

Learning for action-based scene understanding

Cornelia Fermüller^{a,*} and Michael Maynard^{**}

* University of Maryland, Institute for Advanced Computer Studies, Iribe Center for Computer Science and Engineering, 8125 Paint Branch Dr, College Park, MD 20742

** University of Maryland, Computer Science Department, Iribe Center for Computer Science and Engineering, 8125 Paint Branch Dr, College Park, MD 20742

^a Corresponding: cornelia.fermuller@gmail.com

Abstract

Action plays a central role in our lives and environments, yet most Computer Vision methods do not explicitly model action. In this chapter we outline an action-centric framework which spans multiple time scales and levels of abstraction, producing both action and scene interpretations constrained towards action consistency. At the lower level of the visual hierarchy we detail affordances - object characteristics which *afford* themselves to different actions. At mid-levels we model individual actions, and at higher levels we model activities through leveraging knowledge and longer term temporal relations. We emphasize the use of grasp characteristics, geometry, ontologies, and physics based constraints for generalizing to new scenes. Such explicit representations avoid over-training on appearance characteristics. To integrate signal based perception with symbolic knowledge we align vectorized knowledge with visual features. We finish with a discussion on action and activity understanding, and discuss implications for future work.

Keywords: affordances, actions, activities, action-based representations, vector space embeddings

Chapter points

- An action-centric framework for scene and activity interpretation.
- Studies on object affordances and functionalities and their use in the context of action recognition and robot learning.
- Studies on activity recognition as an interplay between cognition and perception.
- The merging of vision and language through embedding spaces.
- Discussion on the future of action and activity understanding through the lens of the action-centric framework.

1. Introduction

The purpose of Computer Vision (CV) is to produce interpretations of images and video which are of use to humans. Action is *important* to model because it is a primary means through which others and ourselves interact with our environment, and it

is largely through interaction that the environment becomes meaningful. Because of the centrality of action in what humans find meaningful in their environment, humans structure their environments around action. So to fully understand human environments requires an understanding of their relation to actual and possible action. Most contemporary CV methods are not action based in their approach - in this chapter we present methods and frameworks which in modeling the observed scene employ an *action based* or *functional* interpretation.

The centering of action in perception aligns with embodied cognition theories [67, 5], which argue that many aspects of cognition take their origin in motor behavior and action. In a computational approach we can leverage action based representation at multiple time scales for a hierarchical approach to scene understanding. At the early hierarchical levels are static components, the objects, humans, and simple movements of the limbs. These are then combined into increasingly more complex notions that involve interactions between scene components. Temporally *actions* chain together in structured ways to constitute *activities*.

The use of action based representations in computational perception approaches is *challenging*. The classic approach to CV is to recognize scene constituents based purely on their appearance. However, the aspects of the scene related to action are often semantic and relational rather than appearance based in nature. To better model interactions, more complicated architectures are required that not only model visual appearance but which leverage a more cognitive understanding of the intermediate semantic and relational structure of action in the input.

Classic end-to-end visual learning becomes intractable with larger input state spaces, as are found with video and action of increasing duration. This is because the variability in visual appearance increases, presenting challenges both in data and in modeling. In order to scale, a more cognitive approach which models not the appearance but the action structure of the activity is necessary.

A primary advantage of the action based approach to scene interpretation is generalization, i.e., the ability to recognize scene quantities beyond those visually observed in the training set. For example, if we can recognize what makes an object usable for cutting, this will allow us to recognize new kinds of cutting tools, such as an Alaskan ulu, although this object has not been in our training set. Similarly, if we can interpret an observed human activity by understanding the interaction of constituents and by understanding the underlying goal, we can be more robust. Individual constituents may be difficult to recognize because of occlusions, size, unfavorable viewing angles or variability in visual appearance and movements, but reasoning about the cognitive plausibility of the activity can allow the recovering from classification errors. Furthermore, action modeling provides the potential to predict far into the future.

This chapter presents CV learning-based approaches and concepts centered on action. We now outline in brief the contents of the rest of this chapter.

Section 2 covers *affordances* - a variety of action-based object description. Affordances have been of great interest to Robot Vision. But also classic, non-embodied CV can benefit from the use of affordances. They reflect how the different objects in a scene can be used, and they are an essential component for action understanding. They carry information on the possible co-occurrences of observed objects, humans, and other scene constituents. Section 2 covers the best known CV works on the topic, which include early studies that reason about affordances via geometric measurements, studies that learn affordance maps using algorithms for object detection and semantic segmentation on depth and geometric feature maps, and studies that combine affordances with other constituents for action recognition.

Section 3 is devoted primarily to our own work on understanding manipulation activities. We argue that activity interpretation should be implemented as a continuous interplay between reasoning and perception processes. Activities are modeled hierarchically. At the lower level are modules for objects, actions, spatial relations, etc, which are merged at the higher level via a grammatical formulation. The grammar and selected modules supporting an action-driven understanding are described.

Section 4 focuses on methods that can achieve a tighter integration of appearance and semantic and relational constraints. We consider the integration within the context of the task of Zero-Shot Learning. We cover first simple methods involving engineered attributes, and proceed through more sophisticated approaches involving merging language and vision through shared embedding spaces, capturing semantic and relational information.

This is followed by a discussion on how these concepts could be applied to action and activity understanding in Section 5, and Conclusions in Section 6.

2. Affordances of Objects

Psychologist James Gibson coined the term “affordance” [21], referring to the action possibilities that an object presents based on humans’ (or animals’) physical capabilities. For instance, a knife affords “cutting,” “stabbing,” “poking,” “slicing,” “throwing,” etc. (to a human). The notion of affordance has recently received great interest in the cognitive science and neuroscience literature, strengthened by brain imaging evidence that showed that observing tools activates motor areas of the brain (for a review see [39]). The concept has been studied in different areas, including developmental psychology, industrial design, sport science, and human computer interaction, and there have been many interpretations and discussions on its meaning. Most distinguish between “affordance” and “function,” with the former meaning properties of objects and the latter referring to the role that an object plays in satisfying some purpose. For example the handle of a cup affords “grasping,” and its interior “containing,” while an electricity plug supports the function to “powering kitchen appliances,” or “charging

devices,” and a water faucet supports the function of “getting drinking water.” However, a formal definition does not exist.

In this section, we first motivate the use of affordances in CV (Sec. 2.1). Then different works from the literature are discussed: Sec. 2.2 is about the earlier approaches, which selected geometric features computed from 3D data to classify affordances of chairs or everyday objects. Sec. 2.3 describes works on learning affordances of objects and their parts using CV recognition algorithms applied to depth data or geometric feature maps. Sec. 2.4 describes approaches using affordances together with other detectors for scene and action recognition, and approaches that learn affordances for embodied agents. Sec. 2.5 concludes with suggestions for future work.

2.1. Why Would Computer Vision be Interested in Affordances?

Looking at objects and scene surfaces from the viewpoint of affordances provides information for visual scene interpretation that is complimentary to the classic cues and aids in robustness and generalizability of learned representations. This information is about the “actionability” that the scene presents at multiple spatial and temporal scales relating to objects, groups of objects, and the complete spatio-temporal scene. Therefore affordances provide information and constraints for scene understanding both in the present and in projecting into the future - thus aiding recognition in addition to prediction, as detailed next.

Models of affordances learned over some objects are transferable to novel object categories. I.e., if our recognition modules can recognize an affordance, they can detect it in objects never seen before, even in a stone that has the right properties. This is because how an object is used depends on physical properties such as its shape, size, material, and weight [27], and we can design processes that pick up these physical properties from images, depth maps, and other modalities, independent of previously encountered object categories. In contrast, classifying objects in images in a conventional end-to-end fashion does not give insight into how visual features such as affordances relate to the object.

Affordances provide valuable information to visual object understanding, such as in understanding the “valid functionality” of objects [25] - e.g., an inverted cup cannot be used to pour into, or similarly a broken chair cannot be used for sitting [22]. Another example is the sub-categorization of the classical visual object categories, such as differentiating between chairs for different uses [61].

Since affordances represent the possible actions that can be performed with an object, they carry valuable information for predicting future actions [33, 53] - because actions relate to each other over time. For example, a bread knife as a whole presents affordances (“graspable”, “cut with”) allowing the action of “slicing bread,” and slicing bread is part of the activity “preparing the bread basket” - an activity consisting

of multiple actions extended through time with temporal dependencies. Knowledge of the possibility of “slicing bread” informs possible subsequent actions such as putting the basket on the table. To summarize, affordances and functionalities at the object level also contain information about possible object interactions, spatio-temporal relations, and activities at longer temporal scales. Modeling these relations to get explicit or implicit relations at multiple time scales and semantic levels of abstraction has value for the task of activity understanding.

The concept of affordances has been central to Robot Vision and to research along the Active Vision Paradigm [4]. The latter advocates that the vision of systems should not be considered a passive process. Biological systems “move their eyes to select what they see” in an active process. Similarly, artificial embodied systems should be able to change the viewpoint of their cameras in order to select what information to gather from their environment, as different viewpoints present different information. Going further, the paradigm also suggests that embodied systems should avoid employing heavy general-purpose vision processes for all purposes, and only process the information necessary to solve the task at hand [15]. Therefore, when a robot or artificial system interacts with objects, often it is more effective to compute what an object can be used for - i.e., compute its affordance and how it can be used - rather than to classify the object according to our language representations. Thus, while the advantages of affordances discussed in this section apply to the classic passive CV formulation, where there is no agent interacting with the environment, a great portion of the research on affordances focuses on Robot Vision.

2.2. Early Affordance Work

Affordances relate to actions. As a consequence they are also grounded in action related physical quantities. For example, an object to sit on or an object to pour into have certain physical quantities, e.g., a certain shape, size, or certain material, etc. All of the earlier approaches utilized such explicit physically meaningful representations in affordance recognition modules.

The first studies used shape and geometry. Stark and Bowyer [61] proposed the first affordance-based approach to object recognition using 3D CAD models as input. A knowledge-graph, similar to a decision tree, was created to classify chairs and sub-categories of chairs (e.g., conventional chair, balans chair, high chair, lounge chair), where the leaves of this graph were procedures for classification of geometric features. These features included relative orientation between surfaces, object dimension, stability, and proximity of surfaces.

Grabner *et al.* [22] detected surfaces that afford “sitting,” by checking the geometry of a 3D human skeleton model in a sitting pose against the object’s geometry. Their features include distance and the intersection of the human’s mesh with the object’s

mesh. The detector was evaluated on Google Warehouse models as well as real 3D data collected with a time of flight camera. For best performance the method was combined with an image based classifier. Similarly, Gupta *et al.* [24] modeled affordances in 3D indoor scenes by detecting the regions of the space which allow a human to use it for one of three functions: “laying down,” “sitting upright,” and “sitting reclined.” Like [24], they also used constraints based on the occupied 3D space and the contact with a human skeleton. However, their method can take as input images, from which it first derives 3D geometry via learning-based regression methods such as [26, 35].

