ECAI 2023 K. Gal et al. (Eds.) © 2023 The Authors. This article is published online with Open Access by IOS Press and distributed under the terms of the Creative Commons Attribution Non-Commercial License 4.0 (CC BY-NC 4.0). doi:10.3233/FAIA230353

A Semantics-Driven Methodology for High-Quality Image Annotation

Fausto Giunchiglia^a, Mayukh Bagchi^a and Xiaolei Diao^{a;*}

^aDepartment of Information Engineering and Computer Science (DISI), University of Trento, Italy. {fausto.giunchiglia, mayukh.bagchi, xiaolei.diao}@unitn.it

Abstract. Recent work in Machine Learning and Computer Vision has highlighted the presence of various types of systematic flaws inside ground truth object recognition benchmark datasets. Our basic tenet is that these flaws are rooted in the many-to-many mappings which exist between the visual information encoded in images and the intended semantics of the labels annotating them. The net consequence is that the current annotation process is largely underspecified, thus leaving too much freedom to the subjective judgment of annotators. In this paper, we propose vTelos, an integrated Natural Language Processing, Knowledge Representation, and Computer Vision methodology whose main goal is to make explicit the (otherwise implicit) intended annotation semantics, thus minimizing the number and role of subjective choices. A key element of vTelos is the exploitation of the WordNet lexico-semantic hierarchy as the main means for providing the meaning of natural language labels and, as a consequence, for driving the annotation of images based on the objects and the visual properties they depict. The methodology is validated on images populating a subset of the *ImageNet* hierarchy.

1 Introduction

Recent work in Machine Learning (ML) and Computer Vision (CV) has highlighted various types of systematic flaws in the development of object recognition benchmark datasets. See, for instance, [6, 37]. Our basic tenet is that these flaws are grounded in the way perception and language interact. As a matter of fact, an early version of this problem was already identified in [48] as the Semantic Gap Problem (SGP), which was described as the "lack of coincidence between the information that one can extract from the visual data and the interpretation that the same data have for a user in a given situation". This problem, still unsolved, was formalized in [14] as the fact that there is a many-to-many mapping between the visual information encoded in an image and the intended semantics of the corresponding linguistic descriptions. Thus, for instance, in Fig.1-(II), Image #251 is labeled Guitar but, just because it gives only a partial view of the instrument, it could have also been labeled Bass or Ukelele, while, at the same time, Guitar is also the label of images #154, #257. The net consequence is that the current annotation process is largely underspecified, thus leaving too much freedom to the annotator's judgment, who can then subjectively select one among the many SGP mappings.

In this paper, we propose vTelos¹, an integrated Natural Language Processing (NLP), Knowledge Representation (KR) and CV methodology whose main goal is to generate high-quality datasets. There are two main underlying intuitions. The first is that the *purpose* of an annotation effort, i.e., what the ML model trained by the dataset should be used for, should be made explicit via precise guide-lines, that we organize around *four main Design Choices*, as follows:

- C1: for each *image*, which object(s) should be considered?
- C2: for each *object*, which set of properties should be considered?
- C3: for each *set of properties*, which label should be selected to properly describe them?
- C4: for each *label*, if *polysemic*,² how do we unambiguously represent its meaning?

The second underlying intuition is that the role of the annotator should be split in two, i.e., that of the Classificationist and that of the Classifier, each associated with different tasks (this terminology being adopted from KR and Knowledge Organization (KO) [43]). Thus, the classificationist, will first define the space of labels to be used (Design Choices C3,C4) and, for each of them, the (visual) properties they describe. Then, in a second phase, the classifier will classify images using the labels defined by the classificationist (Design Choices C1,C2), where Choice C2 is based on the output of Choice C3. Notice how the proposed decision process is reversed with respect to how image annotation is usually performed, as also described above. The key idea is that we should first select the space of labels to be used, with a clear understanding of their intended meaning, in terms of the visual properties they are taken to describe, and then exploit this clarity when selecting, for each input image, the label which most properly describes its content. The proposed approach actually follows the standard practice in KR, for instance when classifying products in an online catalogue [18].

In this process, the intended natural language semantics of natural language terms (as selected during C3, C4) and the relevant visual information of images (i.e., objects and their visual properties, as selected during C1) are natively aligned (during C2), this being the reason why the SGP does not appear. The annotation bias is dealt with not by building an all-encompassing dataset, which would be impossible [4, 15] but, rather, by making explicit the annotation choices, as in [10]. This advantage is further amplified by the fact that, instead of performing choices C3, C4 from scratch, we take, as a predefined source of natural language labels, the WordNet Genus-Differentia lexico-semantic hierarchy [30]. In WordNet, each label is associated with a gloss, which, in fact, is an unambiguous linguistic definition of the properties of the objects it describes. This opens

^{*} Corresponding Author.

¹ From the Greek word Telos which means the end; the object aimed at in an effort; *purpose*, with v standing for visual, as in *visual purpose*.

² For instance, in WordNet, the average number of meanings per noun and per polysemous noun are 1.25 and 4, respectively, where the more common a noun, the higher the polysemy [12].



(I) Multi-Object Images

(II) Single-Object Images

(III) Mislabelled Images

Figure 1: Samples from ImageNet Musical Instrument Sub-Hierarchy (IL stands for ImageNet Label).

up the possibility of full and general alignment between natural language lexical semantics and object recognition, thus enabling a general purpose solution to the SGP many-to-many mapping problem.

There are three main contributions of this paper:

- 1. A classification of the main sources of annotation mistakes;
- 2. A general purpose annotation methodology, i.e., vTelos, for dealing with the SGP many-to-many mapping problem;
- 3. The application of vTelos to the alignment of the linguistic semantics of labels to the visual semantics of images. Notice how multi-lingual lexical resources, see, e.g., [35], allows to go beyond the current English-only main practice and apply vTelos to any natural language.

The paper is organized as follows. Section 2 introduces the sources of the annotation mistakes. Section 3 describes the vTelos four design choices. Section 4 provides the main algorithm which must be followed by annotators in order to implement Choice C2, when taking in input the WordNet hierarchy. Section 5 provides an evaluation of vTelos on (i) inter-annotator agreement, as a proxy of the role of subjective choices, (ii) annotation cost, and (iii) impact on quality of the ML models it generates. The evaluation is carried on a substantial subset of ImageNet [8]. The choice of ImageNet is not by chance. In fact, despite the fact that this information has hardly been used so far, ImageNet has been generated by populating WordNet, and, therefore, it natively contains the definitions encoding the meaning of labels. Furthermore, ImageNet provides a very good testbed for evaluation, having been analyzed in many studies, also focusing on annotation mistakes, see, e.g., [2, 46, 36, 42]. Related work (Section 6) and conclusion (Section 7) complete the paper.

