

Task-driven Processing with Coarse-to-Fine Glimpse-based Active Perception

Oleh Kolner^{1,2}, Thomas Ortner¹, Stanisław Woźniak¹, and Angeliki Pantazi¹

¹ IBM Research, Zurich, Switzerland

² Graz University of Technology, Graz, Austria
olk@zurich.ibm.com

Abstract. State-of-the-art vision models process images in their entirety, lacking the ability to selectively zoom in on relevant regions. This limitation is particularly acute in scenarios where processing must be conditioned on a specific task – such as instance detection, which requires localizing a specific object in a high-resolution, cluttered scene. In such settings, critical details are easily lost as images are often resized to match the model dimensions and computational constraints. We introduce Coarse-to-Fine Glimpse-based Active Perception (CF-GAP), a task-driven front-end that enhances high-resolution processing of existing instance detectors. CF-GAP selectively directs a sequence of limited-view glimpses across the scene, utilizing task information to iteratively refine focus on the most relevant regions. These localized regions are then processed at high resolution by a downstream instance detector. By avoiding full-image processing and eliminating irrelevant confounding information, CF-GAP improves Average Precision (AP) by up to 20% across various state-of-the-art instance detectors on the HR-InsDet and Robotools benchmarks, while further enabling lightweight detectors to outperform their larger counterparts.

Keywords: Active perception · Bio-inspired vision · Instance detection

1 Introduction

State-of-the-art vision models process entire images uniformly, lacking the ability to selectively zoom into task-relevant regions for detailed analysis. Moreover, they require images to be resized to the fixed input dimensions used during pre-training, often obscuring small but important regions. These limitations hinder performance on high-resolution images where task-relevant content is small or cluttered [37, 43, 50]. A further shortcoming is that task-specific information (*e.g.*, a textual prompt or a specific object to be localized) is typically incorporated only after a costly feature extraction stage. However, early integration of such information could significantly reduce the image area requiring intensive computation and mitigate the influence of irrelevant or distracting features.

Human vision, by contrast, is inherently task-driven: we actively seek visual information based on what we aim to accomplish, rather than passively




Fig. 1: Coarse-to-Fine Glimpse-based Active Perception (CF-GAP) iteratively directs a series of glimpses across a high-resolution scene, using task-driven search maps to progressively narrow focus onto the likely search target. The resulting regions of interest are passed at high resolution to the downstream architecture for the final detection.

processing every detail [11, 45]. For example, in an office environment shown in Fig. 1, one often needs to find a specific box or tool – not just any object from those categories. This targeted search corresponds to instance detection, where the goal is to localize a specific object instance given a few visual examples, as opposed to classical object detection, which seeks to identify all objects in a scene. As opposed to typical vision models, humans actively leverage instance-specific features as task cues for efficient visual search [41]. A further hallmark of human vision is the foveal structure of the eye, which provides non-uniform resolution – highest at the center and decreasing toward the periphery – balancing detailed information with a broad field of view [10, 36]. Humans exploit this balance through eye movements (saccades) that follow a coarse-to-fine strategy: macrosaccades, guided by low-resolution cues, direct gaze to promising areas, which are then analyzed in detail through microsaccades [28, 34, 46].

Inspired by saccadic processing, the recently proposed Glimpse-based Active Perception (GAP) model [20] fixates on salient – not necessarily task-relevant – image regions to solve synthetic visual reasoning tasks, demonstrating strong out-of-distribution generalization and, thus, the effectiveness of processing only selected image parts. Building on this, we propose Coarse-to-Fine GAP (CF-GAP), which replaces uniform image processing by selectively directing a sequence of limited-view glimpses (Fig. 1). This sequence is generated through a nested process of coarse and fine glimpring, guided by a task-specific object to be localized, referred to as a search target. Coarse glimpring relies on correlations

between the search target’s features and the downsampled scene to identify an initial coarse glimpse location. There, a fovea-inspired log-polar sensor extracts a limited-view glimpse from the full-resolution scene, magnifying the visual details around the glimpse location while preserving peripheral information at diminishing resolution. Operating on this focused view, fine glimpsing directs the log-polar sensor in a closed loop to iteratively refine the initial location toward a likely search target. The resulting fine glimpse locations define a localized region of interest (RoI) that is passed to a downstream architecture at high-resolution to determine whether the search target is present at those locations.

Our work focuses on the instance detection problem, where visual processing must be conditioned on a specific search target. We design CF-GAP as a frontend module that integrates seamlessly with existing instance detectors as downstream architectures, providing them with targeted, high-resolution input. We combine CF-GAP with several state-of-the-art detectors and evaluate on HR-InsDet [37] and Robotools [21] as two challenging benchmarks. We observe significant performance improvement across all models. Notably, integrating CF-GAP with small models optimized for edge devices allowed them to match, and in some difficult cases, even surpass the performance of their larger versions.

To summarize, our contributions are the following:

1. We introduce a bio-inspired, coarse-to-fine processing scheme to selectively explore high-resolution scenes, replacing exhaustive full-scene processing.
2. We design our model as a front-end module that can be seamlessly integrated with any existing instance detector, enhancing it with targeted, high-resolution input.
3. We show that our approach significantly boosts the performance of state-of-the-art instance detectors, particularly in cluttered and complex scenes.