Hermans *et al.* [27] learned the affordances of everyday objects via intermediate representations that encode visual and physical characteristics. Visual characteristics included color, discrete shape, and texture, and physical characteristics included weight and size. Standard classifiers were used in the pipeline, and the approach was demonstrated on seven affordance classes in the robotics domain.

2.3. Affordance Detection, Classification, and Segmentation

The problem of recognizing affordances associated with objects and scene surfaces is conceptually similar to the problem of object recognition. A number of recent approaches have used tools from object detection, classification, segmentation, and semantic labeling for affordance localization and recognition. However, these techniques usually were not applied to images, but instead either to RGBD data or to feature maps computed from depth data. This section discusses a few such approaches.

2.3.1. Affordance Detection from Geometric Features

This section describes the work of Myers *et al.* [46], the first approach applying modern machine learning tools on geometric features. The section details the approach to affordance detection and discusses computational implications.

The focus of the study were tools used in everyday workspaces, and specifically the detection of tool parts associated with different affordances. A dataset (the RGB-D Part Affordance Dataset) of 105 kitchen, workshop and garden tools was collected. Objects were put on a revolving turntable and recorded with a Kinect camera from a full 360° field of view, about 300 frames for each object, out of which 10,000 RGB-D images were annotated at the pixel level. Fig. 0.1 shows example objects for five of the seven affordances, along with the annotation for one of the objects. It should be noted that affordance is associated with surfaces, for example the inner surface of a cup is “contain” while the outer surface is “wrap-grasp.”

From the raw depth data, shape features were computed patchwise, specifically, the surface normal, principal curvature, shape index, and the HoG-Depth descriptor (histogram of depth gradients). Using these features as input, two classification approaches were proposed: first, a Structured Random Forest (SRF), which creates point-wise classification; and second the S-HMP (Superpixel Hierarchical Matching



Figure 0.1 Sample objects from the RGB-D Part Affordance Dataset, and an example of a full frame image with hand-labeled ground truth (at the lower right). The ground truth labels include rankings for multiple affordances (from [46]).

Pursuit) algorithm [6]. The latter works by first oversegmenting the RGB-D image into superpixels. Then, using a dictionary learning technique, the shape features are sparsely encoded at multiple scales per superpixel. Finally, the features are max-pooled over the super-pixels and classified via an SVM. Example results are shown in Fig. 0.2 for both the S-HMP and the SRF method method, where the gray value encodes the probability for the affordance assignment.

There are two computational aspects to the approach discussed above, that deserve special attention. First, sometimes overlooked, the assignment of affordances to object surfaces in general may not be unique. The same object part may be used for multiple purposes. Assigning affordances is thus a multi-class labeling problem. In [46] this issue was addressed by having multiple annotators rank how close other affordances were with respect to the essential affordance, from which an ordinal scale for affordance assignment at testing was derived.

Second, a main advantage of the approach is its good generalization to new objects and surfaces. Referring to Figure 0.2 (Bottom), one can see that the bottom of a cup is classified with the affordance “Pounding” and the edge of spatula with the affordance “Cutting.” This is because the shape of these objects indicates these properties. However, shape by itself would not be sufficient for classification in a practical system. One would have to add additional properties, the most obvious is material. This would allow to decide that a paper cup cannot be used for pounding, or an object with a soft edge cannot be used for cutting.

2.3.2. Semantic Segmentation, and Classification from Images

Many of the works that followed [46] employed neural network approaches using geometric feature maps as input. Affordance detection at pixel level, thus became a semantic labeling problem. However, different from [46], these approaches often

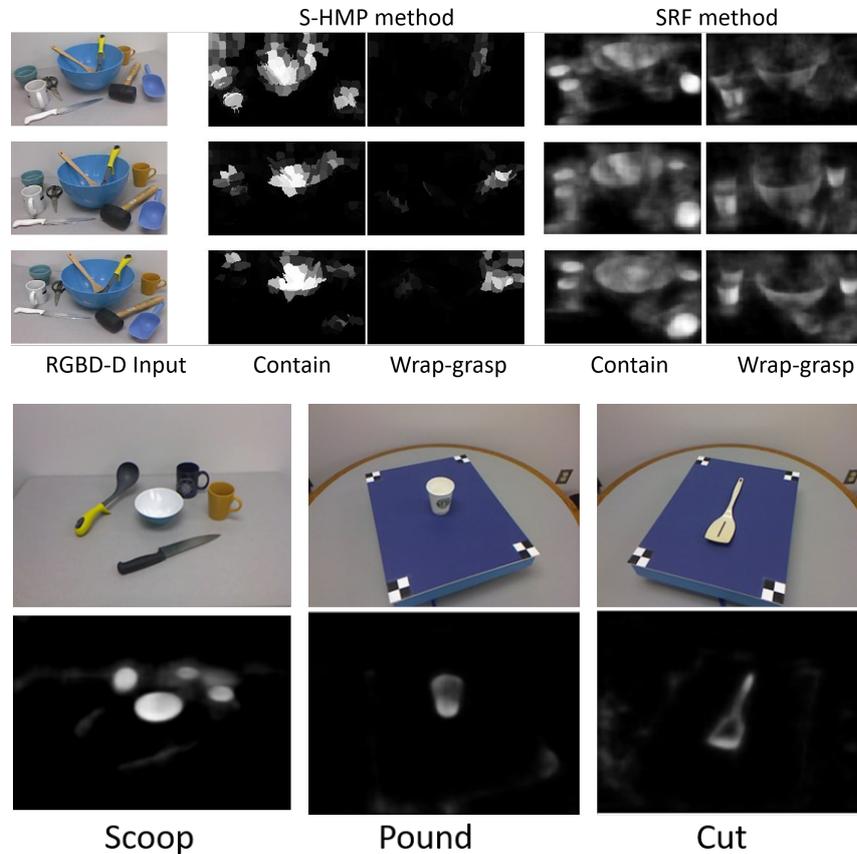


Figure 0.2 (Top) Results of affordance detection across three input RGB-D frames using the S-HMP and the SRF method over a cluttered sequence for the target affordances “contain” and “wrap-grasp”. Brighter means higher probability of the target affordance. (from [46]) (Bottom) Demonstration of the generalization of the method for new objects for the SRF-method: the bottom of the cup was detected with high probability for “Pounding” and the edge of the spatula with high probability for “Cutting.”

used 2D images as input. In a preprocessing step, depth maps or feature maps were regressed via neural networks. Furthermore, some considered natural images with multiple objects, and employed object detection algorithms to localize the objects, before assigning affordances.

For example, Nguyen *et al.* [47] created a dataset with ten object categories and nine affordance categories from the household and workshop domain. It consists of both RGB-D scans and natural images (a subset of ImageNet [55]) - for the latter depth maps were created using the CNN approach of [37]. The images were

annotated with bounding boxes and affordances at the pixel level. The paper’s method first applied an object detector, then within each region computed affordances using a modified VGG-16 network trained for semantic labeling, and finally the affordance values were post-processed with a CRF.

Srikantha and Gall [60] used the dataset of Koppula and Saxena [32], which features rich contextual information in terms of human-object interactions, and curated it with pixel-level affordance annotations. The work explored different levels of supervision for semantic segmentation, using a deep convolutional neural network within an expectation maximization framework to take advantage of weakly labeled data like image level annotations or keypoint annotations, as well as human pose as context.

Roy and Todorovic [54] worked with the indoor scenes from the NYU dataset [58]. Their approach first infers the depth map, surface normals, and coarse-level semantic segmentation using a multi-scale CNN as mid-level cues, which are then jointly fed as inputs to another multi-scale CNN for prediction of the affordance maps.

Ye *et al.* [77] designed a method for localizing and recognizing functional areas in indoor scenes. An ontology, as shown in Fig. 0.3 (Left), was defined to categorize image regions according to their affordance or functionality. Categories include: “open with spherical grasp” (such as a door knob), “open with wrap grasp or drag to open” (such as an oven door), “turn on electricity” (such as light switch), etc., as shown in the second last column of the Figure. The dataset has 500 images featuring kitchens from the SUN dataset [71], which were curated. The method first runs a CCN-based detector trained to detect the region and then a classifier based on a VGG architecture. Fig.0.3 (Right) shows example results.

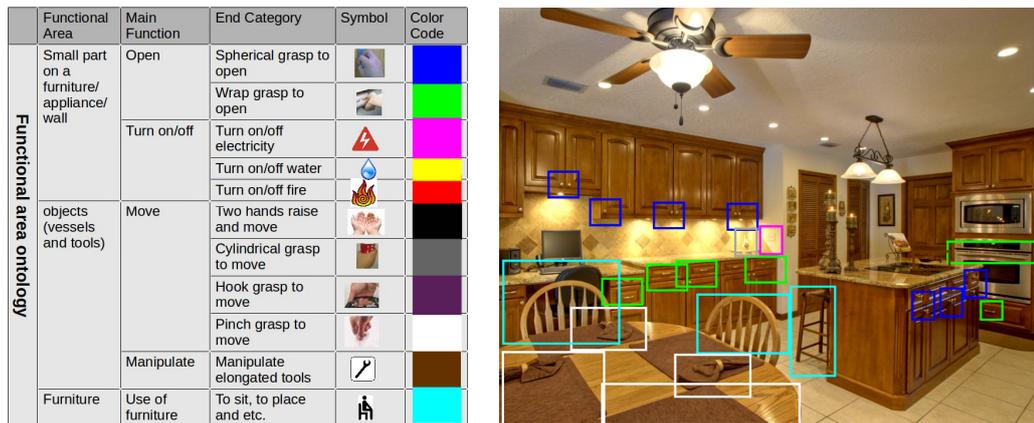


Figure 0.3 (Left) Functionality Ontology (Right) Sample detection results (from [77])

2.4. Affordance in the Context of Action Recognition and Robot Learning

In this section we highlight a few approaches that used affordances in conjunction with other quantities for scene and action understanding. We then discuss approaches that addressed the learning of affordances by robots.

2.4.1. Action Recognition

Affordances encode features of the possible interactions a human can have with the environment. Thus, naturally they provide a glue between different quantities in the scene in space-time, such as between different objects, or between objects and actions. A number of studies have built on this idea, and used affordance relationships as a context for activity and action recognition and prediction. These methods employed various models to encode relations between the different quantities involved, including CRFs, MRFs, And-Or-Graphs, and probabilistic state automata.