2 Sources of annotation mistakes

Let us take an *annotation mistake* to be a label that does not correspond to the (classificationist) *intended semantics* of the image. The sources of annotation mistakes are the cause of the SGP many-to-many mapping as they raise the possibility of choices that are different from those of the classificationist. Annotation mistakes have three main recurring sources, as follows (see, e.g., [52]³ for an extensive evaluation-based analysis of the problem):

- 1. *Mislabelled Images* (MI), where the flaw arises because of mislabelling.
- 2. *Multi-Object Images* (MOI), where the flaw arises from the occurrence of multiple objects in the same image and the consequent possible non-correspondence between the classificationist and the classifier assigned labels.
- Single-Object Images (SOI), where the flaw arises from an ambiguity in how to label the single object in the image.

Mislabelled images are an extreme case of SGP many-to-many mapping, and are motivated by factors such as carelessness, missing knowledge or excessive speed, that is, all well-known crowdsourcing problems, see for instance [5]. A concrete example is in Fig.1-(III), wherein Image #383, labelled as *Acoustic Guitar*, actually depicts a birthday cake.

Multi-Object Images (Fig.1-(I)) constitute more than one-fifth of ImageNet's total images. Here the source of subjectivity is the occurrence of multiple objects, coupled with the fact that the objects selected often do not correspond to the most likely main object. One such example is the image of a *Stage* (Image #193), having ImageNet label *Electric Guitar*. Thus, on one side, empirical evidence from cognitive psychology [44] makes clear that the main object chosen by humans possesses the highest '*cue validity*', i.e., the one which is visually the most salient, while, on the other side, many ImageNet labels correspond to *random features* which don't generalize to object recognition in the wild (see [52] for more details).

Single-Object Images (Fig.1-(II)), are by far the most important source of annotation mistakes. The main reason is in the underlying interaction between *visual polysemy* [48] and *linguistic polysemy* [30]. Single-Object images present three types of ambiguity.

In the first case (SOII), the source of subjectivity is intrinsic to the Image itself. It occurs when an object in an image is visually polysemic⁴, namely when its visual "semantics is described only partially" [48] and, consequently, its linguistic description is not unique. An example is Image #251 representing a dreadnought-shaped instrument, which is labelled *Guitar* in Fig.1-(II), but could also be labelled, e.g., *Bass* or *Ukelele*. Notice that images of this type usually encode a typical partial view with no possibility of discrimination among different objects. [52] observes that not only do humans assign an alternative label for 40% of the images, but they also assign up to 10 different labels.

The second source of subjectivity in single-object images (SOIL) is intrinsic to the meaning of Labels. There are two cases. In the case of polysemic labels (i.e., of one-to-many mapping between labels and images), different objects, each corresponding to one of the possible meanings, are associated with the same label[16]. One such example is the two images labelled as *Dulcimer* (Image #262, #266) in Fig.1-(II)) where the two objects are essentially two different musical instruments, i.e., the first *Dulcimer* being native to the American and the second to the Chinese culture, and are visually very different. The study in [52] highlights how these ambiguous classes generate overlaps in the image distributions with a negative impact on performance. The second case labels which are synonyms (i.e., of many-to-one mapping between labels and images) wherein a single set of

³ https://arxiv.org/abs/2005.11295 is an extended version.

Role	Annotation Choice	Purpose	Source of Subjectivity	Solution
Classificationist	C3: Label generation	Choose label and its properties	MI, SOII	Linguistic Genus-Differentia
Classificationist	C4: Label disambiguation	Choose identifier	SOIL	Alinguistic Identifiers
Annotator	C1: Object localization	Choose object	MOI	Bounding Polygons
Annotator	C2: Visual classification	Choose label via its properties	MI, SOII, SOIL, SOIA	Visual Genus-Differentia

 Table 1: The vTelos four choices. MOI and MI stand for Multi-Object and Mislabelled Images, respectively. SOII, SOIA stand for Single-Object Images where the source of subjectivity is the Image, the Label and the Annotator background knowledge, respectively.

images, will wrongly be split into two sets. One such example is that of *Koto* which can also be synonymously labelled as, for instance, a *Kin* or a *Jusangen*[13].

The third source of subjectivity in single-object images (SOIA) relates to the Annotator incomplete background knowledge. Thus, for instance, Image #257 is labelled as *Guitar* in Fig.1-(II), while the image depicts an Acoustic Guitar. This type of generic label appears whenever correct labelling requires deep domain specific knowledge. [6] discusses the bias which arises because of this problem in the case of minority languages, cultures, and also with images about low income people. Notice that this is not mislabelling. The issue is that the label captures some but not all of the visual properties of the object.

Overall, the sources of annotation mistakes can be summarized by saying that MOI images generate a one-to-many mapping from images to labels, those of type SOIA generate a many-to-one mapping, while those of type MI, SOII, and SOIL generate both types of mappings. These four types of mistakes are independent of one another and can, therefore, occur together in the same image, thus generating multiple occurrences of the SGP many-to-many mapping problem.

3 Annotation design choices

As from the introduction, the vTelos annotation process is articulated into four main choices C1-C4. Let us analyze them in detail following the order in which they are made.

- 1. Choice C3: Label generation. Choose the space of labels used to annotate images. Here the meaning of each label is *defined* in terms of linguistically defined properties encoding a selected set of visual properties. One such example is the definition of a *guitar* as being a "*string instrument with six strings*". This eliminates any source of *linguistic ambiguity* intrinsic to *informally defined labels* and it allows to say that two labels are *synonyms*.
- Choice C4: Label disambiguation. Choose a unique concept identifier for each of the (one or more) meanings of each label. This eliminates any source of *linguistic ambiguity* intrinsic to *polysemous labels* (see footnote ²).
- Choice C1: Object localization. Choose one of the (one or more) objects in an image, thus eliminating any possible source of object ambiguity.
- 4. Choice C2: Visual classification. Choose the relevant visual properties of the object under consideration and annotate the image with one of the labels selected in Choice C3. This choice is conceptually similar to the usual annotation choice. The key difference is that what is selected is not the label but, rather, the linguistic properties that describe the visual properties of the object. Thus for instance, if the object occurring in the image is a guitar, the annotator would not choose the label Guitar but, rather the property having six strings. The label Guitar would then be automatically associated to the label because of the definition of guitar generated by Choice C3. This eliminates any source of visual ambiguity.