2 Related Work

Early instance detection methods trained CNN-based object detectors with a separate class per object instance, using the cut-paste-learn (CPL) framework [6], where objects are pasted into random backgrounds. State-of-the-art instance detection methods [24,37,38] adopt a two-stage pipeline. First, all candidate objects in the scene are detected using foundation models pre-trained on large-scale data, in particular Segment Anything Model (SAM) [19,32] and GroundingDINO [23]. Second, the candidates are matched against visual examples of the search target using DINOv2 [29] features, with the best match returned as the final prediction. Recent extensions improve the matching stage by fine-tuning DINOv2 [38] or training a weight adapter [24] on specific object instances. Critically, all the aforementioned methods are from the ground up not task-driven, since the first stage of proposal detection involves exhaustive analysis of the entire image to detect all possible objects irrespective of the search target. Our approach leverages the information about the task of detecting a specific search target to steer detection toward only promising regions.

Multiple approaches drew high-level inspiration from saccadic processing and showed compelling results across various tasks, including image classification [8, 27, 44], object detection [2, 13, 14], visual exploration [31], and visual reasoning [20, 42]. However, none of them allows for conditioning by task information (*e.g.* by a search target) at inference to guide the image exploration. While our approach is built upon GAP [20], the original framework lacks a coarse-to-fine, task-driven search strategy.

Fovea-inspired vision with non-uniform resolution was explored for guiding the visual search only in simple images [1, 3]. Various methods explored fovea-like log-polar imagery and its rotational and scaling equivariance properties for image classification, object detection, and image correspondence [7, 9, 18]. Another approach [35] proposed a sophisticated tokenization for vision transformers (ViTs) [5], splitting the image into patches of various sizes resembling foveated image structure. While prior work used fovea-inspired imagery mainly to encode a fixed view, we use it instead to navigate the coarse-to-fine glimpsing process.

3 Method

Conceptually, CF-GAP can be viewed as an active process of steering a virtual log-polar sensor across a scene to acquire high-quality information related to the search target (Fig. 1). More specifically, CF-GAP directs a sequence of glimpses – limited views of an image taken at specific locations – to pinpoint a likely search target location. This location defines the center of a RoI that is passed to the downstream architecture, allowing it to isolate and verify the candidate object, without exhaustively processing the entire scene. CF-GAP operates in two nested stages of coarse and fine glimpsing guided by distinct search maps. The search maps are 2D heatmaps that highlight regions with a high probability of containing the search target. The search map extraction follows a common scheme from [47, 51]: a scene encoder and a search target encoder produce features that are compared via convolution (Fig. 2A). The coarse and fine glimpsing stages differ in how this scheme is instantiated, as described below, with further technical details provided in Appendix A.

3.1 Coarse glimpsing

CF-GAP begins with computing a coarse search map from a downsampled version of the scene and sample images of the search target taken from different viewpoints (referred to as search target examples). The scene and search target encoders (Fig. 2A) are instantiated with the lightweight MobileNet-V3 [12]. The features of the search target examples are averaged across spatial dimensions into a single feature vector, which is then convolved over the scene features, producing the coarse search map. This map drives the iterative selection of coarse glimpse locations: at each iteration, the location of the highest value is selected via the winner-takes-all (WTA) strategy, followed by inhibition-of-return (IoR) that masks its surrounding region to prevent repeated selection, similar to [15, 20].

Since coarse glimpse locations are derived from low-resolution imagery, they represent only rough estimates of the search target’s position. Therefore, each coarse glimpse location initiates a closer inspection with fine glimpsing.

3.2 Fine glimpsing

Unlike coarse glimpsing, which operates on a static downsampled scene, fine glimpsing operates on glimpses of both the scene and the search target. It is a closed-loop process: at each iteration, a fine search map is generated from the current scene glimpse, and its 2D centroid determines the next fine glimpse location. The log-polar sensor then moves to that location and extracts a new scene glimpse, from which the next iteration proceeds. The sensor employs a fovea-like log-polar transformation to extract high-resolution detail near the glimpse location




Fig. 2: (A) General architecture for extracting coarse and fine search maps. (B) Scene encoder compresses the scene glimpse into a compact set of latent embeddings via cross-attention, where a small number of learnable embeddings attend to the full spatial input ($N \ll H \times W$). The compressed representation is then processed through self-attention. A second cross-attention decodes the resulting representation back to the original spatial dimensions, where a distinct set of learnable value embeddings defines the output feature space. Q, K, and V denote queries, keys, and values of each attention block. (C) Search target encoder shares the cross- and self-attention blocks with the scene encoder (marked by the same color) and iteratively attends to multiple search target glimpses to produce compact target features to be compared with the glimpse feature map.

while preserving distant context at progressively lower resolution. The resulting focused view makes fine search map generation robust against peripheral distractors, and we analyze its advantages over cartesian cropping in Sec. 5.

Since the fine search map extraction (Fig. 2A) operates on log-polar glimpses, using CNN-based scene and search target encoders (as for the coarse search map) becomes problematic. Specifically, CNNs assume that the input image has uniform resolution, treating all regions equally regardless of their position within that image. By contrast, log-polar glimpses preserve fine detail near the glimpse location while compressing the periphery into much fewer pixels. We therefore propose modules based on Perceiver [16], illustrated in Fig. 2B-C.