Kjellström *et al.* [30] investigated the problem of learning action-object interactions from demonstration, which they define as affordances. Hand actions were classified in the context of the manipulated objects using a CRF that gets as input object and hand features. Objects were modeled using hand-crafted features, and actions were modeled by the hand’s global velocity, orientation, and joint angles, which were computed from the output of a 3D hand reconstruction and tracking method.

Koppula *et al.* [34] considered the problem of learning sequences of subactivities performed by humans and their interactions with objects. They jointly modeled the human activities and object affordances in a Markov Random Field where the nodes represent objects and sub-activities, and the edges represent the relationships between object affordances, their relations with sub-activities, and their evolution over time. Affordance-subactivity relations were computed from relative geometric features between the object and the human’s skeletal joints, and affordance relations between objects from spatial relations. The description was demonstrated for a PR2 robot in performing assistive tasks. In [33], Koppula and Saxena added to the Markov Model also possible future states in order to predict the next action.

Qi *et al.*[52] used a Spatial-Temporal And-Or Graph (ST-AOG) to represent the structure of activities and predict future actions in RGBD video input. Their model is hierarchical: sub-activities are modeled by the human action, the objects, and their affordances in spatial graphs, and a stochastic grammar defined over the sub-activities encodes the activity. Dutta and Zielinska [11] also considered the problem of predicting the next action based on object affordances and human interaction. They employed spatio-temporal based probabilistic state automata to model the interactions. In addition to the action class, they also computed the possible action trajectory. Depending on where an object is relative to the human it has different affordances, and depending

on the affordance, its orientation and distance to the possible action trajectories were encoded as heatmaps.

2.4.2. Affordance Learning in Robot Vision

Affordance has been a central concept in the field of neurorobotics, which aims to understand the cognition of a system whose body is embedded in the environment. In this research, robots acquire increasingly more complex skills using perception and interaction with the environment. Through interaction robots learn affordances, and build upon these hierarchically an understanding of actions, activities and the environment. This research in developmental robotics was enabled by the development of robotic platforms, best known among them, the humanoid robot, iCub [40].

In [17] the authors discussed three broad stages in the development of a robot: first learning a body image, second learning the interactions with external objects, and third learning to interpret object-object interactions. Affordances are central to the latter two stages. The humanoid robot through pushing and pulling actions in different directions learned to interact, and by observing affected objects’ movements learned affordances, such as whether a spherical object is rollable and a cuboid is slide-able. Finally, the robot also learned to mimick an observed action.

Similarly, the authors of [45] defined the three main stages in the architecture of a developing humanoid robot as sensory–motor coordination, world interaction, and imitation. Affordances play a central role for world interactions. In this approach, the system started with basic vision and motor skills from which more complex vision and motor skill were acquired using clustering algorithms. Then, during interaction, effects were observed using perception, such as the changes in object position, velocity, and tactile sensing. A Bayesian networks was used to learn affordances, which in this case were encoded as probabilistic relations between actions and percepts (object features and effects). The system was demonstrated to imitate the actions of humans by performing movements with similar effect.

Ugur *et al.* [65] also demonstrated a robot learning object affordances through interaction and self-observation. In a first step the robot discovered commonalities in its action-effect experiences by discovering effect categories. Building upon these, in a second step, affordance predictors for different behaviors were obtained by learning the mapping from the object features to the effect categories. Ugur and Piater [66] went a step further and studied mechanisms that produce hierarchical structuring of affordance learning tasks. Guided by intrinsic motivation, the robot started with easy tasks, and building on its knowledge of interactions progressively learned more complex tasks by selecting to explore the object and action most different from previously explored ones. For the experiments the robot could compute the visual features of object dimension, surface patch shape, and surface normals, and its actions were poking from three different directions and stacking. In earlier stages it explored the poking

actions to observe their effects on single objects. Building upon these, it then explored in a second stage the stacking of two objects and resulting effects.

2.5. Discussion on Affordance Learning

This section discussed approaches to affordance learning, many of which fall into the domain of Robot Vision and have been conducted with few examples and limited amounts of data. So far, deep learning approaches have not been much used for affordance understanding. The major reason is the lack of large annotated datasets in this domain, necessary for deep learning.

However, we expect that as research shifts away from supervised to unsupervised and self-supervised approaches, we will see learning approaches building on the concept of affordances and observed interactions between humans and objects. This will be facilitated by datasets, such as the EPIC Kitchens dataset, which features a variety of manipulation actions in natural scenes [8].

Affordances and functionalities at the object level also encode information about possible object interactions, spatio-temporal relations and possible activities at longer temporal scales. We have discussed in Section 2.4 approaches using affordances for action modeling. However, in future work, we could model these relations to get explicit or implicit relations at longer time scales for the problem of activity understanding, the topic of Section 3.

Finally, when creating mappings from perception to action for robot learning, we may ground them in affordances. Humans can learn manipulation actions using their perception only. When we see somebody performing actions with a tool unfamiliar to us we can understand the tool’s affordance and perform the same action. Similarly, we could approach robot motor learning using perception and action in a tight loop, grounding them in affordances, something that has not yet been done. The robot would learn the task by observing the action and affordances and issuing commands (based on its existing skill set constrained by affordances) to generate the action approximating the observed one, and then adapt gradually to improve performance. The suggested research tasks then amount to developing self-supervised learning and reinforcement-learning approaches grounded in affordance-based representations.

3. Functional Parsing of Manipulation Actions

This section describes work – largely from our group – on the interpretation of manipulation activities. Inspired by the embodied cognition paradigm [67], this work considers the understanding of human activities a process that involves perception, cognition, and the motor system. The major components are a formalism to combine the different modalities, and CV modules for obtaining semantically meaningful descriptors of action.

3.1. The Active Interplay between Cognition and Perception

Understanding human actions and activities is the most challenging task currently studied in CV. It is not a task of vision only. Humans can understand what others are doing, because they have models of actions and activities. They understand the goals of actions, and this allows them to interpret their observations despite the large variations in which actions can be executed and variations in visual conditions. Knowledge of some form comes into the interpretation process quite early.

We observe that human behavior is active and exploratory. We continuously shift our gaze to different locations in the scene. We recognize objects and actions and this in turn leads us to fixate at new locations. In this process, perception continuously interacts with cognition at different levels of abstraction: to guide attention, to make predictions, to constrain the search space for recognition, and to reason over what is being perceived. We call this interaction between perception and higher level processes the *Cognitive Dialogue* [2], as it amounts to an iteration of questions and answers, with the cognitive or linguistic processes asking questions about the what and where of quantities in the scene, and the visual processes performing localization, detection, recognition, and reconstruction. A possible simple way to selecting the next question would be by using information-theoretic criteria [78].

The reasoning can be implemented through knowledge-based engineering [1] or the use of language. There has been much interest in CV to introduce additional higher-level knowledge about image relationships into the interpretation process. While many studies get this additional information from captions or accompanying text, others (as discussed in Section 4) use advanced language processing to obtain additional high level information. In current research, most commonly, the Word2vec space [41] (see Sec. 4.3) is used as language representation, which encodes similarity about linguistic concepts. Alternatively, one could use older, hand-crafted resources encoding lexical semantics, for example the Word-Net database [42], which relates words through synonymy (words having the same meaning, like “argue” and “contend”) and hypernymy (“is-a” relationships, as between “car” and “vehicle”), among many others. Verbnets [56], which organizes verb classes, is particularly interesting for action understanding.

3.2. Grammars of Action

Various mechanisms have been used to encode relationships between the different semantic concepts, that is between actors, objects, verbs, spatio-temporal relations, and attributes. Section 2, discussed the use of and-or-graphs and Markov models. Others include Markov Logic Networks [64], and planning tools [23]. In this section we describe work on grammars, which can capture the composition of observed activities as sequences of scene constituents, and their recursive structure.

The main motivation for the use of grammars originates from the idea that actions observed in a video have syntactic structure. By considering the goal of actions, the video can be broken into meaningful segments, and these segments together can be organized in the form of a simple grammar. Thus, interpreting the action that is taking place in a video is like understanding a sentence that we read or hear. To parse the video into the primitive actions that constitute complex tasks, the segments of the video are mapped to particular symbols involving objects, tools, movement, and spatial relations. Importantly, the action grammar temporally segments a video at contact, that is, when the hand touches or releases an object, or objects merge or separate. At these points in time a new sub-action starts. In applying the grammar for the analysis of a video, a parse tree is produced, which we call the *activity tree*. Figure 0.4 illustrates the concept. From a video recording of a person performing the activity “cutting a plank,” a graph is created, in which nodes of hands, objects, and tools merge into a common node, whenever they touch, or nodes split when the objects and hands separate. Referring to the figure, to compute the quantities involved, different processes (shown in the four sub-areas of the video frames) extract the human’s body, the hands, objects, and geometric relations.

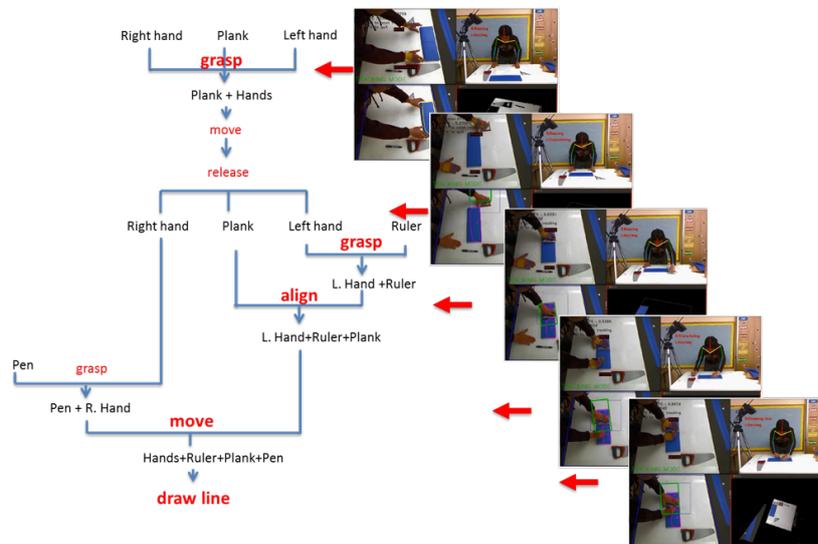


Figure 0.4 Illustration of activity description. The camera monitored a person cutting a plank. Four parallel processes computed essential components: (left up) detection of the hand and classification of the grasp type, (right up) gross motion via skeleton fitting, (left down) segmentation of the object, (right down) 3D shape description of the scene.

