

Integrating Object Affordances with Artificial Visual Attention

Jan Tünnermann, Christian Born & Bärbel Mertsching

GET Lab

University of Paderborn

Pohlweg 47–49, 33098 Paderborn, Germany

{tuennermann, born, mertsching}@get.uni-paderborn.de

Abstract. Affordances, as for example grasping possibilities, are known to play a role in the guidance of human attention but have not been considered in artificial attention systems so far. Extending our earlier work, we investigate the combination of affordance estimation and visual saliency in an artificial visual attention model. Different models based on saliency, affordance estimation, or their combination are suggested and evaluated via their predictions for a change detection task with human observers. Furthermore, we discuss the potential of Growing Neural Gases as a framework for consistently integrating bottom-up saliency, affordance-based and top-down attention mechanisms.

Keywords: Attention, Saliency, Affordance

1 Introduction

With the transition of robots from specialized automata performing predefined tasks to general autonomous agents, the requirements to perceive, reason about, and interact with their environment have drastically increased. A recent development in robotics is to model aspects of environmental psychology, which deal with the interaction between humans and their surroundings. A popular concept from this field, the *affordance of objects*, was introduced by J. J. Gibson in 1977 [1]. In this holistic view, objects possess certain affordances, i.e., objects or their parts can afford certain actions. A common example is a mug, whose handle affords grasping.

This idea has been transferred to technical systems to enhance their performance preparing and executing actions that correspond to such affordances (see e.g., [2–4]). One result of such studies—as well as of earlier general object recognition studies (e.g., [5, 6])—is that the affordance concept in technical systems not only supports direct grasping actions, it can also benefit object recognition and semantic scene perception. The general idea is that in many cases an object is better defined by the actions which the object supports, than with its visual attributes. Coming back to the example of a mug, even though colors and shapes may differ widely, mugs in general afford grasping (possibly by some kind of handle), containing liquid and drinking [7]. Therefore, recent research

integrates affordance estimation with object recognition (e.g., [8, 7]) and the semantic interpretation of scenes and objects (e.g., [9, 10]).

Visual attention is another concept which is being transferred from cognitive psychology to technical systems. The main idea is to filter relevant from irrelevant information very early in processing, and distribute processing resources accordingly. Attention can be guided bottom-up by saliency—local contrasts with respect to features such as intensity, color, or local orientation [11]—or in a top-down manner by incorporating knowledge, such as task demands [12] or the “gist of the scene” [13].

Findings from psychology suggest that affordances influence visual attention. This has been shown in reductions of reaction times when affordances were used to guide attention towards target locations [14, 15] and effects on event related signals in electrophysiological and brain imaging research [16].

The influence of affordance on visual attention can be regarded as a way to bias the distribution of processing resources towards objects that afford actions. These constitute potential targets of actions or reasoning, even though a specific behavior towards them may not be planned at this stage.




Fig. 1. **a:** A test scene. Inset: exemplary grasps fitted on a sparse 3D scene representation. **b:** Grasp points projected into 2D. White patterns represent grasps towards reachable locations, whereas red patterns indicate locations out of reach. “A” affordance-based, “S” saliency-based, and “A+S” combined selection. **c–e:** Underlying affordance (c), saliency (d), and combined (e) maps.

Recently, we investigated whether an artificial visual attention model can better predict human performance in an attention guided task when the prediction is based on affordance estimation instead of saliency estimation [17]. We found that this indeed is the case. Utilizing a change detection task, we compared a popular saliency model [11] with an attention model based on grasping hypotheses (estimated from a sparse scene representation [18], see figure 1a). The phenomenon of change blindness due to short scene interruptions renders the detection of changes in objects difficult, but the performance is increased when attention is allocated to the changing part of the scene [19]. For the evaluation of psychologically inspired computer vision systems, the change blindness

paradigm has the great advantage that images with natural scenes can be used, whereas many other psychophysical tasks require the use of highly artificial synthetic stimuli. We found that human observers performed better in reporting the changes that were made to objects selected by the affordance-based model than when those selected by the saliency model were changed.

In the present study, we extend our previous work [17] to investigate a model based on combined affordance and saliency, as opposed to the models that used either individual component. If this enhances the prediction of change detection, it may better reflect the allocation of attention. This may be the case, for example, when there are several objects of high affordance and their individual saliencies tip the scale for selection priority.

