that  $\mathcal{L}^{KG}$  will also consider unsupervised, self-supervised and weakly-supervising learning algorithms, hence, bigger sets of images such as Youtube-8M [1], YFCC100M [38] and LAION-400-MILLION<sup>6</sup> are also to be integrated. In this light, it will be much more challenging to build a learning architecture as it has to deal with ever-growing training/validation/test sets.

Towards this vision, instead of focusing on learning architectures, i.e., writing new models instead of understanding the nature of the learning tasks and its semantic relationships to the potential data being learnt, we focus on how to represent semantic relations, "know-how", context and domain knowledge, then make them queryable so that the training algorithms can exploit such a human-machine understandable knowledge to build the desired models by just specifying the expected outputs in a declarative fashion. For instance, to build a model for scene understanding,  $\mathcal{L}^{KG}$  can specify how an image scene can be represented as a scene graph constructing from edges (e.g., RDF statements) representing relationships among objects which can be learnt via the models to detect visual relationships, e.g [22], called VRD models. A training algorithm to build such VRD models will have to rely on datasets such as Visual Genome [15] or VrR-VG[19] which have various semantic relationships with datasets for training object detection/segmentation and classifications such as COCO and ImageNet. Interestingly, VisionKG [18] illustrated that such relationships can also be connected to natural language data sources such as Wordnet [27] and Conceptnet [35], Wikidata [39], and Freebase [3]. Along the same line, various datasets can help to facilitate knowledge embedding associated with natural language ones such as CLIP [30], VisualComet [28] and VCR [41]. On the other hand, there are many text-based datasets that can be enriched with visual data such as [36], [9] and [21]. To this end, the next challenge for our framework is how to leverage such rich correlated information among datasets and learning tasks to automate the training algorithms to make it faster, more efficient and more robust in building AI component powered by  $\mathcal{L}^{KG}$ .

Our such knowledge-centric approach hints that companies and organizations should focus on preparing knowledge-driven AI development pipeline. This shift shapes organizations' HR strategy on how to build an AI team that embrace the potential capabilities of  $\mathcal{L}^{KG}$ . While some companies will still require hiring large cohorts of rare and expensive data engineers and scientists, knowledge engineers can offer alternative resources to scale up AI development pipelines. The joint force among them can leverage the knowledge discovery in such pipelines to automate various phases of the development. For instance, the work such as Taskonomy [40] and Taskology [23] can facilitate the processes of building taxonomies/ontologies of the learning tasks so that the relationships among them can be exploited to build more efficient multi-task learning pipelines. Moreover, our knowledge-driven approach will lean towards a similar automation approach of the Software 2.0 [32], DynaBench [13], DynaBoard [24] whereby the knowledge engineers and domain experts can specify application specifications, tests and benchmarks along with domain knowledge as the input to automate the AI

<sup>5</sup><https://vision.semkg.org>

<sup>6</sup><https://laion.ai/laion-400-open-dataset/>

development phases, such as data cleaning, data integration, label creation and model search, model selection and model testing.

## Acknowledgements

This work was funded by the German Research Foundation (DFG) under the COSMO project (ref. 453130567), the German Ministry for Education and Research via The Berlin Institute for the Foundations of Learning and Data (BIFOLD, ref. 01IS18025A and ref. 01IS18037A), and the German Academic Exchange Service (DAAD, ref. 57440921). We thank Jannik Strötgen for his valuable feedback on this work.