We now discuss several grammar approaches in Section 3.2.1, and then discuss

in Section 3.2.2 whether such grammar representations are sufficiently expressive to capture action and activity structure, and sufficiently parsimonious to be preferred over other representations.

3.2.1. Different Implementations of the Grammar

The descriptions are based on context-free grammars, originally introduced in [50]. Summers-Stay *et al.* [62] implemented the idea for parsing assembly actions from RGBD video using only one symbol for all actions, and Yang *et al.* [75] enhanced the description introducing grasp into the description and differentiating hand-to-object contact and object-to-object contact. The grammar describes actions at a level of abstraction that is useful for both video interpretation and robot execution. In [76] this was demonstrated with some examples. By automatically parsing videos that feature cooking instructions from the Youcook dataset [9], actions were parsed, which then were performed by a Baxter robot equipped with the necessary motion capabilities.

Abstract descriptions of actions/verbs are necessary to achieve generalization, and perform what is called in the current terminology, few-shot learning or zero-shot learning (see Section 4). The basic action grammar reduces the description of actions to only the sequence of “touching relations,” that is when the hand touches an object, two objects touch, the hand releases an object, or two objects or pieces of a single object separate [10].

Wörgötter *et al.* explore this concept to formulate an ontology [68] considering one-handed actions. At the first level, actions are categorized according to the sequence of relations that the hand and one or two objects can have, into six classes, which are: rearrange, destroy, break, take-down, hide, construct. From there they iterate possible actions, and they come up with about 30 fundamental manipulations. Yang *et al.* [73] proposed a related concept. They meta-classified actions according to the consequence an action has on an object, that is, what happens geometrically or topologically to an object. They proposed six categories – divide an object, merge two parts, transfer an object, deform an object, object appears, object disappears from the scene – and they also provided algorithms that combine tracking with segmentation to detect topological changes to detect the essential events in video.

3.2.2. Are Grammars Expressive and Parsimonious Descriptions?

An important question is whether the grammatical representations actually are sufficiently rich to allow for classification of many activities. The authors of [69] performed psychophysical and computational experiments to answer this question. They described, as above, actions by the sequence of contacts, using five quantities: the hand, the ground, and three objects. In addition they considered ten spatial relations, i.e., above, below, between, etc., to differentiate between altogether 35 different configurations or actions. A subset of ten of these actions (put, shake, stir, take, uncover,

chop, cut, hide, lay, and push) was performed in a virtual environment, but instead of the actual objects, cubes were used. Experiments found that humans can recognize these actions, and so can their algorithms. Even more, the description was found to be very powerful for prediction; subjects on average only required 56% of the action duration to recognize the action. Thus, it appears that a description relying only on contact and spatial relations is very powerful for visual recognition.

3.3. Modules for Action Understanding

Individual vision processes are required to recognize discrete components, which then can be combined in higher level reasoning processes - such as the grammars from Section 3.2 - to achieve activity recognition and prediction. The descriptors which we discuss in this section differ from those heavily covered in the literature (for an evaluation of successful concepts in current approaches see [57]). Specifically, in Section 3.3.1 we discuss representations of grasp, and in Section 3.3.2 we discuss explicit representation of geometry.

3.3.1. Grasping: An Essential Feature for Action Understanding



Figure 0.5 (Left) Rest or Extension on the handlebar vs. (Right) Firm power cylindrical grasp of the handlebar. (from [74])

The grasp type provides crucial information about actions. As a motivational example, consider the two scenes in Fig. 0.5 from the VOC challenge [14]. Standard CV systems have object and human detectors to recognize the bicycle and the cyclist and pose detectors to confirm that these two cyclists are riding a bike. But humans can tell that the cyclist on the left side is not racing (since his hands are in a “Rest or Extension” grasp), whereas the one on the right is intent on racing (since the hands firmly hold the handlebar in a “Power Cylindrical” grasp).

We cover here two papers, the first employs a basic ontology of grasps types in action understanding tasks, the second studies subtle changes in grasp for differentiating between similar manipulation actions, and develops learning approaches for online action prediction and regression of associated finger forces.

The recognition of grasp type provides essential information for a more detailed analysis of action. Researchers in several areas, including robotics, developmental

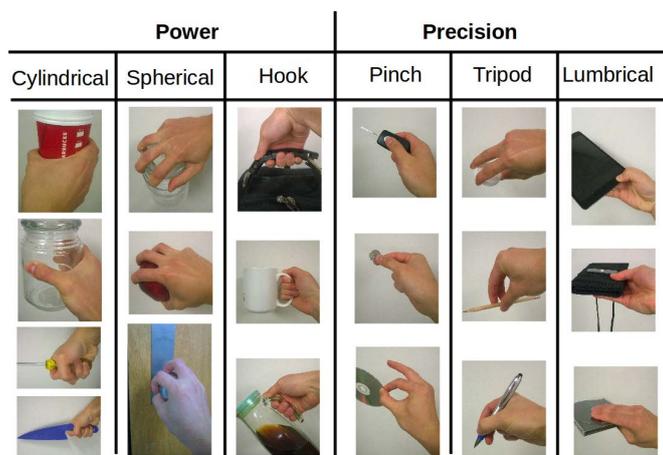


Figure 0.6 Basic classification [7] of active grasps with examples. At the highest level, grasps are categorized into power and precision grips. Power grips are used when an object is held with force, and can be classified as cylindrical, spherical, and hook. Precision grips provide fine movement and accuracy, and are subdivided into pinch, tripod, and lumbrical.

medicine, and biomechanics, have developed grasping taxonomies that represent a hierarchy of the most common hand postures used for object grasping, with each taxonomy based on the needs of the tasks in the field. In [74] a basic classification of the main functional grasps [7] in manipulation tasks was used and then demonstrated as a useful feature in two tasks: for segmentation of activities involving fine motor actions, and for characterizing action intention, i.e., whether the task is casual, or requires skills or forces.

Cognitive studies showed that an actor’s intention shapes his/her movement kinematics during movement execution [3]. For example, when subjects grasped a bottle for pouring, the middle and the ring fingers were more extended than when they grasped the bottle with the intent of displacing, throwing, or passing it. Inspired by these findings, in [16], the authors developed a recurrent neural network architecture that monitors hands for predicting actions. Specifically, they considered sets of actions with the same object, such as “squeezing”, “flipping”, “washing”, “wiping” and “scratching” with a sponge (see Fig. 0.7). They analyzed the system, which predicted

in real-time the ongoing action to determine at what point in time the classification became accurate, and they also performed a psychophysical experiment, evaluating human performance on the same task. At 10 frames after the contact of the hand with the object, the system and the humans started understanding the action (75% classification accuracy for the sponge actions), and at 25 frames the judgement was very good (95% accuracy for the sponge actions). The visual architecture was an RNN using as input tracked image patches around the hand from which VGG-16 features [59] were computed.

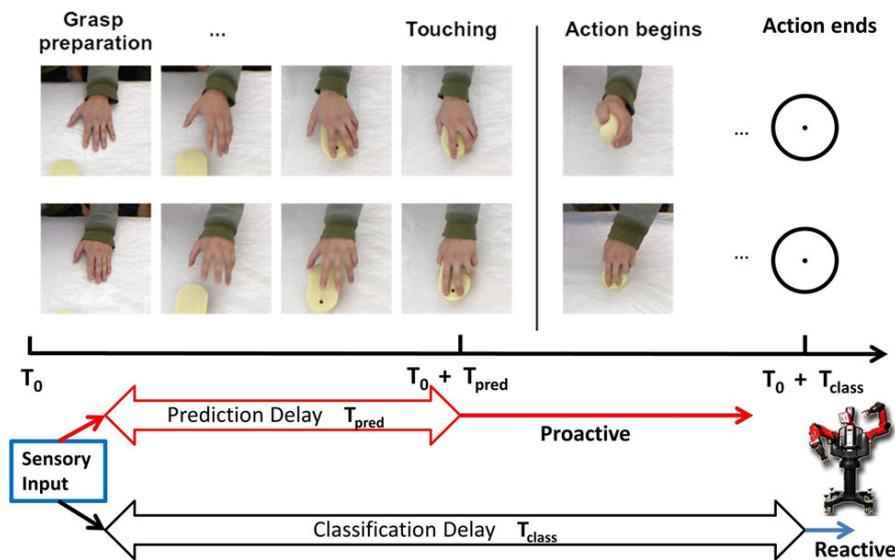


Figure 0.7 Examples demonstrating that early movements are strong indicators of the intended manipulation actions. Early prediction of action significantly reduces the delay in real-time interaction, which is fundamentally important for a proactive system.

In addition the paper also demonstrated association of vision with forces. Data was recorded from subjects that performed the same action with both hands. Sensors on the fingers of one hand recorded the forces, and the other hand was recorded visually. A recurrent neural network was trained to regress from vision to forces. It was then shown that using as input video only, when the visual classifier was combined with regressed forces, improved performance could be achieved. The concept appears promising. As shown, learning the mapping from vision to forces creates a bi-modal space that can aid visual recognition. Furthermore, there are immediate implications for robotics. Currently robots rely on haptic devices or force and torque sensors to learn tasks. If we can predict the forces exerted by the human demonstrator visually, it would allow us to teach robots much more efficiently.

3.3.2. Geometry to Robustify

The use of geometry is important because it provides robust information for scene description and additional information for recognition. Geometry is computed using reconstruction processes, which are low-level (requiring only the image features and knowledge of camera positions), but no machine learning or training data is required. With the advent of cheap RGB-D sensors more than a decade ago, reconstructing scene geometry has become much easier and more accurate, and thus these sensors have become the standard vision sensors in robotics. Their use facilitates computing accurately and fast the distances to control the robot’s movement, as well as computing the geometry and shape of objects to aid scene interpretation. This section discusses three geometric methods: accurate tracking of non-rigid object transformations and detection of topological changes; computation of pairwise spatial relations of objects over time; and computation of object symmetry and its use for better foreground-background segmentation.

Building on an efficient point cloud library [79], Zampogiannis *et al.* [80] developed a technique for accurately tracking non-rigid object transformations and detecting the topological changes that is, contacts and separations of body parts and objects needed for the grammatical description (see Section 3.2). The gist of the method lies in a warp field estimation that considers forward and backward warps between consecutive frames to detect regions of the deformed geometry that undergo topological changes.

Activity descriptions can also profit from descriptors of spatial relations between objects. In [81], the authors introduced a representation for manipulation actions based on the evolution of the spatial relations between objects in the scene. The method was implemented by tracking objects in RGBD video, and reasoning over the spatial relations of observed object pairs. The resulting descriptor amounts to a sequence of spatial relation predicates of the form in, left, right, front, behind, below, above, touch, and it was shown to be sufficiently expressive for distinguishing between four different actions.