The present paper is an update of the report we presented at the First Workshop on Affordances: Affordances in Vision for Cognitive Robotics [20] and is structured as follows: In section 2, we report a new change detection experiment that uses the stimulus material already used in [17] but a different task. This is to test this task and get a preliminary insight how saliency and affordance are related in this stimulus material. In section 3, we describe a prototypical model of combined saliency and affordance which is used to create new test images to evaluate the combination of saliency and affordance (in section 3.2) in the change-blindness task. Finally, in section 4 we discuss Growing Neural Gases as a framework for an artificial attention model that we believe has the potential to integrate bottom-up saliency, affordance-based attention and top-down attention in a consistent architecture and improve on several disadvantages of current attention systems.

The implications are discussed in section 5.

2 Saliency and Affordance in Change Detection

2.1 Experiment I - Motivation

In our previous study [17], we employed a “single-shot” change detection task with natural images (see e.g., [21]) to measure the participants’ distribution of attention towards salient or affording objects. The single-shot paradigm shows a single change from the original to the altered image which are shown only briefly (usually between 100 and 500 ms) and a blank screen is shown between the two images. The presentation usually lasts for less than a second and participants respond afterwards, when the image is already gone. Thus, there is only a single binary hit-or-miss measurement per change. Furthermore, the same images cannot be repeated and therefore the amount of trials is limited to the number of available images. Their creation is quite an effort, due to editing in the changes (object removals in our case). Because of the limited number of trials and the binomially distributed response, a large number of subjects is required (40 - 80) to obtain reliable results. In our previous study there were two possible changes for every image—saliency or affordance—and thus two subjects were required to obtain an affordance and saliency measurement for every image. In the experi-

ment we report in section 3.2, three changes were possible. This further increases the number of subjects required.

Hence, in this experiment, we test the so called “flicker paradigm” (see e.g., [19]): the presentation is similar as described above, but it is repeated until the change is reported. Therefore, a more informative measure, namely the time it takes the subject to detect the change, can be obtained. This not only reduces the number of participants required, but may also allow to relate the degree of affordance and saliency to the response time. Therefore, the objective of this first experiment is, using the stimulus material from [17], to investigate whether the effect that affordances are more important than saliency in change detection can be replicated using the flicker paradigm. Furthermore, a first insight in the influence of the saliency and affordance values on the response time is provided.

Participants: Twelve volunteers (average age of 26.82, $SD = 3.6$) participated in this experiment. All had normal or corrected-to-normal vision and not seen the images before.

Stimuli: The stimulus material reported in [17] was used. This consisted of 28 natural scenes, mostly pictures of office environments that contained a number of objects in the reachable action space and some in background areas which would not be reachable by the observer of the scene. For every image, two changed versions were created by locally altering the image (locally blending in an identical image in which the object had been removed at that location). In one altered image of the same scene, an object selected by the saliency model by Itti et al. [11] had been removed, in the other altered image, one object selected by an affordance-based prediction (density of grasping possibilities per image area; this corresponds to the affordance stream described in see section 3.1) had been removed. Refer to [17], for more details regarding the stimulus generation and the actual pictures.

Design and Procedure: The experiment was conducted on a 12.1” touch-screen laptop¹. The participants were instructed to normally sit in front of the laptop and to adjust the distance and display angle for optimal viewing and touching.

Participants were presented with every original image paired with one of the possible changes. Thus, a single participant saw 14 original–affordance and 14 original–saliency pairs in random order. The number for which each changed image appeared with the affordance or saliency change was balanced over all subjects.

In contrast to the task used in [17], where participants saw the change only once (from original to the changed version), here the images cycled back and

¹ Note that the change blindness effect is very robust and does not require highly accurate timing that can only be established with CRT monitors or specialized equipment, which is the case for many other psychophysical paradigms.

forth between the original and the changed image until the subject discovered the change and responded by touching the screen at the location of the change. If no response was made within one minute, the current trial was aborted and the next trial started. The timing of the image sequence was: “1000 ms initial blank”–“300 ms original image”–“300 ms blank”–“100 ms changed image”–“300 ms blank”. The sequence after the initial blank was repeated until the participant responded. The changed image was shown for a shorter time as we are interested in the distribution of attention in the original image, reducing the risk that highly salient (or affording) objects that become visible in the changed image substantially influence the response. For every trial, the response time was recorded.