These four choices allow us to deal with the annotation mistakes described in Section 2. Consider Table 1, where the first column reports the two vTelos roles (with *Annotator* can also be read as *Classi-fier*), the second, the design choice, the third, the purpose behind the choice, the fourth, the source of subjectivity which is addressed and the last, the solution that is enforced.

In Table 1, the first and the last column are crucial to the vTelos approach for three reasons. The first is the splitting of the annotation tasks into two roles, i.e., classificationist and annotator, where the first is in charge of preparing a well-specified annotation task for the second. The second reason is that, opposite to the annotation standard practice, in vTelos, Choices C3,C4 are performed before Choices C1,C2. This is in order to force the annotator to select a label within of a (possibly very large) pre-selected controlled vocabulary. Following vTelos this vocabulary should be built by a domain expert. The third reason is that classificationist and classifier coordinate through the use of the Visual and the Linguistic Genus-Differentia (see the last column in Table 1). The idea behind the notion of Linguistic Genus-Differentia is that the meaning of a label (e.g., Guitar) is defined in terms of a linguistic description composed of two parts, the (linguistic) Genus, which describes what the object has in common with other objects (in the case of a guitar the fact that it is a string instrument), and the (linguistic) Differentia which describes which properties differentiate the object from the other objects with the same genus (in the case of a guitar, the fact that it has six strings). In turn, given a label whose meaning is defined in terms of Linguistic Genus-Differentia, any object denoted by that label, e.g., a guitar, must have a (visual) Genus which is described by the linguistic genus of the label and similarly for the (visual) differentia. We often talk of linguistic and visual differentia in terms of linguistic and visual properties.⁵

For what concerns the classifier, the goal of Choice C1 is to select the image main object (see Section 2). Consider for instance, in Fig.1 (II), the objects in the MOI Image #171, i.e., a Koto but also a Flute, Music Stand etc., all labelled as Koto. Choice C1 solves the problem of the MOI source of subjectivity. The possible strategies are: (i) removing the image just because of multi-object, (ii) splitting the image into sub-images, one per object, or (iii) isolating objects, e.g., via bounding polygons, in this case allowing for multi-label images. Notice how Choices C1, C2 are crucially based on the distinction between object localization and visual (image) classification, where the former is an inherently visual activity where an object is localized but not recognized, while the latter is an inherently semanticsdriven activity which determines the relevant properties of an object. This distinction, while being well known, is often not enforced in many state-of-the-art approaches, which collapse object detection and recognition (see [20, 45]).

Choice 2 enables the classifier to select the label whose genus and differentia align with the visual genus and differentia of the object in the image. There are only two possibilities. Either such a label does

⁵ The distinction between the visual properties of an object and the linguistic properties which are used, in natural language, to define the label used to name the object, is based on the *Teleosemantics* theory of meaning [31, 32], as adapted to CV in [17].



Figure 2: The ImageNet Musical Instrument hierarchy. Each node is associated with a unique linguistic identifier (id), an ImageNet label (L), a Genus (G), and a Differentia (D). Nodes are populated with the images from Fig.1 where red and green ids correspond to the ImageNet and vTelos classification. The images in the yellow box are discharged by vTelos.

not exist, in which case the object does not belong to that class, or it does exist, in which case the opposite holds. The case of multiple possible labels cannot occur as two different objects with the same genus cannot have the same differentia. This is the key guideline enforced by the classificationist to avoid the occurrence of the SGP many-to-many problem, modulo mistakes on the side of the annotator, i.e., the selection of a label that does not describe the contents of the image. Thus, for instance, a positive example is Image #238 in Fig.1-(II), where the ImageNet gloss of the label Keyboard Instrument forces the selection of the visual property of being a "a musical instrument that is played by means of a keyboard". Dually, Choice 2 forces the rejection of all images which do not satisfy the meaning of the label. Thus, for instance, vTelos will reject the MI Image #383, as the ImageNet label Acoustic Guitar is defined as "a guitar with no input jack"; but the picture does not represent a guitar and, more in general, not even a musical instrument.

4 **Annotation process**

If one looks at the examples provided in Section 3, s/he will find out that each single linguistic differentia is quite fine-grained, which is motivated by the need of making it precise and not ambiguous. S/he will also notice that, on the other hand, the linguistic genus is quite coarse-grained and ambiguous, being in fact an ImageNet label. This apparent contradiction is solved by observing that the meaning of any genus can be recursively generated in terms of a second more general genus plus multiple differentiae down to the genus itself. Thus, for instance, the meaning of the genus String Instrument of the label Koto can be reconstructed as being a Device, with sound mechanism, with taut strings. This leads to the consequence that Choices C3, C4, as carried out by the classificationist, generate a genus-differentia classification hierarchy which suitably organizes the meaning of labels.

The mainstream approach to the generation of this type of hierarchy is the *faceted classification* methodology [43], a tried-and-tested paradigm of developing classification schemas devised specifically to classify any type of information resource. A key innovation of this paper is the suggestion of using, as genus-differentia classification hierarchy, the WordNet lexico-semantic hierarchy [30] and its evolutions, as developed in the work in Computational Linguistics. Thus,

Algorithm 1 $L = V$ is Classify(1)								
H = Tree of (N);								
$N = \langle id, D \rangle;$								
Input: Image <i>I</i> ;								
Output: Image label L;								
1: $R = getRoot(H);$								
2: if $askDif(I, R.D) = False$ then								
3: return $L = "Discharged";$								
4: end if								
5: $L = R.D;$								
6: while $hasChild(R) = True$ do								
7: while $C \in getChild(R)$ do								
8: if $askDif(I, C.D) = True$ then								
9: $R = C; break;$								
10: end if								
11: end while								
12: $L = R.D;$								
13: end while								
14: return <i>L</i> ;								

(T)

for instance, in WordNet we have the following definition of *piano*:

<i>Word</i> : p	iano
Sense ID: 03934354; piano#1	Sense ID: 04998511; piano#2
Synset = piano, forte-piano,	Synset = piano, pianissimo,
Gloss = "a keyboard instrument	Gloss = "(music) low loudness'
that is played by depressing keys"	

From this definition, we know that the label piano has two unique meanings, "piano#1" and "piano#2", called senses, each associated with a unique sense identifier, e.g., 03934354. Different synonymous words for the same sense are clustered into synsets. Further, each synset is associated with a gloss which is a natural language definition of the concept expressed by the synset in terms of genusdifferentia. For instance, the synset {piano, forte-piano, ...}, associated to the sense "piano#1", is defined by a gloss which encodes its genus "keyboard instrument" and its differentia "depressing keys". Fig.2 reports an abstraction of the WordNet Musical Instrument hierarchy. In turn, in ImageNet, this hierarchy has been populated with 3660 images. Notice how in Fig.2 it is possible to reconstruct the definition of the label Koto provided above.