Another concept is to exploit general knowledge of object shape properties. For example, symmetry detection can help with segmentation both in 2D [63] and 3D [12]. Imagine looking at a cluttered scene. Since most objects we work with are symmetric, either bilateral or rotational, we can “fill in the back of the object” that is not visible, and this aids the segmentation and recognition.

3.4. Discussion on Activity Understanding

Activity Understanding is a very challenging problem. End-to-end solutions do not scale well to activity understanding because of the large variations in appearance at high levels of abstraction and temporal extension.

We discussed hierarchical approaches, and we detailed one higher level description – action grammars. Action grammars can segment activities at times of contact and capture the recursive structure of action sequences, similar to that found in language. We described experiments demonstrating the expressive power of action grammars. We also described processes at the lower level, which have not received much attention in CV, but which are essential for supporting an action-based approach. These include affordances, grasp-type and the use of geometry to aid temporal and spatial segmentation and descriptions of spatial relations.

In this section we emphasized the necessity to utilize meaningful action-based representations at different levels of the hierarchy. To elaborate further, it is very important that these representations need to be robust, because of the many challenges involved in activity understanding. We can achieve robustness in part through the use of geometry - geometry does not require memorization, and can be estimated from low-level measurements. Thus, we should introduce geometry into the pipeline whenever possible before starting with recognition. Beyond geometry, any concept that provides universally true information is meaningful for activity understanding. We may model physical laws. We may include model ontologies to aid generalization, for example by grouping verbs according to the effect they have on objects [73]. We also may include processes that model causality; actions constrain each other causally, some combinations are not possible physically. These representations capture more knowledge and better constraints for activity interpretation.

The integration of vision and cognition or language is hard. This is because of the semantic gap, i.e., the disparity between the symbolic or linguistic representations and the visual representations based on signals. We want an integration of the two which is not brittle. Thus, we need to avoid setting thresholds or converting to purely symbolic representation too early within the pipeline. This is because if vision fails to return the right quantities, then imprecisions compound through further abstraction, resulting in failed reasoning. The next research challenge is to study learning approaches that relate perception with higher-level reasoning for a deeper integration. Section 4 covers deep learning approaches useful for such integration. Currently such approaches are primarily constrained to object recognition, and to an extent to action recognition. Activity understanding would benefit from these methods.

4. Functional Scene Understanding through Deep Learning with Language and Vision

Here we consider the merging of vision and language and associated representations - the merging of “signal” and “symbol”. The merging of multiple representations is important due to the suitability of different representations for capturing different characteristics of the world. Here we seek to allow information from lower level represen-

tations of appearance and information from higher level representations of relations and semantics to complement each other.

Systems involving symbolic and continuous representations often have hard boundaries, below which the system is continuous, and above which the system is symbolic. Precisely where this boundary is set varies, but generally does not fall below the level of abstraction reflected in human language as language is a primary source of symbolically represented world knowledge.

Symbolic representation is more important for action and activity understanding than for other CV tasks such as object detection as the nature of this task is more abstract and less appearance based. Action has temporal structure at multiple scales, and is structured around satisfaction of conditions - the defining characteristics of action are semantic and relational.

Many computer vision tasks can benefit from the integration of vision and language. However, one task which is ideal for the study of the integration of the two is Zero Shot Learning (ZSL) - this is because unlike other tasks it cannot be solved without the introduction of non-visual knowledge such as is reflected in language.

ZSL is a task with two sets: a training set, and a test set. The categories of these sets break into “seen” and “unseen” categories. The training set consists of only “seen” categories, while the test set contains “unseen” categories as well as, optionally, “seen” categories. To illustrate, there could be a ZSL task which includes “run” and “stand” in seen categories, and “walk” in unseen categories. The task would then be to learn to visually recognize and properly categorize walking, when walking has never previously been visually encountered, but running and standing have been visually encountered.

There are multiple approaches to ZSL. Early work on ZSL focuses on *attributes* - visually recognizable characteristics with differential class (e.g., object classes or action classes) associations. Attributes are generalizable in that attribute detectors trained only on the seen set are able to detect attributes in samples from both the seen and unseen set.

More recent work on ZSL tends to focus on *semantic embedding spaces*. These are Euclidean vector spaces where semantic categories, such as is reflected in language, are associated with vectors - or points in space. These vector representations of words are of significantly lower dimension than naive 1-hot encodings (vectors where each dimension corresponds to a class, and all but 1 dimension have 0 values), and have the quality that words which are similar in semantics are associated with similar vectors nearby in the embedding space.

With semantic embedding spaces, symbolic representations are vectorized in such a way that the semantic relations of the symbol categories are preserved. This allows for integration of symbolic semantics into deep architectures, whose internal representations consist of vectors. This integration amounts to properly aligning the visual vector representations with the vectorized symbolic representations.

Different ZSL methods use different shared embeddings - some embed visual features into the semantic space, some embed the semantic space into the visual feature space, and some embed both into a third shared space. Once both visual and semantic representations lie in the same space, categorizing visual input is a matter of finding the nearest semantic label in this space.

State-of-the-art contemporary ZSL methods often rely on CNNs for visual features and produce shared embedding spaces with pre-trained semantic embedding spaces such as word2vec [38, 70]. Some shared embedding based ZSL methods structure and train their models in an end-to-end fashion, which is more challenging, but provides dividends in performance [82].

The remainder of this section is structured as follows: In Section 4.1 we detail a simple use of attributes for ZSL; in Section 4.1 we detail a more nuanced *relative attribute* formulation; in Section 4.2 we cover the use of shared semantic spaces in ZSL; in Section 4.3 we cover basic approaches to semantic vector space construction; in Section 4.4 we cover the incorporation of knowledge in the form of graphs for ZSL action classification.

4.1. Attributes in Zero-Shot Learning

The use of attributes - including action centric attributes such as affordances covered in Section 2 - in recognition allows the construction of classifiers which are humanly interpretable and specifiable. Attributes mitigate the issue of opacity through use of an explicit pre-defined mid-level representation below the level of class categories.

Use of attributes allows an easy mechanism through which to learn visual representations from the available training data and to transfer those representations onto the classes for which no training data is available. Attributes are general in that they have presence across multiple class categories, and through taking multiple attributes with different distributions across classes allow representation of select classes.

The use of attributes in ZSL is as follows: Attribute detectors are trained over the seen set and associated attribute labels. These detectors are generalizable across both the seen and unseen set. Due to the different class coverage (e.g., object or action class) of different attribute categories, different combinations of attributes can represent different classes - class detectors can then be instantiated on top of attribute detectors. This instantiation follows a specification of which attributes are associated with which classes. When given a specification for unseen categories, detectors can be built even though no visual samples of the unseen categories have been encountered.

Conventional use of attributes in computer vision is binary: an attribute is represented as either present, or absent. This limits the representational power of attribute representations. However, binary representations can be generalized to scalar representations where each attribute is associated with a scalar degree rather than a binary

category. This is both more flexible representationally, and allows the inclusion of attributes which do not so cleanly fall into a binary categorization. For example, while an attribute of “indoors / outdoors” often is clearly binary, an attribute of “moving fast / moving slow” has a more even distribution over gradations in visual input.

[49] presents one approach to generalizing binary attributes to scalar attributes. A challenge in generalizing from binary to scalar attributes is inconsistency in annotations, as different annotators may have different understandings of what different attributes’ degrees correspond to in the scalar representation. They resolve this challenge by requesting that annotators not assign scalar values to attributes, but rank images in terms of attribute degree. After images are ranked, scalar attribute values can be derived from their annotated relative attribute degrees.

For binary attributes, attribute detectors can be trained using conventional classifiers, but producing scalar attributes requires other methods. [49] train a ranking function over images for attributes over the seen set, and use that ranking function in the production of a scalar value.

Relative attribute representations allow for greater flexibility in class specifications. With consideration to the attribute of “moving fast / moving slow” one can specify that an unseen category “running” is faster than the seen category of “walking”, or that the unseen category of “standing” is slower than the seen category of “walking”. This is done without a need to define either a binary specification or intuit a scalar value with which to describe the unseen categories.

4.2. Shared Embedding Spaces

The generalization of binary attributes to scalar attributes increases their representational power. However, the increased representational power came at the cost of increased difficulty in annotating attributes and specifying classes. Relative Attributes [49] introduced one solution to these challenges.

Attributes can also be abstracted away. It need not be the case that visual representations move through humanly understandable attributes. Abstracting attributes away has a couple advantages:

- Avoid the imprecision introduced by passing through engineered, rather than learned, representations.
- Avoid the overhead and imprecision of annotating attributes.

The question is then how to construct a system such that training for classification of seen classes results in a classifier which works not only for seen classes, but for unseen classes as well, without leveraging an engineered mid-level representation which allows the specification of unseen classes in terms of that mid-level representation.

One approach that can be used is to define unseen classes in terms of their similarity relations to seen classes, where these relations are learned from text corpora. Exten-

sive work exists in Natural Language Processing (NLP) producing *semantic vector spaces* which represent similarity relations among words - word2vec [41] is one popular example. The intuition is that terms in these spaces are located in proximity to other terms with which they share semantic similarity. Trained semantic spaces are publicly available - these can be taken and used without a need to produce them from scratch. See Section 4.3 for details on construction of such spaces.

Terms with semantic similarity often share similarity in the visual space as well - e.g., “jogging” is both semantically and visually part way between “running” and “walking”. And so, it is often the case that if the visual similarity to known categories can be determined, then the semantic similarity relations can be established as well. Then, from these semantic relations we can infer semantic categories of visual inputs.

To illustrate: Consider a set of seen classes including “running” and “walking”, and a set of unseen classes including “jogging”, a semantic vector space capturing semantic proximity between these categories, and a computer vision architecture, such as a CNN, which produces visual feature representations of input. Then, consider an input sample of unknown class. Say that the visual representation of this input as produced by the CNN is part way between the visual representation for “running” and for “walking”. We then go to the semantic language space and see that the label that is located part way between “running” and “walking” is jogging and assign this label to the input.

Most often, comparing similarities between samples of unknown class and known classes is not done in the visual space. It is more common to map the input into the semantic space and perform comparisons to known classes there. This requires embedding one space into another.

The mechanism of embedding one space into another can be as simple as a linear transformation applied over one space, which is then trained over a similarity loss between two spaces. DeVise [18] is a good example of an architecture using this method. This involves two pretrained representations - visual features taken from, for instance, a CNN trained over classification, and word vectors in an embedding space, which may be produced through means discussed in Section 4.3.