Each glimpse is partitioned into non-overlapping patches projected into a D -dimensional space, $\mathbf{F}_s \in \mathbb{R}^{H \times W \times D}$ and $\{\mathbf{F}_m\}_{m=1}^M$ denote the corresponding representations of the scene glimpse and M search target glimpses, respectively. The scene encoder (Fig. 2B) first compresses the scene glimpse into a compact set of N latent embeddings via cross-attention, where a small number of learnable queries attend to the full spatial input, and then processes them through self-attention:

$$\tilde{\mathbf{h}}_s = \mathcal{C}(\mathbf{Q} = \mathbf{e}_q, \mathbf{KV} = \mathbf{F}_s), \quad \mathbf{h}_s = \mathcal{S}(\mathbf{QKV} = \tilde{\mathbf{h}}_s), \quad (1)$$

where $\mathbf{e}_q \in \mathbb{R}^{N \times D}$ is a set of learnable embeddings and $N \ll H \times W$. The compression via cross-attention alleviates the quadratic cost of the subsequent self-attention. The self-attention, in turn, integrates fine local details near the glimpse center with the broader peripheral context. Finally, the latent embeddings are decoded back to the original spatial dimensions through a second cross-attention, producing the glimpse feature map $\mathbf{F}_s^*$ as the final output:

$$\mathbf{F}_s^* = \mathcal{C}_{\text{dec}}(\mathbf{Q} = \mathbf{F}_s, \mathbf{K} = \mathbf{h}_s, \mathbf{V} = \mathbf{e}_o). \quad (2)$$

The decoding cross-attention uses a distinct set of learnable embeddings $\mathbf{e}_o \in \mathbb{R}^{N \times D}$ as values to decouple the feature space used for latent compression from the one used for the output feature map. This is intended to let the output feature map be optimized specifically for comparison with features extracted from the search target.

The search target encoder (Fig. 2C) follows the same compress-and-process pattern with shared cross- and self-attention blocks, but instead of decoding back to spatial dimensions, it iteratively attends to M search target glimpses to produce a set of compact target features. At each iteration m , the cross-attention receives the search target glimpse $\mathbf{F}_m$ as keys and values and the previous output as queries, followed by self-attention:

$$\tilde{\mathbf{h}}_T^{(m)} = \mathcal{C}(\mathbf{Q} = \mathbf{h}_T^{(m-1)}, \mathbf{KV} = \mathbf{F}_m), \quad \mathbf{h}_T^{(m)} = \mathcal{S}(\mathbf{QKV} = \tilde{\mathbf{h}}_T^{(m)}), \quad (3)$$

with the query at the first iteration initialized to $\mathbf{h}_T^{(0)} = \mathbf{e}_q$ and the final output $\mathbf{h}_T = \mathbf{h}_T^{(M)}$, $\mathbf{h}_T \in \mathbb{R}^{N \times D}$ used as the set of target features. These target features are then correlated with the glimpse feature map via convolution, yielding N correlation maps that are averaged into a single fine search map.

3.3 Downstream architecture

At the end of each fine-glimpsing loop, CF-GAP provides the downstream architecture with three inputs. First, a fixed-size RoI cropped around the last fine glimpse location. Second, the fine glimpse location itself, which, depending on the downstream architecture, serves either as a spatial prompt (*e.g.* for SAM-like models) or as a bounding box filter to constrain object detection. Third, multiple visual examples of the search target from different viewpoints, used for matching with the detected candidate object. After a predefined number of coarse and fine glimpsing iterations, the best-matched candidate object is returned. Importantly, CF-GAP is agnostic to the choice of downstream architecture. During evaluation, we employ several state-of-the-art detectors as described in Sec. 4.

4 Experiments

Instance detection consists of individual tasks, each defined by a few visual examples of a specific object to be localized in an input scene. Unlike classical object detection, which identifies all instances of an object category, instance detection is conditioned on a particular object instance. We consider two benchmarking datasets. The first, HR-InsDet [37], contains 100 object instances with 24 visual examples each, and 160 high-resolution scenes spanning 14 indoor scenarios. For training, the dataset provides 200 images with random backgrounds to synthesize training data via the cut-paste-learn strategy [6], where search targets are resized and pasted onto arbitrary backgrounds. The dataset is split into subsets by the level of clutter and occlusion – *easy* and *hard* – and by object size – *small*, *medium*, and *large*. We report results for each subset following the HR-InsDet evaluation protocol. The second benchmark, Robotools [21], contains 20 object instances and 1581 test images from 24 indoor scenarios. Unlike HR-InsDet, Robotools prohibits using its 20 search targets for training, thereby evaluating generalization to novel objects. Accordingly, we train CF-GAP using only objects from HR-InsDet. We report average precision (AP) at Intersection-over-Union (IoU) thresholds from 0.5 to 0.95 in steps of 0.05, as well as AP₅₀ at an IoU threshold of 0.5.