The use of lexico-semantic hierarchies as genus-differentia classifications has three main consequences. The first is that lexicosemantic hierarchies can be reused, at no cost, for the generation of high-quality datasets (thus generalizing the idea behind the development of ImageNet). The second is that the follow-up application of Choices C1, C2 allows to generate datasets where the semantics of language and vision are natively aligned thus enabling a general solution to the SGP many-to-many mapping problem. The third is that, given a set of images that populate the root node, the only information needed to label an image is its visual differentia. As the following will make clear, this last fact is the one which causes a somewhat radical departure from the current mainstream approach to image annotation.

Choice C2, as carried out by the classifier, takes in input the hierarchy generated by Choices C3, C4 and associates each single image to a node in the hierarchy. The resulting algorithm is reported in Algorithm 1. This algorithm takes in input a hierarchy H, e.g., the WordNet/ ImageNet lexico-semantic hierarchy in Fig.2, and an Image I and produces in output the image label L (Lines 3, 12). After a first initialization where it checks whether the input Image can be classified, i.e., it is in the scope of the root label R (as defined

by differentia from a more abstract genus) (Lines 1-4), Algorithm 1 proceeds down the tree. This is achieved via a first *while* (Line 6) which recursively descends the hierarchy and a second *while* (Line 7) which, at each level in the hierarchy, checks whether there is a sibling matching the image visual differentia (Line 8). Algorithm 1 applies the *GetSpecific* principle [18], i.e., it stops whenever none of the siblings satisfies the test in Line 8.

Figure 2 shows the results of applying Algorithm 1 to the images in Fig 1. Here, the images which cannot be classified in any of the nodes below the root node are collected in the yellow box and labeled Discharged. Let us see how the various types of annotation mistakes from Section 2 are handled. We assume that MOI images have been previously split in multiple sub-images, one per object, with numbering increasing from left to right. In fact, as from Table 1, MOI images are dealt with by Choice C1. We also assume that SOIL labels have been handled when building the classification hierarchy. Thus, for instance in Fig. 2, the label Dulcimer, coherently with the ImageNet differentia, has a single meaning describing American Dulcimers. Taking into account Chinese dulcimers would require extending ImageNet by adding a new sense to the label Dulcimer. Synonyms are natively handled via synsets. MI images are moved to the label which suitably described their visual contents. Thus, for instance, Image #364 labelled Dulcimer in Fig.1 (III), is discharged. SOII images are always classified higher in the hierarchy, if compared to the ImageNet position, this is because the partial view of the object does not allow for its full discrimination. Thus, for instance, Image, #251 labelled Guitar in Fig.1 (II), is labelled as a musical instrument. With SOIA images the situation is the opposite of SOII images, that is, they are classified deeper in the hierarchy. Thus, for instance, Image #257 labelled as Guitar in Fig.1 (II) can be correctly labelled as an Acoustic guitar based on the visual property "no input jack".

Based on the above, it is possible to classify the results of the annotation of an image as follows.

- Correct. The image annotation is correct;
- *Incorrect*. The image annotation is wrong. This case can be further refined as follows:
 - *Generic*. The image annotation is more general than the correct label, i.e., it does not exploit all its visual properties;
 - Restricted. The image annotation is more specific than the correct label, i.e., it has been generated by refining the correct label using visual properties which are not visible;
 - Misplaced. The image annotation is unrelated to the correct label, i.e. it has been generated by using visual properties which are alternative to those needed to select the correct label.

MI images can take any type of incorrect label. SOIA images always have generic labels. SOII images always have restricted labels.

The last observation is that lexico-semantic-resources are a very good starting point but more work needs to be done to make them fully usable for image annotation. This is in fact ongoing work. Fig.2 exemplifies two cases. The first is the root node whose differentia is based on sound and, as such, not recognizable visually. This applies to all differentia describing senses. The second is the differentia of the label *Acoustic Guitar*. This differentia, articulated as a missing visual feature, will always empty the parent node, in this case, the label *guitar*.

5 Evaluation

As from Section 1, all the experiments focus on ImageNet. We focus only on the annotator choices C1, C2, where choices C3, C4 are automatically enforced by the fact that we use ImageNet as background knowledge. To this extent, in Choice C2, we assume the visual differentia to be as defined by the Imagenet gloss. We perform four evaluations. In the first, we evaluate the inter-annotator agreement, as a proxy of the decrease of the role of subjective choices. We consider both expert and non-expert annotators. In the second we evaluate the annotation cost. In the third, we evaluate the increase of accuracy generated on various state-of-the-art ML model. The last experiment is an ablation study.

5.1 Inter-annotator agreement

We selected 450 of the 3660 images populating the ImageNet categories in Fig.2, 50 images per category. To minimize data bias, these 450 labelled image dataset, let us call GT1 (where GT stands for *Ground Truth*) was randomly generated. Then, a second dataset GT2 was generated by an expert annotator EA₁ by relabelling the 450 images of GT1 using Choice C2 only, as described in Section 4. We then selected 16 volunteer *non-expert annotators*, with no previous annotation experience. The annotators were university students and collaborators, with an average age of 29 years. To maximize diversity, we selected people from different disciplines and from 3 continents and 7 countries. The annotators were organised in 2 groups of 8 people, each annotator relabelling GT1 and generating a dataset as follows:

- GT3, as created by each member of Group 1 The annotators were asked to perform the annotation in two steps. First, they had to annotate images based on Choice C2. Here they were allowed to create a new category "*Discharged*" for all the images they could not classify. Then, in the second step, they were asked to label the categories generated in the first step S2 with their "most suitable" label, not necessarily an ImageNet label (which was not provided).
- GT4, as created by each member of Group 2. This group was instructed to annotate images using the ImageNet labels plus the label "Discharged".