One method of embedding visual features into the semantic space of the word vectors is then illustrated in Figure 0.8. A layer of nodes is appended to the top of the pre-trained CNN features, and then trained. The loss that this layer is trained over is the similarity (e.g., cosine similarity) between the output of this layer, and the vectors in the semantic embedding space corresponding to the labels of the visual input. This learns a simple linear mapping from visual feature vectors onto text derived semantic features.

More sophisticated mappings than linear translations are often used. Using multiple layers of neurons, in conjunction with non-linear activations, produce non-linear mappings (e.g., [29] use such a method). [31] use an autoencoder with semantic con-

straints to produce the embedding, and find that the constraint of visual reconstruction leads to better generalization to unseen classes in ZSL.

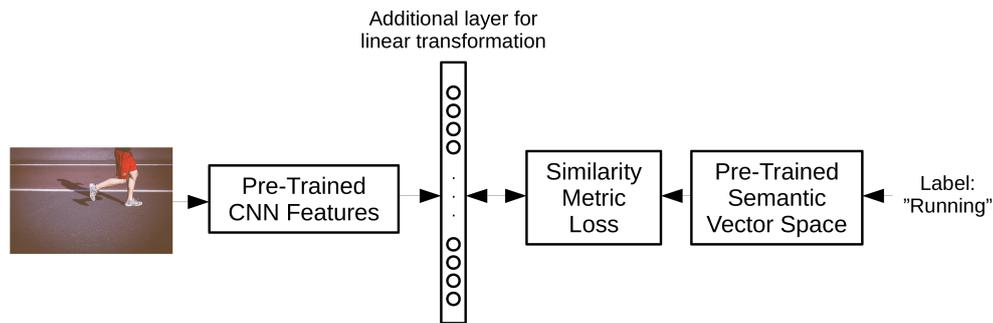


Figure 0.8 A simple approach to embedding visual representations into a semantic space, as is used by methods such as DeVise [18]. Two pre-trained models are used: 1) pretrained visual features, such as produced by a CNN trained over a visual task, and 2) a semantic vector space as produced through methods discussed in section 4.3. A simple linear transformation is produced through appending a layer of nodes on top of the visual features, and training them w.r.t. their similarity to the semantic word vector of the labels corresponding to visual input. Once this linear transformation is learned, it constitutes a mapping from the space of visual features into the semantic vector space.

Once both visual input and semantic representations are placed in the same space, then determining the class of novel visual input is as simple as representing visual input in that space and then finding the nearest semantic label in that same space.

4.3. Construction of Semantic Vector Spaces

4.3.1. word2vec

Here we wish to construct a vector space into which words are embedded in such a way that their semantics are captured spatially. The end results is a space where, for example, vectors for “jog”, “run”, and “walk” are located in proximity, with “jog” being placed in the middle of the three.

How can a word’s semantics be defined? One answer is through the interactions that that word has with other words in text corpora - semantics of words can be defined in relation to other words. How do we model words’ relations to other words? One simple approach is co-occurrence - if two words occur in proximity, they are taken to be related. The more frequently that they co-occur together, the more strongly they are related. This is the principle on which methods embedding words into vector spaces such as word2vec are based.

We start from the simplest vector representation for words - a 1-hot encoding,

where each vector position corresponds to one word in a vocabulary V . This is a high-dimensional, inefficient representation, where there are no meaningful spatial relations among words. We can embed words into a lower dimensional space with said desirable properties through the solving of one of two related tasks:

1. Predicting a target word based on context
2. Predicting context based on a target word

Here “context” C is defined as the set of terms “nearby” the target term. These are defined as the other terms present in an n -gram associated with the target term, without consideration for word order.

Each of these tasks can be solved using simple architectures, and the solving of these tasks produce in the process lower dimensional representations which can then be taken and used for other tasks.

A method - *Continuous Bag Of Words* (CBOW) - for solving Task 1 [41] is shown in Figure 0.9(a). Input consists of multiple words of context, each represented as a 1-hot vector of dimension $|V|$. These vectors are summed to produce a vector of size $|V|$. This is fed through a single layer of N neurons, where N is the dimension of the embedding space. On top of this we have one more layer, of size $|V|$, whose job it is to predict, in 1-hot representation, the word associated with the context consisting of the terms fed as input to the first layer.

A method - *Skip Gram* - for solving Task 2 [41] is shown in Figure 0.9(b). Input consists of a single term, represented as a 1-hot vector of dimension $|V|$. This is fed into a single layer of dimension N , where N is the dimension of the embedding space. On top of this layer we have an output layer, consisting of $|C|$ sets of $|V|$ nodes, each associated with one term of context C in the n -gram associated with the input term.

Both architectures are trained through serially feeding in n -grams extracted from large text corpora, and applying a soft-max loss to the final layer to enforce alignment with expected terms.

After these architectures are trained, the hidden layer then constitutes a mapping of V from a 1-hot representation of size $|V|$ to an embedded representation of size N where terms are spatially located in proximity to terms with similar semantics.

4.4. Shared Embedding Spaces and Graphical Models

ZSL of action can benefit from additional structure beyond what simple mapping into a shared embedding space can provide. Additional structure can be represented in the form of graphs. Multiple works [19, 29, 72, 20] employ graphs for the purpose of recognition of actions, processing these through use of a Graph Convolutional Network (GCN) to produce vector representations of action categories suitable for ZSL.

[19] evaluate 3 different graphical representations, the last of which is applicable to Few-Shot Learning rather than ZSL due to its use of visual features from a small num-

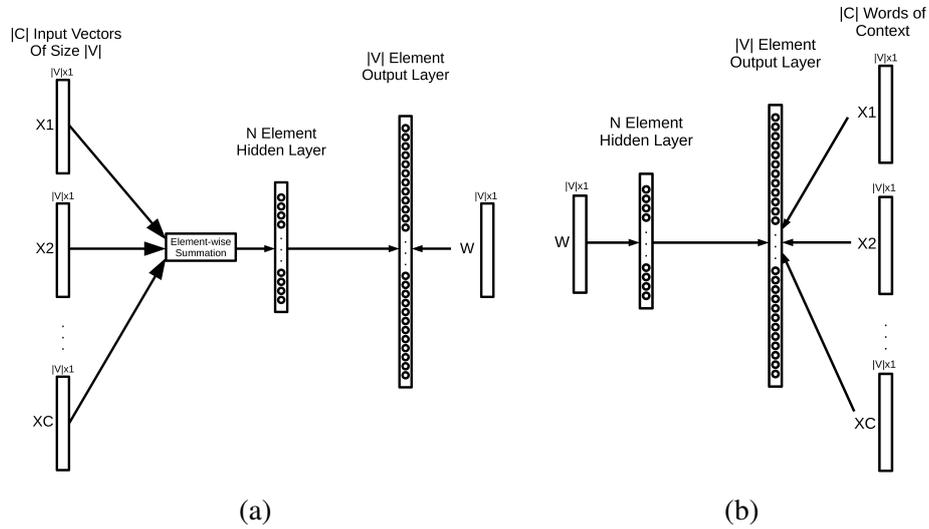


Figure 0.9 (a) The Continuous Bag Of Words method to solving the task of predicting a target word based on that words context [41]. $|C|$ input vectors of context are fed through two layers, the first a hidden layer of size N , the second an output layer of size V . Loss is computed w.r.t. the word W . After training the output layer can be discarded, and the hidden layer used as a translation from 1-hot word encodings into an embedded space of size N . (b) Skip Gram method to solving the task of predicting the context of a term [41]. Word W is fed through two network layers, the first a hidden layer of size N , the second an output layer of size V . Loss is computed w.r.t. $|C|$ words of context. After training, the output layer can be discarded, and the hidden layer used as a translation from 1-hot word encodings into an embedded space of size N .

ber of samples. [29] construct a graph based on Subject Verb Object triplets derived from knowledge corpora. All of these are processed through a GCN.

Throughout [19] sentence2vec [48] is used rather than word2vec, as authors find that action categories are better represented through phrases than through single words which can be confounded by multiple meanings. The first graph is composed of nodes taking values of sentence2vec embeddings of action category phrases. These nodes are linked together based on cosine similarity between vector representations - the top N closest neighbors per node are given edges. The second graph associates verbs and nouns - derived from Part-Of-Speech tagging of phrases describing actions - with each action class, incorporating nouns as a strong connection between seen and unseen categories. The third graph incorporates visual representations derived from a small number of samples of the unseen categories - this graph is thus applicable to Few-Shot Learning, rather than ZSL. The reason for the incorporation of visual representations is that categories which are similar in the semantic space may nevertheless have distinct visual appearances - authors give the example of “pommel horse” and “horse walking”,

which have similar word embedding representations but dissimilar visual appearances.

[29] construct a graph consisting of three types of nodes: noun nodes, object nodes, and action nodes. Action nodes are linked to the verbs and nouns which the associated action involves.

Values of graph nodes are generally initialized with values derived from semantic vector spaces - this is important as it establishes the initial relations which the GCN iterates over. This iteration incorporates relations defined by the edges in the graph, and allows information to transfer from node-to-node along edges. E.g., in [29] action nodes, which are initially set to zero vectors, acquire a representation determined by the vectors of neighboring noun and verb nodes, which have been initialized with semantic vectors.

Like previously detailed methods, these methods learn a mapping from visual features (taken from a CNN pre-trained on a separate task) into a shared embedding space - though here that space is shared with representations produced by the GCN. In [29] that mapping is produced through two layers of neurons with sigmoid activations, resulting in a non-linear mapping from the visual features into the shared semantic space. Similar to previous work, these layers are trained through applying a loss measuring the similarity between the visual features after embedding, and the vectors associated with the labels of the input as produced by the GCN.

The mapping from pre-trained visual features onto the GCN produced action vectors can then be trained over the seen training set. To obtain action predictions when applying the network to novel actions during testing, nearest neighbor can be performed between the visual features and the action vectors.

5. Future Directions

Here we discuss implications and future work implied by the action based framework outlined in the rest of this chapter. Action has implications for *tasks and datasets* - it enables conceptual modeling conducive to generalization and longer term temporal prediction. Action benefits from modeling of *concepts* beyond those from conventional CV. As CV progresses the scalability of fully supervised methods becomes an increasing issue, and action based methods help mitigate this issue while benefiting from *semi- and un-supervised* paradigms. Finally action helps enable an integration of *cognition* and symbolic modeling into *perception*, including across multiple perceptual modalities.