Baselines. The strongest baselines, OTS-FM [37], IDOW [38], and NIDS-Net [24], employ pre-trained foundation models to process the entire scene, first detecting bounding boxes for all object-like regions (proposals). A feature extractor then generates embeddings for each proposal, which are matched to the search target’s examples via Stable Matching [25]. All methods use either SAM [19] or GroundingDINO [23] for proposals and DINOv2 [29] for feature extraction. IDOW extends OTS-FM by fine-tuning DINOv2 on search targets from HR-InsDet. However, we exclude IDOW from our evaluations as its fine-tuned weights are unavailable, precluding integration with our CF-GAP front-end. We do include NIDS-Net, which follows a similar but higher-performing approach: it trains a weight adapter for DINOv2 and additionally uses SAM

to mask out backgrounds within each proposal. The core baseline set comprises OTS-FM_{SAM} and OTS-FM_{GroundingDINO}, depending on proposal detector, and NIDS-Net, which uses GroundingDINO for proposal detection. We additionally consider two efficient SAM variants as OTS-FM backbones: MobileSAM [48], a distilled version of SAM, and Segment This Thing (STT) [35], which uses a fovea-inspired tokenization that partitions the image into patches of increasing size with distance from a given location. These are denoted OTS-FM_{MobileSAM} and OTS-FM_{STT}. Since STT requires a location input for the tokenization, it can only be evaluated in combination with CF-GAP and is thus excluded from the main pairwise comparisons between standalone baselines and their CF-GAP extensions. Unless stated otherwise, all baseline results are reproduced using the publicly available code.

Setup. We integrate CF-GAP with each baseline as its downstream architecture. For the OTS-FM_{MobileSAM} and OTS-FM_{SAM} baselines, CF-GAP changes the input structure: standalone baselines receive the full high-resolution scene together with a coarse grid of 2D point prompts to detect proposals at each grid location. By contrast, CF-GAP provides only a small RoI along with fine glimpse locations as 2D point prompts at the end of each fine glimpsing loop, improving both efficiency and focus. As the OTS-FM_{GroundingDINO} baseline does not support point-based prompting, the fine glimpse locations are used to filter out detected bounding boxes that do not contain them.

In HR-InsDet and Robotools, scenes are sized at 6144×8192 and 1920×1080 pixels, respectively, whereas baseline models require resizing below 2048×2048 . CF-GAP supports flexible input sizes; for faster experimentation, HR-InsDet scenes were resized to 4096×5460 while Robotools scenes were kept at original resolution. Coarse glimpsing operates on scenes downsampled by a factor of 2; fine glimpsing operates at full resolution. The log-polar sensor diameter is set to 4096 pixels, and the resulting glimpses are resized to 245×245 . Each scene undergoes N_c coarse glimpses, each followed by N_f fine glimpses, with $N_c=30$ and $N_f=3$ by default unless stated otherwise. Further details are provided in Appendix B.

5 Results

5.1 Benchmarking results

As shown in Figs. 3 and 4, CF-GAP consistently improves all baseline models, demonstrating the effectiveness of the coarse-to-fine glimpsing. In evaluations on HR-InsDet dataset, the biggest benefits are observed in scenes with small-sized objects and hard scenes, *i.e.* scenes with high clutter and partial occlusions. OTS-FM_{MobileSAM}, the smallest baseline, benefits the most, achieving up to 20% AP improvement on hard scenes. With CF-GAP, it surpasses both OTS-FM_{SAM} and OTS-FM_{GroundingDINO} on the difficult subsets while performing comparably on simpler ones. On Robotools, CF-GAP also consistently improves AP of all baselines. Lightweight OTS-FM_{MobileSAM} paired with CF-GAP achieves AP₅₀ comparable even with the heavier baselines; its AP, however, stays below theirs,




Fig. 3: Performance on HR-InsDet dataset across two dataset groupings – by search target sizes, and by scene types.




Fig. 4: Performance on Robotools dataset, testing generalization to novel objects unseen during training. Baseline results in dashed bars are taken from [24, 38]

indicating less precise bounding boxes due to its weaker detector backbone. We provide more extensive tabular comparisons in Appendix C.

5.2 Analysis and Ablations

Naive high-resolution baseline. To demonstrate the importance of CF-GAP, we compare it against a straightforward alternative for processing high-resolution images. Baseline models cannot handle scenes at their original resolution because their ViT-based backbones decompose images into a fixed number of patches to limit the quadratic cost of self-attention. A naive solution is to split each scene into smaller overlapping patches and let the baseline treat each patch as a separate RoI. Tab. 1 reports results for this approach, where the full-sized (6144×8192) scenes are split into 1024×1024 patches with 50% overlap, yielding 165 RoIs per scene – over $5 \times$ more than the 30 RoIs produced by CF-GAP. OTS-FM models without fine-tuned matching do not benefit from such patching, as the larger number of candidate proposals across all RoIs leads to increased matching errors. NIDS-Net, whose fine-tuned features better discriminate the search target, does benefit – particularly for small objects – but still falls behind

Table 1: Comparison to patched baselines that split the scene into smaller, separately processed patches. Results on the two most challenging HR-InsDet subsets.