The motivation for focusing EA1 and Group 1 on Choice C2 was that of highlighting the use of the differentia in place of the label. The results are reported in Table 2. Notice how this table reports the number of images associated with each label for all four datasets and for each annotator working on GT3 and GT4. As it can be seen, the number of images (and the images) associated with the same label are different. We measure the inter-annotator agreement of GT3 and GT4 using Krippendorff's alpha [25, 1] (last row). Krippendorff's alpha is 0.5973 for GT3 and 0.5047 for GT4, that is, GT3 is around 18% higher than GT4. This is an important improvement but still not good enough.⁶ The cause of this, not fully satisfactory, result was the Multi-object images which are more than 58% of the total (2125 out of 3660 images). We have therefore performed a second experiment where Choice C1, C2 were enforced and where MOI images were proposed multiple times with a bounding box around the object to be annotated. This experiment was performed by Group 1 (generating eight datasets GT2*) and two expert annotators EA1, EA2 (generating two datasets GT3*). Table 3 reports the results obtained. Here,

⁶ As from [24][Table 2, p. 417], a value of alpha in the range (0.6, 0.8] indicates substantial agreement while an index in the range (0.8, 1] indicates near-perfect agreement.

			GT3 (Annotation via Differentia labels)									GT4 (Annotation via Category labels)								
Id.	GT1 GT2	GT2	Differentia	NA1.1	$NA_{1.2}$	NA1.3	$NA_{1.4}$	NA1.5	NA1.6	NA _{1.7}	NA1.8	Categories	NA _{2.1}	NA _{2.2}	NA2.3	NA _{2.4}	NA2.5	NA _{2.6}	$NA_{2.7}$	NA2.8
1	50	41	with Sound Mechanism	33	12	27	25	28	29	12	18	Musical Instrument	16	42	100	17	27	20	19	26
1_1	50	123	with Taut Strings	46	97	71	133	112	83	62	79	Stringed Instrument	162	78	0	115	144	106	161	74
1_1_1	50	34	with 6 Strings	37	13	40	34	58	34	19	31	Guitar	77	18	41	40	13	9	0	8
1_1_1_1	50	40	with No Input Jack	66	77	53	54	26	67	68	50	Acoustic Guitar	13	81	66	43	70	70	82	71
1_1_1_2	50	43	with Input Jack	54	61	67	43	28	45	82	78	Electric Guitar	74	65	86	44	70	71	68	74
1_1_2	50	31	Elliptical body with 3 or 4 strings	61	37	36	32	29	30	37	30	Dulcimer	0	6	31	31	0	16	6	46
1_1_3	50	27	with 13 Strings	42	32	42	16	18	44	53	45	Koto	0	47	47	39	0	41	0	43
1_2	50	47	with Keyboard	49	46	41	49	43	47	50	46	Keyboard Instrument	41	49	0	44	40	47	52	44
1_3	50	47	with Embouchure	62	60	60	63	50	55	60	59	Wind Instrument	61	57	55	62	54	61	60	59
	0	17	Unrecognised	0	15	13	1	58	12	7	14	Unrecognised	6	7	24	15	32	9	2	5
all	450	450	Krippendorff's alpha	0.5973						Krippendorff's alpha	a 0.5047					-				

Table 2: Inter-annotator agreement evaluation with Choice C2 only. The images in GT1 have ImageNet labels. GT2, GT3, GT4 images have labels selected by the expert annotator, and by the 8+8 non-expert annotators in Groups 1 and 2, respectively. $NA_{j,i}$ stands for annotator *i* of Group *j*. Id. is the label identifier, as from Fig. 2. The last row reports the value of Krippendorff's alpha.

Index	Catagoria	GT2*		GT3* (Only Single-Object images)								
muex	Categories	EA ₁	EA_2	NA1.1	$NA_{1.2}$	NA1.3	$NA_{1.4}$	$NA_{1.5}$	$NA_{1.6}$	NA1.7	NA1.8	
1	with Sound Mechanism	17	17	16	2	5	11	11	11	4	4	
1_1	with Taut Strings	42	42	21	41	32	43	41	39	28	32	
1_1_1	with 6 Strings	21	20	21	8	23	19	26	24	11	15	
1_1_1_1	with No Input Jack	21	22	31	34	29	28	18	32	33	29	
1_1_1_2	with Input Jack	22	22	24	32	31	21	20	21	39	39	
1_1_2	Elliptical body with 3 or 4 strings	13	13	24	13	13	14	12	13	15	14	
1_1_3	with 13 Strings	12	12	12	7	10	8	9	10	13	9	
1_2	with Keyboard	33	33	33	33	29	34	29	31	34	33	
1_3	with Embouchure	10	10	20	20	21	23	14	15	19	21	
	Unrecognised	11	11	0	12	9	1	22	6	6	6	
all	Krippendorff's alpha	0.9877		0.7595								

 Table 3: Inter-annotator agreement with Choices C1, C2 (full vTelos), generating two GT2* and eight GT3* datasets.

Anno.	Acoustic Guitar	Dulcimer	Koto
NA _{1.1}	Guitar	Appalachian Dulcimer	Biwa
NA _{1.2}	Guitar	IDK (I Don't Know)	IDK
NA1.3	Wooden guitar	IDK	Zither
NA _{1.4}	Wooden guitar	3 or 4 String Musical Instrument	13 String Koto
NA _{1.5}	Hawaiian Guitar	3-4 Stringed Elliptical Instrument	13-String Instrument
NA1.6	6 Stringed Instrument no Jack	Fretted Stringed Instrument	13 Stringed Instrument
NA _{1.7}	Classic Guitar	Elliptical Stringed Instrument	Rectangular Stringed Instrument
NA _{1.8}	Non Powered Guitars	Short-Stringed Music Instruments	Japanese Stringed Instrument

Table 4: Category Labels suggested by Group 1 annotators.

Krippendorff's alpha goes up to 0.7595 among non-expert annotators and 0.9877 between the two experts, namely up to near-perfect agreement.

An interesting final observation is that in Group 2, but not in Group 1, some annotators were unable to annotate some images (see, e.g., the "0"s in Table 2). Two such examples are "Dulcimer" and "Koto".⁷ This observation is also confirmed in Table 4, which reports some example labels provided during the second part of the Group 1 experiment. The observation here is that, despite properly labeling images via differentia, some annotators did not know the names of some categories, or used wrong or more generic labels, or in one case used a more specific label.