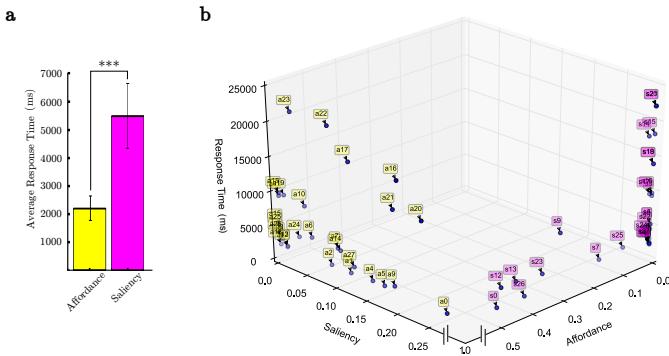


Fig. 2. a: Average response times for affordance- and saliency-based changes (error bars show the SEM; *** $p < 0.001$). **b:** Average response times (in milliseconds) of each change plotted into the affordance–saliency space with the saliency and affordance values at the changed locations in the respective maps. Points with pink labels (“sX”) are saliency-based changes, whereas yellow labels (“aX”) mark affordance-based changes in images X. Because the most salient and most affording objects had been selected for a change, they are set to maximum (1.0) in their respective dimension. The value for the other dimension was obtained by averaging the respective map under the removed object’s bounding rectangle.

Results and Discussion: Figure 2a shows the average response time to saliency- and affordance-based changes. Affordance based changes are reported significantly faster, $t(11) = -5.03, p < 0.001$, confirming our earlier results [17] from the single shot hit-or-miss task. This provides further evidence for the importance of affordances in the deployment of attention. The one minute time limit was reached only four times in the 336 changes presented over all subjects.

Furthermore, based on visual inspection, the distribution of responses times per image in the affordance–saliency space (see figure 2b), suggest that no long response times are found when a change has also high activity in the respective

other dimension. Especially for the saliency-based changes, long response times occur where the affordance is close to zero. Therefore, an attention model that combines saliency and affordance could show a better performance than models based on each individual component. Such a combined model is described in the following section and tested experimentally in section 3.2.

3 Combining Saliency and Affordance

3.1 A Combined Model of Saliency and Affordance

The proposed combined model of saliency and affordance is presented in figure 3. The processing begins with the acquisition of a stereo image pair. The saliency maps and the eventually combined attention map are based on the left image. Therefore, the left image is segmented into homogeneously colored regions ① (see [22] or [23] for segmentation methods for region-based attention on which this model is based). These regions can be considered proto-objects, which at pre-attentional stages are used for the feature and saliency computations and later integrate the different saliency dimensions and the affordance estimation.

In the saliency stream, the regions are used to generate feature magnitude maps for *color*, *orientation*, *eccentricity*, *symmetry* and *size* (s2). The feature *color* is obtained as the average color of all pixels of a region. *Orientation*, *eccentricity* and *symmetry* are calculated based on 2D central moments of the spatial distribution of a region’s pixels. *Size* is the number of pixels in a region. For details on the calculation of feature magnitudes, please refer to [24]. As a next step in this stream, saliency maps are calculated for each feature dimension individually (s3). This is done by applying a voting style procedure, where each region collects votes from its neighbors, regarding the dissimilarities in every feature dimension. Details on how the difference between two regions in a specific feature dimension is measured can be found in [24].

In the affordance stream, the left image is used to generate *2D Texlets* which are small local texture patches (a1). Using stereo disparities (a2), the *2D Texlets* are transferred into 3D space to form *3D Texlets* (a3). There, position-based *k*-means clustering is used to form small groups of neighboring *Texlets*. Planes are fitted through these *groups* to form so called *Surflings* (a4), which are then further grouped (when close to each other and similarly oriented) to generate *Surfaces* (a5).