Model	AP	
	small	hard
OTS-FM _{MobileSAM}	12.4	22.0
Patched OTS-FM _{MobileSAM}	18.4	23.2
CF-GAP + OTS-FM _{MobileSAM}	29.3	41.7
OTS-FM _{SAM}	14.6	28.0
Patched OTS-FM _{SAM}	18.9	24.2
CF-GAP + OTS-FM _{SAM}	32.4	47.1
OTS-FM _{GroundingDINO}	28.8	37.2
Patched OTS-FM _{GroundingDINO}	22.4	28.1
CF-GAP + OTS-FM _{GroundingDINO}	39.6	50.2
NIDS-Net	32.4	39.9
Patched NIDS-Net	50.1	49.1
CF-GAP + NIDS-Net	52.1	63.2

its CF-GAP extension, especially in hard, cluttered scenes. These results confirm that while exhaustive patching can in principle recover lost high-resolution detail, CF-GAP is both more effective – achieving higher AP, and more efficient – passing over $5\times$ fewer RoIs to the downstream architecture.

Glimpse locations. In addition to providing targeted RoIs to downstream architectures, CF-GAP also provides fine glimpse locations that specify where objects have to be detected within each RoI. For SAM-based models, this improves the efficiency by replacing the dense grid of point prompts with only a few glimpse locations. For GroundingDINO-based models, the fine glimpse locations act as spatial filters that exclude candidate objects that do not overlap with them, substantially reducing the number of candidates and thereby improving matching effectiveness. This is confirmed in Tab. 2, which reports higher performance when the downstream architecture receives both the RoIs and the glimpse locations, compared to receiving only the RoIs.

Table 2: Ablation of inputs provided to the downstream architecture. Results are shown for CF-GAP+NIDS-Net evaluated on HR-InsDet.

Downstream architecture input	AP				
	small	medium	large	easy	hard
w/o locations	41.8	73.1	86.0	74.5	49.5
with locations	52.1	79.7	86.1	78.2	63.2

Fine glimpsing. Given that CF-GAP consists of coarse and fine glimpsing processes, it is important to demonstrate the need for the latter. Tab. 3 shows that fine glimpsing is most beneficial in hard scenes and scenes with small-sized objects. This is because the low resolution of the coarse search map can yield

glimpse locations that are offset from the actual object, disrupting downstream architectures that are prompted to detect objects at specific locations. Fine glimpsing corrects for this spatial error (see Fig. 7 for examples). In simpler cases where objects are easy to find, fine glimpsing provides marginal benefit.

Table 3: Ablation of fine glimpsing. Results are shown for CF-GAP+NIDS-Net evaluated on HR-InsDet.

Fine glimpsing	AP				
	small	medium	large	easy	hard
no	44.6	77.3	85.5	75.6	56.2
yes	52.1	79.7	86.1	78.2	63.2

Log-polar glimpses. Another study shows the advantage of extracting log-polar glimpses. Compared to a naive approach of extracting a small cartesian crop, the log-polar representation eliminates the need to tune the crop size, which would otherwise be highly sensitive to the proximity of the initial coarse glimpse location to the object and to the object’s size. For example, the crop size can be either too small, providing too little information, or too big, providing too




Fig. 5: Qualitative comparison between crop-based and log-polar glimpses across different sizes. Red dots mark the initial coarse glimpse location, offset from the search target, and green triangles mark the subsequent fine glimpse location based on fine search maps extracted from each glimpse. Unlike crop-based glimpses (top row), which change drastically with size, log-polar glimpses (bottom row) remain visually stable across different diameters, as peripheral differences are compressed into the far-right part of each log-polar image. The middle row shows the regions covered by the log-polar glimpses.




Fig. 6: Performance comparison for different glimpse types and their sizes. Results are for CF-GAP+NIDS-Net evaluated on the two most challenging HR-InsDet subsets.

much distraction (Fig. 5, top row). By contrast, due to logarithmically diminishing resolution at the periphery, the log-polar glimpses consistently maintain the focus on nearby regions regardless of the area captured in the full-resolution scene. This is visually apparent in the bottom row of Fig. 5, where log-polar glimpses remain visually similar regardless of their size (diameter), as opposed to crop-based glimpses in the top row. As a result, the fine search map extraction becomes more robust against the distracting peripheral information as indicated by the resulting fine glimpse locations in Fig. 5. Fig. 6 confirms this quantitatively, showing more stable performance over a broader range of glimpse sizes for log-polar compared to cropped-based glimpses. In addition, in Appendix D, we empirically justify our architecture for processing log-polar glimpses (Fig. 2B-C) by comparing it to a CNN-based model, showing that the latter is less effective.

Factoring out the downstream architecture. Instance detection can be decomposed into two stages: 1) finding a candidate region likely to contain the search target, and 2) matching it against the search target’s examples for recognition. Since CF-GAP primarily improves the first stage, we measure how often the search target is found in the scene regardless of whether it is successfully recognized during matching. We note that we cannot report more standard average recall metrics, as CF-GAP does not directly detect bounding boxes. Focusing on the most challenging HR-InsDet subsets, Tab. 4 shows that CF-GAP locates small objects and objects in hard scenes more frequently than standalone baselines. These results also reflect the upper-bound AP achievable under a perfect matching stage. The last row of Tab. 4 further highlights the importance of fine glimpsing for precise localization.

Table 4: Percentage of scenes in the two most challenging HR-InsDet subsets, where the search target was found, but not necessarily correctly recognized.