5.2 Annotation cost

As from Section 4, vTelos requires that each image is recursively compared with the node one level down in the lexico-semantic hierarchy till it meets a differentia which it does not satisfy. So the question arises of what is the cost of the multiple checks. The preliminary observation is that, while containing an excess of 100K concepts, WordNet is a broad and flat hierarchy whose paths have a depth in the range of around 7 to 15 nodes (with different depths for different versions) and where the nodes higher in the hierarchy are very abstract and have no visual content. A 4-node deep hierarchy, like the one in Fig. 2, is therefore a good proxy for evaluating the annotation cost. We have therefore compared the time consumed by two non-expert annotators in the generation of two GT3 and GT4-like datasets restricted to contain one hundred images. The results show that the average annotation time per image is 5.565 seconds for GT3 and 3.473 seconds for GT4, that is, a time increase for vTelos of around 60% which one would guess it would become a little more than double with 8-node level deep hierarchies. The time increase correlates to the images classified in intermediate nodes. The fact that it is a low increase seems to correlate to the fact that each single choice is simpler given that: (i) the same image is seen more than once, (ii) the choice is restricted within a precise set of labels, (iii) which have similar meanings, because selected following the pattern induced by the tree structure, thus ultimately improving the automation of the process.

5.3 Machine Learning Accuracy

We applied eight state-of-the-art ML methods trained on two datasets, called below "ImageNet" "vTelos, the first containing the original 3660 images, the second containing the same 4605 images where MOI images are transformed in SOI images. 80% of the images were used as training set and the remaining 20% were the test set.

All the experiments have been implemented with PyTorch with identical settings. During the training, we performed the same data augmentation (horizontal flipping and random scaling with multiple times) on each dataset and randomly sampled 224×224 crops from augmented images. All experiments have been optimized using Adaptive Moment Estimation with learning rate initialized to 0.0002, momentum initialized to 0.9, and weight decay to 10^{-8} . All models have been pre-trained on ImageNet with ResNet101 [21]. The results are reported in Table 6. It can be noticed that all models trained on vTelos get significant improvement in accuracy, especially ResNet which achieves up to 23.44% improvement.

Finally, if we look at the heat maps in Fig. 3, we can see how the models pay attention to the visual regions described by the linguistic differentia, see for instance the keyboard in the case of the label *Keyboard Instrument*.

5.4 Ablation study

We conducted two ablation studies, the first focusing on Choice C1, the second of Choice C2 for all the eight systems tested in the previous experiment. The results are reported in Table 5 where the first and the last row replicate the results in Table 6. The key observation is that, in all cases, the improvement in performance due to Choice C2 with respect to that due to Choice C1 is substantial (e.g., 5 times with AlexNet, 2.9 times with ZFNet, 13.5 times with RAN, 8 times with ResNet) thus providing further confirmation of the intuition suggested by the heat maps in Fig. 3. Taking into account the relevance of Choice C1 which transforms Multi-Object Images into Single-Object images, we take this as an important indicator of the fact that *labeling images using differentia labels is the way to go*.

⁷ Koto is a Japanese instrument.

Methods	AlexNet		ZFNet		VGG		GoogleNet		ResNet		DenseNet		RAN		SENets	
Methous	Acc.	Impro.(%)	Acc.	Impro.(%)	Acc.	Impro.(%)	Acc.	Impro(%)	Acc.	Impro.(%)	Acc.	Impro.(%)	Acc.	Impro.(%)	Acc.	Impro.(%)
ImageNet	0.543	-	0.612	-	0.655	-	0.734	-	0.593	-	0.724	-	0.713	-	0.728	-
only Chioce C1	0.551	1.47%	0.623	1.80%	0.672	2.55%	0.745	1.55%	0.610	2.86%	0.733	1.24%	0.718	0.74%	0.746	2.54%
only Chioce C2	0.583	7.37%	0.644	5.23%	0.732	11.82%	0.829	12.94%	0.729	22.93%	0.782	8.01%	0.770	8.04%	0.807	10.87%
Choice C1 & C2 (vTelos)	0.596	9.76%	0.657	7.35%	0.743	13.44%	0.835	13.76%	0.732	23.44%	0.793	9.53%	0.784	9.96%	0.811	11.40%

Table 5: Ablation study. "Acc." is the accuracy, "Impro.(%)" is the improvement in performance over ImageNet.

Mathada	Accura	acy	Improvement
Wiethous	ImageNet	vTelos	mprovement
AlexNet [26]	0.543	0.596	9.76%
ZFNet [56]	0.612	0.657	7.35%
VGG [47]	0.655	0.743	13.44%
GoogleNet [49]	0.734	0.835	13.76%
ResNet [21]	0.593	0.732	23.44%
DenseNet [23]	0.724	0.793	9.53%
RAN [53]	0.713	0.784	9.96%
SENets [22]	0.728	0.811	11.40%

Table 6: Classification results for ImageNet and vTelos.



Figure 3: Attention heat maps.

6 Related work

In this work, we build on top of the ImageNet work [8]. With respect to this work, the main innovations are as follows. First, the idea of using (differentia) properties, and not class names, for labeling images. Second, the idea of defining an annotation methodology, rather than just producing a data set. Third, exploiting the faceted classification approach as a powerful technique for further improving the annotation quality. Finally, the fact that our approach allows for the generation of natural language aware annotations. As a consequence, the relabeling of the ImageNet images will open the door to the possibility of a dataset labeled in up to 1000+ languages, including a large number of minority languages.

As from the introduction, the problem of dataset quality has been known for a while, with [51] being an early paper focusing on it and [48] being the work which identified the SGP as the main cause of the problem. Recent work has started focusing on ways for dealing with the problem of dataset quality. Thus, [40] advocates more careful practices in the development of datasets that are attentive to their limitations and impact. [54] advocates for fairer datasets and provides a vertical methodology for balancing the people subtree of ImageNet. [55] shows how to improve the quality of ImageNet by using a classifier trained in a higher-quality dataset; this work is complementary to ours as it leaves open the problem of how to generate the higher quality dataset. As far as we know, vTelos is the first methodology for generating high-quality datasets and it does so by using a highly

inter-disciplinary KR/CV/NLP approach.