Grasping hypotheses in 3D space are generated by fitting a simulated gripper (see figure 1b) to elements of the scene considering the surfaces generated in the process outlined above (a6). For details on the generation of this feature hierarchy and the grasping hypotheses, see [18]. The result of this process, which we make use of in this study, is the estimated contact points of the gripper on the surfaces. Note that in the present study the simulated gripper performs simple two-fingered grasping and does not include a bio-mechanic model of grasping. The latter can be used to filter out grasping hypotheses which may not be performable by the observer. The only criterion which is used to filter grasping

hypotheses with regard to performability is a distance limit of 70 cm, which roughly corresponds to human reaching distance.

The combined attention map is then obtained by integrating the individual saliency maps (a4) (according to the combination strategy described in [24], which assigns higher weights to channels with sharp saliency peaks) and combining this contribution from the saliency stream with the contribution from the affordance stream using the regions which have been obtained during the initial segmentation. While the saliency contribution is already in region-form, the grasping hypotheses (represented by the contact points) are projected from 3D into 2D. All points which fall into a certain region (a7)+2 are summed and normalized by the region size to yield the affordance value of a region. Because the contact points can often be found on the edges of objects (in their 2D projection), instead of considering a single point in this process, each back-projected point is expanded to 5×5 points in a square region surrounding the initial location, which contribute decreasingly weighted with regard to their distance from the original location. This can be seen in figure 1b.

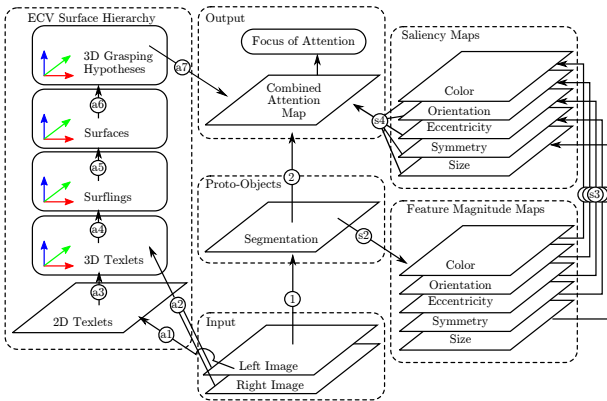


Fig. 3. Structure of the proposed model. The type and flow of data is described in the main text referring to this figure.

When the estimated affordance is combined with the saliency, this can be done in different ways. In principle, feature combination strategies as described in [11] can be applied. For instance, the map providing a small number of strong peaks could be preferred. Nevertheless, given that there is not yet any knowledge on how affordance and saliency are supposed to combine, we perform a naive combination. This is done by normalizing the values into the same dynamic range ($0 \dots 1$) and then calculating a weighted average of affordance and saliency for each region.

For the experiment reported in section 3.2, equal weights are used to obtain the combined prediction. Weighting saliency one and affordance zero yields the pure saliency-based prediction while the reverse results in the purely affordance-based prediction.

3.2 Experiment II - Evaluation

This second experiment is based on predictions from the newly introduced model that combines affordance and saliency estimation. It is intended to investigate whether predictions based on combined saliency and affordance better reflect human attention than the individual components.

Participants: Thirty volunteers (average age of 28.46, $SD = 5.86$) participated.

Stimuli, design and procedure: Stimulus material, experimental design and procedure mostly correspond to the description of the first experiment in section 2.1. The only differences were the use of a new image set (see figure 4a) with an additional third possible change based on the combined prediction with equally weighted affordance and saliency contributions. Furthermore, the saliency-based prediction was obtained with the region-based saliency model [24], which constitutes the saliency channel in the combined model, in contrast to the first experiment where the model by Itti and colleagues [11] was used.

Due to the fact that three predictions (affordance, saliency, combined) are required for each image, and the images focus mainly on the action space where saliency and affordance are both expected to be relatively high, sometimes the same object was selected by two or all three predictions. In such a case, the scene was slightly rearranged by unsystematically shifting objects or the camera system, and the scene was rerecorded, until an image with three distinct predictions was obtained.