Model	Target found, %	
	small	hard
OTS-FM _{SAM}	42.1	56.3
NIDS-Net	73.3	75.2
CF-GAP	81.9	89.6
CF-GAP w/o fine glimpsing	69.5	74.6




Fig. 7: Coarse and fine glimpse locations for two scenes where baseline models failed to detect the search target (left: HR-InsDet, right: Robotools). Glimpse locations concentrate on a small fraction of the scene rather than spreading uniformly, and fine glimpse locations visibly correct the spatial imprecision of coarse ones.

Qualitative inspection. We provide visualizations of coarse and fine glimpse locations for a couple of scenes in Fig. 7. Note that the glimpses are not spread out across entire scenes meaning that only a subset of the entire scene will be passed to downstream architectures. This, in turn, implies the reduction of irrelevant and potentially distracting information. The visualizations in zoom-in panels also illustrate how fine glimpsing corrects for the spatial imprecision of coarse glimpsing. More examples are provided in Appendix E. Failed cases are shown and discussed in Appendix F.

5.3 Computational Cost

While CF-GAP significantly improves performance of the baselines, it incurs additional computational cost of repeatedly running the downstream architecture after each coarse glimpse. To fully leverage the power of CF-GAP, one has to select the downstream architecture wisely. In particular, it is costly to use a heavy, inefficient model designed and trained to handle complex scenery with numerous objects and various visual intricacies. In fact, this is unnecessary since CF-GAP provides targeted, high-quality information stripped of irrelevant details. Hence, a less powerful but more efficient downstream architecture can suffice. To illustrate this, we compare SAM with its two more efficient versions – MobileSAM and STT – as OTS-FM’s backbones combined with CF-GAP. The comparison is made both in terms of performance and efficiency, with the latter being represented via end-to-end FLOPS for all models. As can be seen in Fig. 8, CF-GAP




Fig. 8: CF-GAP with efficient OTS-FM_{SAM} variants evaluated on the two most challenging HR-InsDet subsets. Each dot corresponds to a specific number of coarse glimpses (1 to 30). Other baselines are omitted as no efficient variants are available.

makes the more efficient versions surpass the original OTS-FM_{SAM} model in terms of both efficiency and performance. Although, to the best of our knowledge, there are no efficient versions currently available for GroundingDINO, we expect to observe similar trends to those of SAM and its efficient alternatives.

More generally, with each coarse glimpse invoking the downstream architecture, CF-GAP directly trades efficiency for performance. Fig. 9 traces this trade-off along the number of coarse glimpses N_c . The comparison between the standalone baseline and its CF-GAP extension at the matched compute budget corresponds to the case of using a single coarse glimpse. In addition, we break




Fig. 9: Impact of the number of coarse glimpses on performance. Results are shown for CF-GAP+NIDS-Net evaluated on the two most challenging HR-InsDet subsets. The stars mark the performance of standalone NIDS-Net.

Table 5: Cost breakdown for CF-GAP paired with downstream architectures. N_c and N_f correspond to the number of coarse and fine glimpses, respectively.

Component	GFLOPs
Coarse search map, SM_c (MobileNet)	60 (per scene)
Fine search map, SM_f (Perceiver-based encoders)	1 (per glimpse)
CF-GAP internal cost: $SM_c + N_c \times N_f \times SM_f$	$60 + N_c \times N_f \times 1$
Downstream architecture (DA)	
{OTS _{STT} OTS _{MobileSAM} OTS _{SAM} NIDS-Net}	{80 130 5800 800}
Total cost: [CF-GAP internal cost] + $N_c \times$ [DA cost]	

down the compute cost in Tab. 5, showing that the total cost is dominated by N_c invocations of the downstream architecture. We report the cost in FLOPS, since wall-clock runtime depends on implementation-specific optimizations beyond the scope of this work. We also note that two further costs, shared by all downstream architectures, are omitted from the table: encoding the search target examples and the candidate objects with DINOv2, both of which vary across datasets and scenes.

Lastly, Fig. 9 shows the cumulative percentage of hitting the search target (in gray): in $\sim 50\%$ of scenes, CF-GAP finds search targets within the first 8-10 glimpses. This hints that the computational cost can be reduced, given a more robust matching stage that could terminate the glimpsing process once the search target is recognized.

6 Discussion

Our results demonstrate that CF-GAP consistently boosts existing instance detectors, with the largest gains in the most challenging settings of small-sized objects and cluttered scenes. Combining CF-GAP with instance detectors as downstream architectures exhibits a functional dichotomy of looking and seeing: CF-GAP *looks* for task-relevant regions and directs the downstream architecture as a *seeing* component to analyze them in high resolution. This division of labor, in turn, allows for employing lighter, distilled models such as OTS-FM_{MobileSAM} as a downstream architecture, achieving competitive performance compared to larger models. Moreover, one can also use large downstream architectures such as OTS-FM_{STT} with advanced fovea-inspired tokenization techniques that allow to retain their expressivity while making them very efficient. Hence, the combination of our lightweight CF-GAP with such downstream architectures paves the way to efficient yet powerful task-driven models.