Tasks and datasets: Activities span long time spans. Thus when seeking solutions to visual activity understanding, we face problems much more challenging than those we encounter in current action recognition tasks. Objects, actions, affordances, and other scene constituents relate to each other semantically over multiple time scales, and we need to find ways to model these relations. We think that this capability is

not well demonstrated with the task of recognizing actions. Instead, we should pick tasks that demonstrate generalization and a conceptual understanding of action (as opposed to a purely appearance based understanding), such as zero-shot learning and the prediction of future actions, and translation from one viewpoint to another (e.g., from first person to third person). Today’s CV research is largely driven by the collection of new datasets and definition of new challenges - existing datasets do not sufficiently cover long term and conceptual modeling of activities. The datasets ideally should have recordings from multiple views, because this opens possibilities for interesting research, for example, to solve the challenge of transferring knowledge between the first person and the third person view. Lastly, most datasets feature indoor scenes. It will be interesting to collect outdoor scenes and analyze them, as discussed above, by looking at the relations between affordances, interactions and long-term relations. We could also attempt an action-based analysis on data from the autonomous driving domain.

Concepts for long-term activity understanding: Activity understanding requires action-based concepts at multiple time scales. We discussed in this chapter such quantities at the single image and short-term time scales, including affordances, hand grasps, geometric relations, and we emphasized the use of geometric reconstruction processes because of their robustness (see discussion in Section 3.4). The next step will be to add further constraints and include robust constraints for modeling temporal-relations at longer scales. We could make use of ontologies in categorizing objects and actions. On verbs we can impose classifications based on action effects [73], ergonomic principles, or force and location related constraints. Longer-term relations include causal relations, and constraints on possible and impossible action sequences. We can also model physics constraints, and use physics engines - but in order to integrate these into deep architectures we need to include these constraints into vector spaces that relate perception to cognition (Section 4).

Reducing supervision: Early approaches on integrating language with vision (see Sections 4.1 and 4.4) have relied heavily on supervision. For example, visual attribute recognition or object recognition has been implemented via supervised learning. Graph-based models using shared embeddings for ZSL need to incorporate in advance the categories of the “unseen” set which they may encounter during test time. Naturally, evolving approaches will find their way into the action-based framework, including unsupervised and self-supervised learning, transfer learning, meta-learning, and eventually never-ending learning [43]. For example, in constructing visual ontologies, we don’t want to rely fully on supervised visual learning of meta-verb classes. One solution to such modeling is dictionary learning [83], but so far these approaches have been limited to simple actions. We will need methods that scale to more complex human manipulation actions. Generative Adversarial Networks (GANs) and Variational Autoencoders (VAEs) have been shown successful in modeling fine-grained ac-

tion images and videos. We could, for example, use VAEs to learn the underlying distribution of data into a discretized latent space that encodes the meta-categories.

Involving the whole brain: Humans’ understanding of the world is grounded in our motoric cognition and all our senses. Similarly, our models should include other sensory modalities like auditory, tactile, or proprioception signals in addition to modeling of vision and cognition representations. Different modalities provide complementary information, and because of this allow for different ways of organizing concepts. In addition to questions of learning with different modalities, we also need methods for accessing stored concepts when given perception from any modality. Studying the integration of diverse modalities will also benefit from studying memory. A framework known as Vector Symbolic Architectures (VSA) [51, 13], which includes Hyperdimensional (HD) Computing (methods making use of very high dimensional vector spaces) [28], has been proposed as a theoretical model for artificial intelligence. HD Computing combines advantages of neural-based AI approaches with systematic compositionality and rule-like behavior from classical symbolic AI [36]. In this framework concepts are encoded into vector spaces, and algebraic operations are defined on these vector spaces. These operations include the addition of related concepts and the binding of vectors of different origin - for example this could be sound and vision, or vision and motor [44]. These operations maintain the separability of one modality from the other. We may build on this framework and integrate it with neural network approaches. The goal will be to retain an explicit memory encoding different perception modalities, while maintaining the capability of recalling information from any modality.

6. Conclusions

The purpose of Computer Vision is to produce interpretations which are of use to humans. Action is central to our understanding of the world, yet is underutilized in contemporary CV. This chapter covered scene and activity understanding that has the concept of action and interaction at its center. We covered action based approaches to scene understanding involving modeling at multiple temporal scales, starting from object interpretation in terms of affordances at the instantaneous level, up to basic actions, then up to full activities at the longer temporal scale. We described the well developed area of affordance learning, and described works on activity understanding combining cognitive and linguistic approaches with humanly interpretable modules essential in characterizing activities and segmenting video temporally. We discussed methods for the integration of visual representations with knowledge, both engineered and derived from text corpora. We covered the integration of vision, centered on action, with graphical constraints, and discussed future directions including creating new challenges and datasets, adding concepts for encoding long-term relations, adapting semi- and un-supervised learning approaches, and incorporating memory as a central component to the action-based framework.

ACKNOWLEDGEMENTS

The support of the National Science Foundation under grants BCS 1824198 and OISE 2020624 are gratefully acknowledged.

REFERENCES

1. Somak Aditya, Yezhou Yang, Chitta Baral, Yiannis Aloimonos, and Cornelia Fermüller. Image understanding using vision and reasoning through scene description graph. *Computer Vision and Image Understanding*, 173:33–45, 2018.
2. Yiannis Aloimonos and Cornelia Fermüller. The Cognitive Dialogue: A new model for vision implementing common sense reasoning. *Image and Vision Computing*, 34:42–44, 2015.
3. Caterina Ansuini, Andrea Cavallo, Cesare Bertone, and Cristina Becchio. Intentions in the brain: the unveiling of Mister Hyde. *The Neuroscientist*, 21(2):126–135, 2015.
4. Ruzena Bajcsy. Active perception. *Proceedings of the IEEE*, 76(8):966–1005, 1988.
5. Lawrence W Barsalou. Grounded cognition. *Annu. Rev. Psychol.*, 59:617–645, 2008.
6. Liefeng Bo, Xiaofeng Ren, and Dieter Fox. Unsupervised feature learning for rgb-d based object recognition. In *Experimental Robotics*, pages 387–402. Springer, 2013.
7. Mark R Cutkosky. On grasp choice, grasp models, and the design of hands for manufacturing tasks. *IEEE Transactions on Robotics and Automation*, 5(3):269–279, 1989.
8. Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, et al. Scaling egocentric vision: The EPIC-Kitchens dataset. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 720–736, 2018.
9. Pradipto Das, Chenliang Xu, Richard F Doell, and Jason J Corso. A thousand frames in just a few words: Lingual description of videos through latent topics and sparse object stitching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2634–2641, 2013.
10. Eadom Dessalene, Chinmaya Devaraj, Michael Maynard, Cornelia Fermüller, and Yiannis Aloimonos. Forecasting Action through Contact Representations from First Person Video. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
11. V Dutta and T Zielinska. Action prediction based on physically grounded object affordances in human-object interactions. *Proceedings of the 11th International Workshop on Robot Motion and Control*, 2017.
12. Aleksandrs Eciņs, Cornelia Fermüller, and Yiannis Aloimonos. Cluttered scene segmentation using the symmetry constraint. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 2271–2278, 2016.
13. Chris Eliasmith. *How to Build a Brain: A Neural Architecture for Biological Cognition*. Oxford University Press, 2013.
14. M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The Pascal Visual Object Classes (VOC) Challenge. *International Journal of Computer Vision*, 88(2):303–338, June 2010.
15. Cornelia Fermüller and Yiannis Aloimonos. Vision and Action. *Image and Vision Computing*, 13(10):725–744, 1995.
16. Cornelia Fermüller, Fang Wang, Yezhou Yang, Konstantinos Zampogiannis, Yi Zhang, Francisco Barranco, and Michael Pfeiffer. Prediction of manipulation actions. *International Journal of Computer Vision*, 126(2):358–374, 2018.
17. Paul Fitzpatrick, Giorgio Metta, Lorenzo Natale, Sajit Rao, and Giulio Sandini. Learning about objects through action-initial steps towards artificial cognition. In *IEEE International Conference on Robotics and Automation*, volume 3, pages 3140–3145, 2003.
18. Andrea Frome, Greg Corrado, Jonathon Shlens, Samy Bengio, Jeffrey Dean, Marc’Aurelio Ranzato, and Tomas Mikolov. Devise: A deep visual-semantic embedding model. In *Advances in Neural Information Processing Systems 26*, 2013.
19. Pallabi Ghosh, Nirat Saini, Larry S Davis, and Abhinav Shrivastava. All about knowledge graphs for actions. *arXiv preprint arXiv:2008.12432*, 2020.
20. Pallabi Ghosh, Yi Yao, Larry Davis, and Ajay Divakaran. Stacked spatio-temporal graph convolutional networks for action segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 576–585, 2020.
21. James J Gibson. The theory of affordances. In John Bransford Robert E Shaw, editor, *Perceiving, Acting, and Knowing: Toward an Ecological Psychology*, pages pp.67–82. Hillsdale, N.J. : Lawrence Erlbaum Associates, 1977.
22. Helmut Grabner, Jürgen Gall, and Luc Van Gool. What makes a chair a chair? In *IEEE Conference*