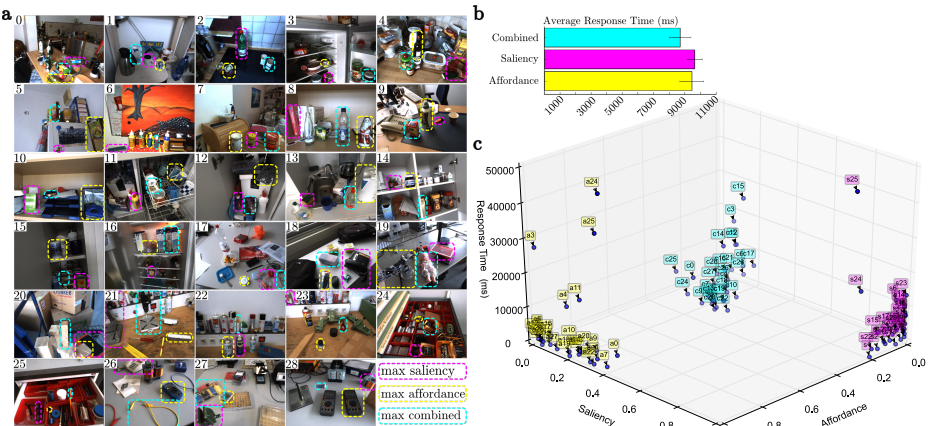


Fig. 4. a: Changed locations marked in the original images (best viewed magnified in the digital version). **b:** Average response times for the changes based on affordance, saliency (region-based), and combined predictions (error bars show the SEM). **c:** Distribution of average response times (in milliseconds) per image and change, visualized as in figure 2b. The numbers correspond to figure 4a.

Results and Discussion: Figure 4b shows the average response times to changes based on the (region-based) saliency, affordance, and combined predictions. According to an one-way repeated measures ANOVA, no effect of prediction type was found, $F(2, 29) = 0.63$, $p = 0.54$. This is in contrast to the result of our first experiment, where the responses to affordance-based changes were significantly faster. Furthermore, the saliency conditions (from experiment 1 and experiment 2), as well as the affordance conditions (from each experiment), differ significantly, $t(40) = 5.82$, $p < 0.001$ (affordance), $t(40) = 4.3$, $p < 0.001$ (saliency) according to Holm-Bonferroni corrected two-sided t-tests.

The long response times in the second experiment indicate that the task was more difficult than in the first experiment. Moreover, the scenes were arranged to contain a large number of affording objects in the action space and thus also the saliency-based selections were mainly such foreground objects, whereas the stimulus material used in the first experiment contained saliency-based changes which were frequently in the background. In addition to the different scene arrangement, the region-based saliency model may less often select locations in the background where the contrasts are low, because such parts are often merged into larger background regions with low saliency (see figure 1d).

Inspecting the distribution the individual changes' average response times in the affordance-saliency space (figure 4c), it can be seen that while the saliency-based changes benefit from increasing affordance (response times decrease along the affordance axis), the same is not the case for the affordance-based changes which do not show a similar pattern. The changes based on the combined prediction cluster at lower saliency and affordance values (which were obtained by averaging the respective maps under the bounding rectangle of the removed combined object). Those, which are particularly low, resulted in longer response times. Notably, the one minute limit was reached twelve times for affordance- and nine times for saliency-based changes, and only once in the combined condition.

4 Growing Neural Gases as an Architecture for Combining Saliency, Top-Down Attention and Affordances

The attention model in the present paper uses separate pipelines for estimating affordances and saliency. Affordances are estimated based on the sparse scene representation by a 3D feature hierarchy [18] and then fed into the region-based attention framework [24, 25], in which saliency is computed, and for which top-down mechanisms are available, e.g. [26, 27].

In the long term, a fully integrated architecture is desirable. In such a framework, feedback loops, which are known to be highly important in biological vision, can be established: results from higher levels of the architecture, such as affordances or top-down information can be used to influence the generation or propagation of the scene representation from lower levels. Furthermore, the integration of different dimensions, such as various saliency dimensions, specific and gist-based top-down influences and object affordances benefits from a common

consistent representation which can preserve the relative strength of each dimension's contribution (a prerequisite for more advanced combination strategies as suggested in [12]).

In this section we discuss Growing Neural Gases (GNG; e.g. [28]) as pre-attentional structures for a fully integrated approach. GNGs have the promising feature that they can be continuously updated as new visual information is acquired. It is not required to completely start over with every new frame, a strong drawback of current region-based architectures, which impairs the stability of pre-attentional structures over time.

The GNG algorithm is an unsupervised learning technique, which uses a growing network model. Nodes within a graph may be connected by edges and possess attributes which represent the properties of the state space in which the learning is performed (e.g. x- and y-positions if two dimensional positions are learned).