## References

- [1] Sami Abu-El-Haija, Nisarg Kothari, Joonseok Lee, Paul Natsev, George Toderici, Balakrishnan Varadarajan, and Sudheendra Vijayanarasimhan. 2016. YouTube-8M: A Large-Scale Video Classification Benchmark. *arXiv e-prints* (2016).
- [2] Christian Bizer, Tom Heath, and Tim Berners-Lee. 2011. Linked data: The story so far. In *Semantic services, interoperability and web applications: emerging concepts*. IGI global, 205–227.
- [3] Kurt Bollacker, Robert Cook, and Patrick Tufts. 2007. Freebase: A shared database of structured general human knowledge. In *AAAI*, Vol. 7. 1962–1963.
- [4] Holger Caesar, Jasper Uijlings, and Vittorio Ferrari. 2018. Coco-stuff: Thing and stuff classes in context. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1209–1218.
- [5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*. 248–255.
- [6] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [7] Christiane Fellbaum. 2010. WordNet. In *Theory and applications of ontology: computer applications*. 231–243.
- [8] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. 2013. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research* (2013).
- [9] C. González, N. Ayobi, I. Hernández, J. Hernández, J. Pont-Tuset, and P. Arbeláez. 2021. Panoptic Narrative Grounding. *arXiv e-prints* (2021).
- [10] Aidan Hogan, Eva Blomqvist, Michael Cochez, Claudia d’Amato, Gerard De Melo, Claudio Gutierrez, Sabrina Kirrane, José Emilio Labra Gayo, Roberto Navigli, Sebastian Neumaier, et al. 2021. Knowledge graphs. *ACM Computing Surveys (CSUR)* 54, 4 (2021), 1–37.
- [11] Rui Huang and Yixuan Li. [n.d.]. MOS: Towards Scaling Out-of-Distribution Detection for Large Semantic Space. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*. 8710–8719.
- [12] Justin Johnson, Ranjay Krishna, Michael Stark, Li-Jia Li, David Shamma, Michael Bernstein, and Li Fei-Fei. 2015. Image retrieval using scene graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 3668–3678.
- [13] Douwe Kiela, Max Bartolo, Yixin Nie, Divyansh Kaushik, Atticus Geiger, Zhengxuan Wu, Bertie Vidgen, Grusha Prasad, Amanpreet Singh, Pratik Ringshia, Zhiyi Ma, Tristan Thrush, Sebastian Riedel, Zeerak Waseem, Pontus Stenetorp, Robin Jia, Mohit Bansal, Christopher Potts, and Adina Williams. 2021. Dynabench: Rethinking Benchmarking in NLP. *CoRR* abs/2104.14337 (2021).

- [14] Pang Wei Koh, Shiori Sagawa, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanus Phillips, Irena Gao, Tony Lee, et al. 2021. Wilds: A benchmark of in-the-wild distribution shifts. In *International Conference on Machine Learning*. 5637–5664.
- [15] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. 2017. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision* 123, 1 (2017), 32–73.
- [16] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. *NeurIPS* (2012).
- [17] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, Tom Duerig, and Vittorio Ferrari. 2020. The Open Images Dataset V4: Unified image classification, object detection, and visual relationship detection at scale. *IJCV* (2020).
- [18] Anh Le-Tuan, Trung Kien-Tran, Manh Nguyen-Duc, Jicheng Yuan, Manfred Hauswirth, and Le-Phuoc Danh. 2021. VisionKG: Towards A Unified Vision Knowledge Graph. In *Proceedings of the ISWC 2021 Posters & Demonstrations Track (CEUR Workshop Proceedings)*.
- [19] Yuanzhi Liang, Yalong Bai, Wei Zhang, Xueming Qian, Li Zhu, and Tao Mei. 2019. VrR-VG: Refocusing Visually-Relevant Relationships. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*. IEEE, 10402–10411.
- [20] 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 *ECCV*.
- [21] Yongfei Liu, Bo Wan, Lin Ma, and Xuming He. 2021. Relation-aware Instance Refinement for Weakly Supervised Visual Grounding. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*. Computer Vision Foundation / IEEE, 5612–5621.
- [22] Cewu Lu, Ranjay Krishna, Michael Bernstein, and Li Fei-Fei. 2016. Visual Relationship Detection with Language Priors. In *European Conference on Computer Vision*.
- [23] Yao Lu, Sören Pirk, Jan Dlabal, Anthony Brohan, Ankita Pasad, Zhao Chen, Vincent Casser, Anelia Angelova, and Ariel Gordon. 2020. Taskology: Utilizing Task Relations at Scale. *CoRR* abs/2005.07289 (2020).
- [24] Zhiyi Ma, Kawin Ethayarajh, Tristan Thrush, Somya Jain, Ledell Wu, Robin Jia, Christopher Potts, Adina Williams, and Douwe Kiela. 2021. Dynaboard: An Evaluation-As-A-Service Platform for Holistic Next-Generation Benchmarking. *CoRR* abs/2106.06052 (2021).
- [25] Robin Manhaeve, Sebastijan Dumancic, Angelika Kimmig, Thomas Demeester, and Luc De Raedt. 2021. Neural probabilistic logic programming in DeepProbLog. *Artif. Intell.* 298 (2021), 103504.
- [26] Frank Manola, Eric Miller, Brian McBride, et al. 2004. RDF primer. *W3C recommendation* 10, 1-107 (2004), 6.
- [27] George A Miller. 1995. WordNet: a lexical database for English. *Commun. ACM* 38, 11 (1995), 39–41.
- [28] Jae Sung Park, Chandra Bhagavatula, Roozbeh Mottaghi, Ali Farhadi, and Yejin Choi. 2020. Visualcomet: Reasoning about the dynamic context of a still image. In *European Conference on Computer Vision*. 508–524.
- [29] Danh Le Phuoc, Thomas Eiter, and Anh Lê Tuấn. 2021. A Scalable Reasoning and Learning Approach for Neural-Symbolic Stream Fusion. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021*. AAAI Press, 4996–5005.