With respect to the annotation process, the crowdsourcing community has focused extensively on the problem of quality, see, e.g., [38, 11]. [5] provides a quite comprehensive characterization of quality in crowdsourcing and an extensive analysis of the state of the art. Very recently some work has started to focus on how to operationally improve the quality of the process, see, e.g., [27, 7]. Most relevant to this paper, [34] provides and exploits a set of metrics, including Krippendorff's alpha, with the goal of monitoring the image annotation process. However, as far as we know, there two are key differences. The first is that vTelos is an end-to-end general methodology. The second is that, differently from the work above, which focuses on how to measure the effects and control the behaviour of annotators, here the focus is on aligning the semantics encoded in images and in NLP descriptions, i.e., in dealing with the SGP many-to-many mappings.

Earlier work on the SGP has focused on how to integrate featurelevel and semantic-level information. Thus, some have proposed the use of ontologies [19], others of high-level features [29, 9], and others to ask users [50]. More recently [39] has proposed to handle the SGP when aggregating multi-level features. All these approaches focus on object labels and not on the differentia.

From a methodological point of view, this work is based on a mix of KR and NLP. In KR, as applied to CV, a fair amount of work, motivated by (Cognitive) Robotic applications has concentrated on identifying the function of objects see, e.g., [3, 41, 28]. None of this work uses the faceted classification approach. Furthermore, a second difference is that this work has concentrated on how to enable a meaningful human-machine interaction and not, as it is the case here, on how to enforce coherence between how humans and machines describe and name objects. A lot of work has also focused on the integration of NLP and CV, see [33] for a recent very extensive survey of the field. While none of this work is focused on the definition of a general methodology for high quality dataset generation, image captioning could be used towards the automatic generation of glosses. This would be relevant when there is a need for new corpora for which human annotation is not an option (for instance because requiring expert annotators).

7 Conclusion

In this paper, we have proposed a general CV/NLP/KR methodology for producing high quality image datasets. The motivation for this work lies in the need of overcoming some of the limitations which exist in the current annotation process. This is only a first step. We believe in fact that the use of NLP/KR techniques, as the main mechanism for labeling examples, is the key to the development of explainable and high performance ML systems.

Acknowledgements

This research has received funding from the European Union's Horizon 2020 FET Proactive project "WeNet – The Internet of us", grant agreement No 823783.

- [1] Ron Artstein, 'Inter-annotator agreement', in *Handbook of linguistic* annotation, 297–313, Springer, (2017).
- [2] Lucas Beyer, Olivier J Hénaff, Alexander Kolesnikov, Xiaohua Zhai, and Aäron van den Oord, 'Are we done with imagenet?', arXiv preprint arXiv:2006.07159, (2020).
- [3] Luca Bogoni and Ruzena Bajcsy, 'Interactive recognition and representation of functionality', *Computer Vision and Image Understanding*, 62(2), 194–214, (1995).
- [4] Paolo Bouquet and Fausto Giunchiglia, 'Reasoning about theory adequacy. a new solution to the qualification problem', *Fundamenta Informaticae*, 23(2, 3, 4), 247–262, (1995).
- [5] F. Daniel, P. Kucherbaev, C. Cappiello, B. Benatallah, and M. Allahbakhsh, 'Quality control in crowdsourcing: A survey of quality attributes, assessment techniques, and assurance actions', ACM Computing Surveys, 51(1), 1–40, (2018).
- [6] Terrance de Vries, Ishan Misra, Changhan Wang, and Laurens van der Maaten, 'Does object recognition work for everyone?', in *CVPR Work-shops*, pp. 52–59, (2019).
- [7] Gianluca Demartini, Kevin Roitero, and Stefano Mizzaro, 'Managing bias in human-annotated data: Moving beyond bias removal', *arXiv* preprint arXiv:2110.13504, (2021).
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei, 'Imagenet: A large-scale hierarchical image database', in *IEEE Conf. computer vision and pattern recognition*, pp. 248–255, (2009).
- [9] M. Elahi, Y. Deldjoo, Farshad Bakhshandegan M., L. Cella, S. Cereda, and P. Cremonesi, 'Exploring the semantic gap for movie recommendations', in *RecSys*, pp. 326–330, (2017).
- [10] L. Erculiani, A. Bontempelli, A. Passerini, and F. Giunchiglia, 'Egocentric hierarchical visual semantics', in *HHAI Conf.*, (2023).
- [11] Ralph Ewerth, Matthias Springstein, Lo An Phan-Vogtmann, and Juliane Schütze, "are machines better than humans in image tagging?"-a user study adds to the puzzle', in *ECIR*, pp. 186–198. Springer, (2017).
- [12] Abed Alhakim Freihat, 'An organizational approach to the polysemy problem in wordnet', (2014).
- [13] F. Giunchiglia, M. Bagchi, and X. Diao, 'Aligning visual and lexical semantics', in *Conf. on Information (iConference)*, (2023).
- [14] F. Giunchiglia, L. Erculiani, and A. Passerini, 'Towards visual semantics', Springer Nature Computer Science (SNCS), 2(6), (2021).
- [15] F. Giunchiglia, S. Kleanthous, J. Otterbacher, and T. Draws, 'Transparency paths-documenting the diversity of user perceptions', in Adjunct Proc. of the 29th ACM Conference UMAP, pp. 415–420, (2021).
- [16] Fausto Giunchiglia, Xiaolei Diao, and Mayukh Bagchi, 'Incremental image labeling via iterative refinement', in *IWCIM@ICASSP*, (2023).
- [17] Fausto Giunchiglia and Mattia Fumagalli, 'Concepts as (recognition) abilities', in FOIS, pp. 153–166, (2016).
- [18] Fausto Giunchiglia, Ilya Zaihrayeu, and Uladzimir Kharkevich, 'Formalizing the get-specific document classification algorithm', in *TPDL*, pp. 26–37. Springer, (2007).
- [19] JS Hare, Paul H Lewis, PGB Enser, and CJ Sandom, 'Mind the gap: another look at the problem of the semantic gap in image retrieval', in *Multimedia Content Analysis, Management, and Retrieval*, (2006).
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, 'Spatial pyramid pooling in deep convolutional networks for visual recognition', *IEEE TPAMI*, **37**(9), 1904–1916, (2015).
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, 'Deep residual learning for image recognition', in CVPR, pp. 770–778, (2016).
- [22] Jie Hu, Li Shen, and Gang Sun, 'Squeeze-and-excitation networks', in CVPR, pp. 7132–7141, (2018).
- [23] G. Huang, Z.ang Liu, L. VanDer Maaten, and KQ. Weinberger, 'Densely connected convolutional networks', in CVPR, (2017).
- [24] John Hughes, 'Krippendorffsalpha: An r package for measuring agreement using krippendorff's alpha coefficient', *The R Journal*, 1(1), (2021). Also: arXiv preprint arXiv:2103.12170.
- [25] Klaus Krippendorff, 'Computing krippendorff's alpha-reliability', Computing, 1, 1–25, (2011).
- [26] A. Krizhevsky, I. Sutskever, and GE Hinton, 'Imagenet classification with deep convolutional neural networks', *Neurips*, (2012).
- [27] K. Kyriakou, P. Barlas, S. Kleanthous, E. Christoforou, and J. Otterbacher, 'Crowdsourcing human oversight on image tagging algorithms: An initial study of image diversity', (2021).
- [28] Hector Levesque and Gerhard Lakemeyer, 'Cognitive robotics', Foundations of artificial intelligence, 3, 869–886, (2008).