While learning, examples are presented to the network. The algorithm determines the node closest to the example according to a distance measure on the respective attributes of node and example. For this winning (closest) node and, with a reduced strength, its topological neighbors (i.e. those directly connected via a single edge) the properties are updated. Edges carry an age value, which is increased with every update step. Additionally, a new edge is inserted between the winning node and the runner-up or, if such an edge already exists, its age is reset. Edges which are too old are deleted. Furthermore, each node carries a problem domain dependent error value, which must have the characteristic that it is reduced when a new node is inserted close to the given node. In GNGs this insertion is done at fixed intervals. Therefore, the node with the highest error value is determined, as well as its topological neighbor with the highest error value. The new node is inserted between and connected with these two nodes, replacing the original connection between them. The error value is redistributed between all three nodes. This reduces the probability that the next insertion is performed in the same area and thereby guides the growing of the network. Further measures, such as an utility term that quantifies the usefulness of a node for the network and may result in deletion of the node, are introduced to avoid the network to grow infinitely.

The algorithm is initialized with two connected nodes, which carry randomly initialized properties and it results in a dynamically adapting set of nodes, with changing neighborhoods. These neighborhoods can form multiple independent graphs.

4.1 Pre-Attentional Structures and Saliency based on GNGs

GNGs can be adapted to generate pre-attentional structures similar to the regions produced in the segmentation step of the region-based framework (see section 3.1). Pixels of the image are chosen uniformly at random and used as examples. These pixels have a position and color, which can be used to train the GNG as described above. Distances of each neuron to the chosen pixel can be

calculated by using, e.g., a weighted euclidean distances. Figure 5a depicts the result of this procedure applied to a simple synthetic test image.




Fig. 5. a: Result of a prototypical implementation of GNG for generating pre-attentional structures. Node colors reflect the color of the represented object. **b:** GNGs may include useful candidates (two-node networks; white) for estimating grasp affordances. Networks with more than two nodes are colored in gray.

For the resulting graphs, feature magnitudes and then saliency can be calculated using the methods described in [24]. Hence, where the positions of a region’s pixels are considered to calculate central moments or other features, the GNG nodes and their positions may be used in the same way.

4.2 Applying Top-Down Information in GNGs

As argued above, mechanisms from region-based attention can be easily transferred to GNG structures. Therefore, template networks can be obtained and considered as described in [26] for region-based attention. Complex templates consisting of several networks represent the analog to multi-region templates introduced in [27].

Furthermore, as the transfer from pixels to the substantially smaller number of neurons provides a simplified problem space, rough heuristics, such as for determining the gist of a scene, may also benefit from a GNG-based representation.

4.3 Integrating Object Affordances in GNGs

On the one hand, pre-attentional structures obtained from GNGs may prove sufficiently stable (as compared to segmented regions that often exhibit unstable shapes) to apply appearance-based affordance estimation in local and global contexts as proposed by [29]. On the other hand, GNGs obtained with the described procedure may provide easily identifiable candidates for graspable elements. Figure 5b shows GNGs obtained from a picture of a cooking pot with a handle extracted from Song et al. [29]’s figure 4b (the pink dot was removed). Highlighting only two-node networks, in agreement with the ground-truth for such handles (see Song et al. [29]’s figure 2a, successfully detects the pot’s handle). Such rough heuristics could be directly useful for generating local graspability estimates which can then be fused with global estimates as in [29], or at least provide candidates for more expensive follow-up processing.

5 Discussion and Conclusion

The results of our first experiment further supports the idea that object affordances are important for the spatial deployment of visual attention. In the second experiment we did not find additional enhancements by combining saliency and affordance. This is in line with another change blindness study [30], where saliency did not further enhance the detection of changes in objects which are shown in unusual contexts. Early attention appears to be strongly influenced by the environment represented in the scene. The second experiment, however, did also replicate the advantage of the affordance-based predictions over the saliency-based predictions. This can originate from the use of a different saliency model, from the use of a different set of stimuli, or a combination of both. The region-based saliency model is not expected to yield predictions so much different in general (see [25] for comparisons with the model from [11]). However, as explained in section 3.2, the arrangements in the test scenes differ from those used in [17] in important aspects:

The scenes used in [25] contained salient—but ungraspable—objects in the background, which were frequently selected by the saliency model. The scenes used in experiment 2 focus on objects in the action space, as the goal was to provide a variety of action space objects, which can exhibit different saliencies. Several scenes, especially scenes 24 and 25 (see figure 4a), contain a large amount of competing targets resulting in extreme response times for saliency as well as affordance changes (see figure 4c). Hence, this stimulus material may not be sensitive enough to reveal differences which may be small for these rather similar action space objects.