Limitations and future work. In its current form, CF-GAP relies solely on texture-based guidance. However, the human visual system is known to leverage high-level semantics about objects, spatial layouts of scenes, and many other features when searching for task-relevant information. In addition, the inhibition-of-return mechanism only suppresses previously visited locations and their immediate neighborhoods, rather than entire task-irrelevant regions. This can cause the glimpsing process to repeatedly revisit the same distractor object. Another limitation is that CF-GAP passes a fixed-size crop as RoI to the downstream architecture, requiring the crop to be conservatively large to accommodate objects of varying sizes. An adaptive mechanism that adjusts the crop based on the content of the task-relevant region would improve both efficiency and precision. Finally, invoking the downstream architecture after every coarse glimpse incurs a computational cost that grows linearly with the number of coarse glimpses. A more robust matching stage that halts the glimpsing process once the search target is confidently recognized would alleviate this cost overhead. Addressing these limitations and extending CF-GAP to other task definitions, such as text-based queries, are promising directions for future work.

References

1. Akbas, E., Eckstein, M.P.: Object detection through search with a foveated visual system. *PLoS computational biology* **13**(10), e1005743 (2017)
2. Ba, J., Mnih, V., Kavukcuoglu, K.: Multiple object recognition with visual attention. *arXiv preprint arXiv:1412.7755* (2014)
3. Cheung, B., Weiss, E., Olshausen, B.: Emergence of foveal image sampling from learning to attend in visual scenes. *arXiv preprint arXiv:1611.09430* (2016)
4. Dai, J., Qi, H., Xiong, Y., Li, Y., Zhang, G., Hu, H., Wei, Y.: Deformable convolutional networks. In: *Proceedings of the IEEE international conference on computer vision*. pp. 764–773 (2017)
5. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An image is worth 16x16 words: Transformers for image recognition at scale. In: *International Conference on Learning Representations* (2021)
6. Dwivedi, D., Misra, I., Hebert, M.: Cut, paste and learn: Surprisingly easy synthesis for instance detection. In: *Proceedings of the IEEE international conference on computer vision*. pp. 1301–1310 (2017)
7. Ebel, P., Mishchuk, A., Yi, K.M., Fua, P., Trulls, E.: Beyond cartesian representations for local descriptors (2019)
8. Elsayed, G., Kornblith, S., Le, Q.V.: Saccader: Improving accuracy of hard attention models for vision. *Advances in neural information processing systems* **32** (2019)
9. Esteves, C., Allen-Blanchette, C., Zhou, X., Daniilidis, K.: Polar transformer networks. *arXiv preprint arXiv:1709.01889* (2017)
10. Freeman, J., Simoncelli, E.P.: Metamers of the ventral stream. *Nature neuroscience* **14**(9), 1195–1201 (2011)
11. Hayhoe, M., Ballard, D.: Eye movements in natural behavior. *Trends in cognitive sciences* **9**(4), 188–194 (2005)
12. Howard, A., Sandler, M., Chu, G., Chen, L.C., Chen, B., Tan, M., Wang, W., Zhu, Y., Pang, R., Vasudevan, V., et al.: Searching for mobilenetv3. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 1314–1324 (2019)
13. Ibrayev, T., Mukherjee, A., Aketi, S.A., Roy, K.: Toward two-stream foveation-based active vision learning. *IEEE Transactions on Cognitive and Developmental Systems* **16**(5), 1843–1860 (2024)
14. Ibrayev, T., Nagaraj, M., Mukherjee, A., Roy, K.: Exploring foveation and saccade for improved weakly-supervised localization. In: *Gaze Meets Machine Learning Workshop*. pp. 61–89. PMLR (2024)
15. Itti, L., Koch, C., Niebur, E.: A model of saliency-based visual attention for rapid scene analysis **20**(11), 1254–1259 (1998), conference Name: *IEEE Transactions on Pattern Analysis and Machine Intelligence*
16. Jaegle, A., Borgeaud, S., Alayrac, J.B., Doersch, C., Ionescu, C., Ding, D., Koppula, S., Zoran, D., Brock, A., Shelhamer, E., et al.: Perceiver io: A general architecture for structured inputs & outputs. *arXiv preprint arXiv:2107.14795* (2021)
17. Kim, D., Lin, T.Y., Angelova, A., Kweon, I.S., Kuo, W.: Learning open-world object proposals without learning to classify. *IEEE Robotics and Automation Letters* **7**(2), 5453–5460 (2022)
18. Kim, J., Jung, W., Kim, H., Lee, J.: Cycnn: A rotation invariant cnn using polar mapping and cylindrical convolution layers. *arXiv preprint arXiv:2007.10588* (2020)

19. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. In: Proceedings of the IEEE/CVF international conference on computer vision (2023)
20. Kolner, O., Ortner, T., Woźniak, S., Pantazi, A.: Mind the GAP: Glimpse-based active perception improves generalization and sample efficiency of visual reasoning. In: The Thirteenth International Conference on Learning Representations (2025)
21. Li, B., Wang, J., Hu, Y., Wang, C., Scherer, S.: Voxdet: Voxel learning for novel instance detection. *Advances in Neural Information Processing Systems* **36**, 10604–10621 (2023)
22. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: Proceedings of the IEEE international conference on computer vision. pp. 2980–2988 (2017)
23. Liu, S., Zeng, Z., Ren, T., Li, F., Zhang, H., Yang, J., Jiang, Q., Li, C., Yang, J., Su, H., et al.: Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In: European conference on computer vision. pp. 38–55. Springer (2024)
24. Lu, Y., Guo, Y., Ruozzi, N., Xiang, Y., et al.: Adapting pre-trained vision models for novel instance detection and segmentation. In: 2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 13341–13348. IEEE (2025)
25. McVitie, D.G., Wilson, L.B.: The stable marriage problem. *Communications of the ACM* **14**(7), 486–490 (1971)
26. Mercier, J.P., Garon, M., Giguere, P., Lalonde, J.F.: Deep template-based object instance detection. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 1507–1516 (2021)
27. Mnih, V., Heess, N., Graves, A., kavukcuoglu, k.: Recurrent Models of Visual Attention. In: *Advances in Neural Information Processing Systems*. vol. 27. Curran Associates, Inc. (2014)
28. Navon, D.: Forest before trees: The precedence of global features in visual perception. *Cognitive psychology* **9**(3), 353–383 (1977)
29. Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., et al.: Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193* (2023)
30. Osokin, A., Sumin, D., Lomakin, V.: Os2d: One-stage one-shot object detection by matching anchor features. In: European Conference on Computer Vision. pp. 635–652. Springer (2020)
31. Pardyl, A., Wronka, M., Wołczyk, M., Adamczewski, K., Trzciński, T., Zieliński, B.: Adaglimpse: Active visual exploration with arbitrary glimpse position and scale. In: European Conference on Computer Vision. pp. 112–129. Springer (2025)
32. Ravi, N., Gabeur, V., Hu, Y.T., Hu, R., Ryali, C., Ma, T., Khedr, H., Rädle, R., Rolland, C., Gustafson, L., et al.: Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714* (2024)
33. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems* **28** (2015)
34. Rolfs, M.: Microsaccades: small steps on a long way. *Vision research* **49**(20), 2415–2441 (2009)
35. Schmidt, T., Newcombe, R.: Segment this thing: Foveated tokenization for efficient point-prompted segmentation. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 29428–29437 (2025)

36. Schwartz, E.L.: Spatial mapping in the primate sensory projection: analytic structure and relevance to perception. *Biological cybernetics* **25**(4), 181–194 (1977)
37. Shen, Q., Zhao, Y., Kwon, N., Kim, J., Li, Y., Kong, S.: A high-resolution dataset for instance detection with multi-view object capture. *Advances in Neural Information Processing Systems* **36**, 42064–42076 (2023)
38. Shen, Q., Zhao, Y., Kwon, N., Kim, J., Li, Y., Kong, S.: Solving instance detection from an open-world perspective. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 9901–9910 (2025)
39. Tian, Z., Shen, C., Chen, H., He, T.: Fcos: Fully convolutional one-stage object detection. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 9627–9636 (2019)
40. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Advances in Neural Information Processing Systems* **30** (2017)
41. Wolfe, J.M.: Visual search. *Current biology* **20**(8), R346–R349 (2010)
42. Woźniak, S., Jónsson, H., Cherubini, G., Pantazi, A., Eleftheriou, E.: On the visual analytic intelligence of neural networks. *Nature Communications* **14**(1), 5978 (2023)
43. Wu, P., Xie, S.: V*: Guided visual search as a core mechanism in multimodal llms. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024)
44. Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A., Salakhudinov, R., Zemel, R., Bengio, Y.: Show, attend and tell: Neural image caption generation with visual attention. In: Bach, F., Blei, D. (eds.) *Proceedings of the 32nd International Conference on Machine Learning*. *Proceedings of Machine Learning Research*, vol. 37, pp. 2048–2057. PMLR, Lille, France (07–09 Jul 2015)
45. Yarbus, A.L.: *Eye Movements and Vision*. Springer US (1967)
46. Yu, X., Zhou, Z., Becker, S.I., Boettcher, S.E., Geng, J.J.: Good-enough attentional guidance. *Trends in Cognitive Sciences* **27**(4), 391–403 (2023)
47. Zeng, A., Florence, P., Tompson, J., Welker, S., Chien, J., Attarian, M., Armstrong, T., Krasin, I., Duong, D., Sindhvani, V., et al.: Transporter networks: Rearranging the visual world for robotic manipulation. In: *Conference on Robot Learning*. pp. 726–747. PMLR (2021)
48. Zhang, C., Han, D., Qiao, Y., Kim, J.U., Bae, S.H., Lee, S., Hong, C.S.: Faster segment anything: Towards lightweight sam for mobile applications. *arXiv preprint arXiv:2306.14289* (2023)
49. Zhang, H., Li, F., Liu, S., Zhang, L., Su, H., Zhu, J., Ni, L.M., Shum, H.Y.: Dino: Detr with improved denoising anchor boxes for end-to-end object detection. *arXiv preprint arXiv:2203.03605* (2022)
50. Zhang, J., Khayatkhoei, M., Chhikara, P., Ilievski, F.: MLLMs know where to look: Training-free perception of small visual details with multimodal LLMs. In: *The Thirteenth International Conference on Learning Representations* (2025)
51. Zhang, M., Feng, J., Ma, K.T., Lim, J.H., Zhao, Q., Kreiman, G.: Finding any waldo with zero-shot invariant and efficient visual search. *Nature Communications* **9**(1), 3730 (2018)
52. Zhou, X., Wang, D., Krähenbühl, P.: Objects as points. *arXiv preprint arXiv:1904.07850* (2019)