- [30] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event, Marina Meila and Tong Zhang (Eds.)*.
- [31] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. (2018).
- [32] Christopher Ré. 2018. Software 2.0 and Snorkel: Beyond Hand-Labeled Data. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. Association for Computing Machinery, New York, NY, USA.
- [33] Christopher Ré, Feng Niu, Pallavi Gudipati, and Charles Srisuwananukorn. 2019. Overton: A Data System for Monitoring and Improving Machine-Learned Products. *CoRR* abs/1909.05372 (2019).
- [34] Maarten Sap, Ronan Le Bras, Emily Allaway, Chandra Bhagavatula, Nicholas Lourie, Hannah Rashkin, Brendan Roof, Noah A. Smith, and Yejin Choi. 2019. ATOMIC: An Atlas of Machine Commonsense for If-Then Reasoning. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019*.
- [35] Robyn Speer, Joshua Chin, and Catherine Havasi. 2017. Conceptnet 5.5: An open multilingual graph of general knowledge. In *Thirty-first AAAI conference on artificial intelligence*.
- [36] Krishna Srinivasan, Karthik Raman, Jiecao Chen, Michael Bendersky, and Marc Najork. 2021. WIT: Wikipedia-Based Image Text Dataset for Multimodal Multilingual Machine Learning.
- [37] Chen Sun, Abhinav Shrivastava, Saurabh Singh, and Abhinav Gupta. 2017. Revisiting unreasonable effectiveness of data in deep learning era. In *Proceedings of the IEEE international conference on computer vision*. 843–852.
- [38] Bart Thomee, David A. Shamma, Gerald Friedland, Benjamin Elizalde, Karl Ni, Douglas Poland, Damian Borth, and Li-Jia Li. 2016. YFCC100M: The New Data in Multimedia Research. (2016).
- [39] Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: a free collaborative knowledgebase. *Commun. ACM* 57, 10 (2014), 78–85.
- [40] Amir Roshan Zamir, Alexander Sax, William B. Shen, Leonidas J. Guibas, Jitendra Malik, and Silvio Savarese. 2018. Taskonomy: Disentangling Task Transfer Learning. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*. Computer Vision Foundation / IEEE Computer Society, 3712–3722.
- [41] Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2018. From Recognition to Cognition: Visual Commonsense Reasoning. *CoRR* abs/1811.10830 (2018).
- [42] Xiangyun Zhao, Samuel Schuster, Gaurav Sharma, Yi-Hsuan Tsai, Manmohan Chandraker, and Ying Wu. 2020. Object Detection with a Unified Label Space from Multiple Datasets. In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XIV (Lecture Notes in Computer Science, Vol. 12359)*, Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm (Eds.). Springer, 178–193.