An important next step in this line of research is a deeper integration of affordance in attention systems. In the present study, affordance and saliency were estimated based on separate low-level scene representations (a 3D feature hierarchy and simple image regions). A more integrated approach (that does not require a step such as the projection of 3D grasp hypotheses into 2D regions) allows more sophisticated combination strategies on the basis of common pre-attentional structure. As suggested by our results, future models that combine saliency and affordance could employ strategies strongly biased towards affordances when these are available in the scene and fall back to saliency only when no other information is available.

The fact that affordance-based advantages are present in 2D images presented to humans which depict foregrounds and background (as in our first experiment, and the experiments reported in [25]), proves that binocularly cues are not necessary for the effect in biological vision. Thus, a 2D dimensional retinotopical structures would provide a useful domain for such a fully integrated approach. We discussed Growing Neural Gases as a framework for this in section 4, which may allow to integrate appearance based affordance-estimation as suggested by Song et al. [29] with bottom-up and top-down attention in future work.

Our current approach focuses on estimating the affordance of scene elements with respect to actions the system is able to perform towards these elements. Following a different approach, work by Varadarajan & Vincze [31] includes

further aspects of environmental psychology, such as the affordance relations between the represented objects, in an attention model. A form of semantic saliency is directly derived from such affordances. The procedure presented in [31] obtains saliency from affordance aberrations, a measure of how unusual an object appears in its local semantic context. In addition to successful application in artificial vision, such models well predict the findings of [30] and others, which indicate that semantic scene context outranks pure stimulus driven saliency in early human vision.

With regard to our specific approach, a fully integrated approach and the implementation of this attention system based on grasp affordances in a robot with an articulated arm and gripper will allow a practical evaluation of the proposed concepts.

References

1. Gibson, J.J.: *Perceiving, Acting, and Knowing: Toward an Ecological Psychology.* (1977)
2. Detry, R., Kraft, D., Kroemer, O., Bodenhausen, L., Peters, J., Krüger, N., Piater, J.: Learning Grasp Affordance Densities. *Paladyn* **2**(1) (2011) 1–17
3. Varadarajan, K.M., Vincze, M.: Affordance Based Part Recognition for Grasping and Manipulation. In: *ICRA Workshop on Autonomous Grasping.* (2011)
4. Varadarajan, K.M., Vincze, M.: Object Part Segmentation and Classification in Range Images for Grasping. In: *15th International Conference on Advanced Robotics.* (2011) 21–27
5. Stark, L., Bowyer, K.: Generic Recognition Through Qualitative Reasoning About 3-D Shape and Object Function. In: *IEEE Conference on Computer Vision and Pattern Recognition.* (1991) 251–256
6. Rivlin, E., Dickinson, S., Rosenfeld, A.: Recognition by Functional Parts [Function-Based Object Recognition]. In: *IEEE Conference on Computer Vision and Pattern Recognition.* (1994) 267–274
7. Castellini, C., Tommasi, T., Noceti, N., Odone, F., Caputo, B.: Using Object Affordances to Improve Object Recognition. *IEEE Transactions on Autonomous Mental Development* **3**(3) (2011) 207–215
8. Gijssberts, A., Tommasi, T., Metta, G., Caputo, B.: Object Recognition Using Visuo-Affordance Maps. *IEEE/RSJ International Conference on Intelligent Robots and Systems* (2010) 1572–1578
9. Yao, B., Ma, J., Fei-Fei, L.: Discovering Object Functionality. In: *IEEE International Conference on Computer Vision.* (2013) 2512–2519
10. Zhao, Y., Zhu, S.C.: Scene Parsing by Integrating Function, Geometry and Appearance Models. In: *IEEE Conference on Computer Vision and Pattern Recognition.* (2013) 3119–3126
11. Itti, L., Koch, C.: Feature Combination Strategies for Saliency-Based Visual Attention Systems. *Journal of Electronic Imaging* **10**(1) (2001) 161–169
12. Navalpakkam, V., Itti, L.: A Goal Oriented Attention Guidance Model. In: *Biologically Motivated Computer Vision.* (2002) 453–461
13. Oliva, A., Torralba, A.: Building the Gist of a Scene: The Role of Global Image Features in Recognition. *Progress in Brain Research* **155** (2006) 23–36