- [29] Hao Ma, Jianke Zhu, Michael Rung-Tsong Lyu, and Irwin King, 'Bridging the semantic gap between image contents and tags', *IEEE Transactions on Multimedia*, **12**(5), 462–473, (2010).
- [30] George A Miller, 'Wordnet: a lexical database for english', *Communications of the ACM*, 38(11), 39–41, (1995).
- [31] Ruth Garrett Millikan, On clear and confused ideas: An essay about substance concepts, Cambridge University Press, 2000.
- [32] Ruth Garrett Millikan, 'Neuroscience and teleosemantics', Synthese, 1– 9, (2020).
- [33] A. Mogadala, M. Kalimuthu, and D. Klakow, 'Trends in integration of vision and language research: A survey of tasks, datasets, and methods', *Journal of Artificial Intelligence Research*, 71, (2021).
- [34] Joseph Nassar, Viveca Pavon-Harr, Marc Bosch, and Ian McCulloh, 'Assessing data quality of annotations with krippendorff alpha for applications in computer vision', arXiv preprint arXiv:1912.10107, (2019).
- [35] Roberto Navigli and Simone Paolo Ponzetto, 'Babelnet: The automatic construction, evaluation and application of a wide-coverage multilingual semantic network', *Artificial Intelligence*, **193**, 217–250, (2012).
- [36] Curtis G Northcutt, Anish Athalye, and Jonas Mueller, 'Pervasive label errors in test sets destabilize machine learning benchmarks. arxiv 2021', arXiv preprint arXiv:2103.14749.
- [37] Curtis G Northcutt, Anish Athalye, and Jonas Mueller, 'Pervasive label errors in test sets destabilize machine learning benchmarks', *arXiv* preprint arXiv:2103.14749, (2021).
- [38] Stefanie Nowak and Stefan Rüger, 'How reliable are annotations via crowdsourcing: a study about inter-annotator agreement for multi-label image annotation', in *IEEE-MIPR*, pp. 557–566, (2010).
- [39] Yanwei Pang, Yazhao Li, Jiabing Shen, and Ling Shao, 'Towards bridging semantic gap to improve semantic segmentation', in *ICCV*, pp. 4230–4239, (2019).
- [40] Amandalynne Paullada, Inioluwa Deborah Raji, Emily M Bender, Emily Denton, and Alex Hanna, 'Data and its (dis) contents: A survey of dataset development and use in machine learning research', *Patterns*, 2(11), 100336, (2021).
- [41] Michael Pechuk, Octavian Soldea, and Ehud Rivlin, 'Function-based classification from 3d data via generic and symbolic models', in AAAI, volume 20, p. 950, (2005).
- [42] IB. Raji, EM. Bender, A. Paullada, E. Denton, and A. Hanna, 'Ai and the everything in the whole wide world benchmark', in *NIPS*, (2021).
- [43] S R Ranganathan, Prolegomena to Library Classification, Asia Publishing House (Bombay and New York), 1967.
- [44] Eleanor Rosch, Carolyn B Mervis, Wayne D Gray, David M Johnson, and Penny Boyes-Braem, 'Basic objects in natural categories', *Cognitive psychology*, 8(3), 382–439, (1976).
- [45] O. Russakovsky, J. Deng, H. Su, et al., 'Imagenet large scale visual recognition challenge', *IJCV*, 115(3), 211–252, (2015).
- [46] A. Salari, A. Djavadifar, RL. Liu, and H. Najjaran, 'Object recognition datasets and challenges: A review', *Neurocomputing*, (2022).
- [47] Karen Simonyan and Andrew Zisserman, 'Very deep convolutional networks for large-scale image recognition', in *ICLR*, (2015).
- [48] A. Smeulders, M. Worring, S. Santini, A. Gupta, and R. Jain, 'Contentbased image retrieval at the end of the early years', *IEEE TPAMI*, (2000).
- [49] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich, 'Going deeper with convolutions', in CVPR.
- [50] J. Tang, Z. Zha, D. Tao, and T. Chua, 'Semantic-gap-oriented active learning for multilabel image annotation', *IEEE TIP*, 21(4), (2011).
- [51] Antonio Torralba and Alexei A Efros, 'Unbiased look at dataset bias', in CVPR 2011, pp. 1521–1528. IEEE, (2011).
- [52] D. Tsipras, S. Santurkar, L. Engstrom, A. Ilyas, and A. Madry, 'From imagenet to image classification: Contextualizing progress on benchmarks', in *ICML*, pp. 9625–9635. PMLR, (2020).
- [53] Fei Wang, Mengqing Jiang, Chen Qian, Shuo Yang, Cheng Li, Honggang Zhang, Xiaogang Wang, and Xiaoou Tang, 'Residual attention network for image classification', in *CVPR*, pp. 3156–3164, (2017).
- [54] K. Yang, K. Qinami, L. Fei-Fei, J. Deng, and O. Russakovsky, 'Towards fairer datasets: Filtering and balancing the distribution of the people subtree in the imagenet hierarchy', in *Proc. FACT Conference*.
- [55] S. Yun, Sj. Oh, B. Heo, D. Han, J. Choe, and S. Chun, 'Re-labeling imagenet: from single to multi-labels, from global to localized labels', in *CVPR*, pp. 2340–2350, (2021).
- [56] Matthew D Zeiler and Rob Fergus, 'Visualizing and understanding convolutional networks', in ECCV, pp. 818–833. Springer, (2014).