- on *Computer Vision and Pattern Recognition*, pages 1529–1536, 2011.
23. Anupam Guha, Yezhou Yang, Cornelia Fermüller, and Yiannis Aloimonos. Minimalist plans for interpreting manipulation actions. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5908–5914, 2013.
 24. Abhinav Gupta, Scott Satkin, Alexei A Efros, and Martial Hebert. From 3d scene geometry to human workspace. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1961–1968, 2011.
 25. Mohammed Hassanin, Salman Khan, and Murat Tahtali. Visual affordance and function understanding: A survey. *arXiv preprint arXiv:1807.06775*, 2018.
 26. Varsha Hedau, Derek Hoiem, and David Forsyth. Recovering the spatial layout of cluttered rooms. In *IEEE International Conference on Computer Vision*, pages 1849–1856, 2009.
 27. Tucker Hermans, James M Rehg, and Aaron Bobick. Affordance prediction via learned object attributes. In *IEEE International Conference on Robotics and Automation (ICRA): Workshop on Semantic Perception, Mapping, and Exploration*, 2011.
 28. Pentti Kanerva. Hyperdimensional computing: An introduction to computing in distributed representation with high-dimensional random vectors. *Cognitive Computation*, 1:139–159, 2009.
 29. Keizo Kato, Yin Li, and Abhinav Gupta. Compositional learning for human object interaction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 234–251, 2018.
 30. Hedvig Kjellström, Javier Romero, and Danica Kragić. Visual object-action recognition: Inferring object affordances from human demonstration. *Computer Vision and Image Understanding*, 115(1):81–90, 2011.
 31. Elyor Kodirov, Tao Xiang, and Shaogang Gong. Semantic autoencoder for zero-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3174–3183, 2017.
 32. Hema S Koppula and Ashutosh Saxena. Physically grounded spatio-temporal object affordances. In *European Conference on Computer Vision*, pages 831–847. Springer, 2014.
 33. Hema S Koppula and Ashutosh Saxena. Anticipating human activities using object affordances for reactive robotic response. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(1):14–29, 2015.
 34. Hema Swetha Koppula, Rudhir Gupta, and Ashutosh Saxena. Learning human activities and object affordances from rgb-d videos. *The International Journal of Robotics Research*, 32(8):951–970, 2013.
 35. David C Lee, Abhinav Gupta, Martial Hebert, and Takeo Kanade. Estimating spatial layout of rooms using volumetric reasoning about objects and surfaces. In *Advances in Neural Information Processing Systems*, pages 1288–1296, 2010.
 36. S.D. Levy and R. Gayler. Vector symbolic architectures: A new building material for artificial general intelligence. In *Proceedings of the First Conference on Artificial General Intelligence (AGI-08)*. IOS Press, 2008.
 37. Fayao Liu, Chunhua Shen, and Guosheng Lin. Deep convolutional neural fields for depth estimation from a single image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5162–5170, 2015.
 38. Devraj Mandal, Sanath Narayan, Sai Kumar Dwivedi, Vikram Gupta, Shuaib Ahmed, Fahad Shahbaz Khan, and Ling Shao. Out-of-distribution detection for generalized zero-shot action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9985–9993, 2019.
 39. A. Martin. The representation of object concepts in the brain. *Annual Review of Psychology*, 58:25–45, 2007.
 40. Giorgio Metta, Giulio Sandini, David Vernon, Lorenzo Natale, and Francesco Nori. The iCub humanoid robot: an open platform for research in embodied cognition. In *Proceedings of the 8th Workshop on Performance Metrics for Intelligent Systems*, pages 50–56, 2008.
 41. Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013.
 42. George A Miller, Richard Beckwith, Christiane Fellbaum, Derek Gross, and Katherine J Miller. In-

- roduction to wordnet: An on-line lexical database. *International Journal of Lexicography*, 3(4):235–244, 1990.
43. T. Mitchell, W. Cohen, E. Hruschka, P. Talukdar, J. Betteridge, A. Carlson, B. Dalvi, M. Gardner, B. Kisiel, J. Krishnamurthy, N. Lao, K. Mazaitis, T. Mohamed, N. Nakashole, E. Platanios, A. Ritter, M. Samadi, B. Settles, R. Wang, D. Wijaya, A. Gupta, X. Chen, A. Saparov, M. Greaves, and J. Welling. Never-ending learning. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI-15)*, 2015.
 44. A Mitrokhin, P Sutor, C Fermüller, and Y Aloimonos. Learning sensorimotor control with neuro-morphic sensors: Toward hyperdimensional active perception. *Science Robotics*, 4(30):eaaw6736, 2019.
 45. Luis Montesano, Manuel Lopes, Alexandre Bernardino, and José Santos-Victor. Learning object affordances: from sensory–motor coordination to imitation. *IEEE Transactions on Robotics*, 24(1):15–26, 2008.
 46. Austin Myers, Ching L Teo, Cornelia Fermüller, and Yiannis Aloimonos. Affordance detection of tool parts from geometric features. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1374–1381, 2015.
 47. Anh Nguyen, Dimitrios Kanoulas, Darwin G Caldwell, and Nikos G Tsagarakis. Object-based affordances detection with convolutional neural networks and dense conditional random fields. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5908–5915, 2017.
 48. Matteo Pagliardini, Prakhar Gupta, and Martin Jaggi. Unsupervised learning of sentence embeddings using compositional n-gram features. *arXiv preprint arXiv:1703.02507*, 2017.
 49. Devi Parikh and Kristen Grauman. Relative attributes. In *International Conference on Computer Vision*, pages 503–510, 2011.
 50. Katerina Pastra and Yiannis Aloimonos. The minimalist grammar of action. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367(1585):103–117, 2012.
 51. Tony A Plate. Holographic reduced representations. *IEEE Transactions on Neural networks*, 6(3):623–641, 1995.
 52. Siyuan Qi, Siyuan Huang, Ping Wei, and Song-Chun Zhu. Predicting human activities using stochastic grammar. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1164–1172, 2017.
 53. Siyuan Qi, Wenguan Wang, Baoxiong Jia, Jianbing Shen, and Song-Chun Zhu. Learning human-object interactions by graph parsing neural networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 401–417, 2018.
 54. Anirban Roy and Sinisa Todorovic. A multi-scale cnn for affordance segmentation in rgb images. In *European Conference on Computer Vision*, pages 186–201. Springer, 2016.
 55. Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.
 56. Karin Kipper Schuler. *VerbNet: A broad-coverage, comprehensive verb lexicon*. PhD thesis, Computer and Information Science Department, University of Pennsylvania, Philadelphia, PA, 2005.
 57. Gunnar A Sigurdsson, Olga Russakovsky, and Abhinav Gupta. What actions are needed for understanding human actions in videos? In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2137–2146, 2017.
 58. Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from rgb-d images. In *European Conference on Computer Vision*, pages 746–760. Springer, 2012.
 59. Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
 60. Abhilash Srikantha and Jürgen Gall. Weakly supervised learning of affordances. *arXiv preprint arXiv:1605.02964*, 2016.
 61. Louise Stark and Kevin Bowyer. Achieving generalized object recognition through reasoning about association of function to structure. *IEEE Transactions on Pattern Analysis and Machine Intelligence*,

- 13(10):1097–1104, 1991.
62. Douglas Summers-Stay, Ching L Teo, Yezhou Yang, Cornelia Fermüller, and Yiannis Aloimonos. Using a minimal action grammar for activity understanding in the real world. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4104–4111, 2012.
 63. Ching L Teo, Cornelia Fermüller, and Yiannis Aloimonos. Detection and segmentation of 2d curved reflection symmetric structures. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1644–1652, 2015.
 64. Son D Tran and Larry S Davis. Event modeling and recognition using markov logic networks. In *European Conference on Computer Vision*, pages 610–623. Springer, 2008.
 65. Emre Ugur, Erhan Oztop, and Erol Sahin. Goal emulation and planning in perceptual space using learned affordances. *Robotics and Autonomous Systems*, 59(7-8):580–595, 2011.
 66. Emre Ugur and Justus Piater. Emergent structuring of interdependent affordance learning tasks using intrinsic motivation and empirical feature selection. *IEEE Transactions on Cognitive and Developmental Systems*, 9(4):328–340, 2016.
 67. Francisco J Varela, Eleanor Rosch, and Evan Thompson. *The Embodied Mind: Cognitive Science and Human Experience*. MIT Press, 1993.
 68. Florentin Wörgötter, Eren Erdal Aksoy, Norbert Krüger, Justus Piater, Ales Ude, and Miniija Tamosiunaite. A simple ontology of manipulation actions based on hand-object relations. *IEEE Transactions on Autonomous Mental Development*, 5(2):117–134, 2013.
 69. Florentin Wörgötter, F Ziaetabar, S Pfeiffer, O Kaya, T Kulvicius, and M Tamosiunaite. Humans predict action using grammar-like structures. *Scientific Reports*, 10(1):1–11, 2020.
 70. Yongqin Xian, Tobias Lorenz, Bernt Schiele, and Zeynep Akata. Feature generating networks for zero-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5542–5551, 2018.
 71. Jianxiong Xiao, James Hays, Krista A. Ehinger, A. Oliva, and A. Torralba. Sun database: Large-scale scene recognition from abbey to zoo. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 3485–3492, 2010.
 72. Sijie Yan, Yuanjun Xiong, and Dahua Lin. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
 73. Yezhou Yang, Cornelia Fermüller, and Yiannis Aloimonos. Detection of manipulation action consequences (MAC). In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2563–2570, 2013.
 74. Yezhou Yang, Cornelia Fermüller, Yi Li, and Yiannis Aloimonos. Grasp type revisited: A modern perspective on a classical feature for vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 400–408, 2015.
 75. Yezhou Yang, Anupam Guha, Cornelia Fermüller, and Yiannis Aloimonos. A cognitive system for understanding human manipulation actions. *Advances in Cognitive Systems*, 3:67–86, 2014.
 76. Yezhou Yang, Yi Li, Cornelia Fermüller, and Yiannis Aloimonos. Robot learning manipulation action plans by “watching” unconstrained videos from the world wide web. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29, 2015.
 77. Chengxi Ye, Yezhou Yang, Ren Mao, Cornelia Fermüller, and Yiannis Aloimonos. What can I do around here? Deep functional scene understanding for cognitive robots. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 4604–4611, 2017.
 78. Xiaodong Yu, Cornelia Fermüller, Ching Lik Teo, Yezhou Yang, and Yiannis Aloimonos. Active scene recognition with vision and language. In *2011 International Conference on Computer Vision*, pages 810–817, 2011.
 79. Konstantinos Zampogiannis, Cornelia Fermüller, and Yiannis Aloimonos. Cilantro: A lean, versatile, and efficient library for point cloud data processing. In *Proceedings of the 26th ACM International Conference on Multimedia*, pages 1364–1367, 2018.
 80. Konstantinos Zampogiannis, Cornelia Fermüller, and Yiannis Aloimonos. Topology-aware non-rigid point cloud registration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.
 81. Konstantinos Zampogiannis, Yezhou Yang, Cornelia Fermüller, and Yiannis Aloimonos. Learning

- the spatial semantics of manipulation actions through preposition grounding. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 1389–1396, 2015.
82. Li Zhang, Tao Xiang, and Shaogang Gong. Learning a deep embedding model for zero-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2021–2030, 2017.
 83. J. Zheng, Z. Jiang, and R. Chellappa. Cross-view action recognition via transferable dictionary learning. *IEEE Transactions on Image Processing*, 25(6):2542–2556, 2016.

Author biographies

Cornelia Fermüller is a Research Scientist at the University of Maryland Institute for Advanced Computer Studies. She holds a Ph.D. from the Vienna University of Technology, Austria and an M.S. from the Graz University of Technology, both in Applied Mathematics. Her research interest has been to understand principles of active vision systems and develop biological-inspired methods, especially in the area of motion. Her recent work has focused on human action interpretation and the development of 3D motion algorithms for extreme conditions using event-based sensors.

Michael Maynard is a PhD candidate in the department of Computer Science at the University of Maryland College Park, advised by Yiannis Aloimonos and Cornelia Fermüller. His background encompasses symbolic Artificial Intelligence, including cognitive architectures, Computer Vision, including action understanding, and methods integrating AI and CV.