14. Roberts, K.L., Humphreys, G.W.: Action-Related Objects Influence the Distribution of Visuospatial Attention. *Quarterly Journal of Experimental Psychology* **64**(4) (2011) 669–88
15. Garrido-Vásquez, P., Schubö, A.: Modulation of Visual Attention by Object Affordance. *Frontiers in Psychology* **5** (2014) 59
16. Handy, T.C., Grafton, S.T., Shroff, N.M., Ketay, S., Gazzaniga, M.S.: Graspable Objects Grab Attention When the Potential for Action is Recognized. *Nature Neuroscience* **1** (2003) 1–7
17. Tünnermann, J., Krüger, N., Mertsching, B., Mustafa, W.: Affordance Estimation Enhances Artificial Visual Attention: Evidence from a Change Blindness Study. (in review)
18. Popović, M., Kootstra, G., Jørgensen, J.A., Kragic, D., Krüger, N.: Grasping Unknown Objects Using an Early Cognitive Vision System for General Scene Understanding. *IEEE/RSJ International Conference on Intelligent Robots and Systems* (2011)
19. Rensink, R.A., O’Regan, J.K., Clark, J.J.: To See or Not to See: The Need For Attention to Perceive Changes in Scenes. *Psychological Science* **8**(5) (1997) 368–373
20. Tünnermann, J., Mertsching, B.: Saliency and Affordance in Artificial Visual Attention. In: *RSS 2014, First Workshop on Affordances: Affordances in Vision for Cognitive Robotics*. (2014)
21. Tseng, P., Tünnermann, J., Roker-Knight, N., Winter, D., Scharlau, I., Bridgeman, B.: Enhancing Implicit Change Detection Through Action. *Perception* **39**(10) (2010) 1311–1321
22. Aziz, M.Z., Shafik, M.S., Mertsching, B., Munir, A.: Color Segmentation for Visual Attention of Mobile Robots. In: *IEEE Symposium on Emerging Technologies*. (2005) 115–120
23. Backer, M., Tünnermann, J., Mertsching, B.: Parallel k-Means Image Segmentation Using Sort, Scan and Connected Components on a GPU. In Keller, R., Kramer, D., Weiss, J.P., eds.: *Facing the Multicore-Challenge III*. (2013) 108–120
24. Aziz, M.Z., Mertsching, B.: Fast and Robust Generation of Feature Maps for Region-Based Visual Attention. **17**(5) (2008) 633–644
25. Tünnermann, J., Mertsching, B.: Region-Based Artificial Visual Attention in Space and Time. *Cognitive Computation* **6**(1) (2014) 125–143
26. Aziz, M.Z., Mertsching, B.: Visual Search in Static and Dynamic Scenes Using Fine-Grain Top-Down Visual Attention. In: *Computer Vision Systems*. (2008) 3–12
27. Tünnermann, J., Born, C., Mertsching, B.: Top-Down Visual Attention with Complex Templates. In: *International Conference on Computer Vision Theory and Applications*. (2013) 370 – 377
28. Fritzke, B.: A Self-Organizing Network That can Follow Non-Stationary Distributions. In: *Artificial Neural Networks – ICANN’97*. (1997) 613–618
29. Song, H.O., Fritz, M., Gu, C., Darrell, T.: Visual Grasp Affordances from Appearance-Based Cues. In: *IEEE ICCV Workshops*. (2011) 998–1005
30. Stirk, J.A., Underwood, G.: Low-Level Visual Saliency Does not Predict Change Detection in Natural Scenes. *Journal of Vision* **7** (2007) 1–10
31. Varadarajan, K.M., Vincze, M.: Semantic Saliency Using k-TR Theory of Visual Perception. In: *21st International Conference on Pattern Recognition*. (2012) 3676–3679