

Cost-Aware Fine-Grained Recognition for IoTs Based on Sequential Fixations

Hanxiao Wang, Venkatesh Saligrama, Stan Sclaroff, Vitaly Ablavsky
Boston University

{hwxw, srv, sclaroff, ablavsky}@bu.edu

Abstract

We consider the problem of fine-grained classification on an edge camera device that has limited power. The edge device must sparingly interact with the cloud to minimize communication bits to conserve power, and the cloud upon receiving the edge inputs returns a classification label. To deal with fine-grained classification, we adopt the perspective of sequential fixation with a foveated field-of-view to model cloud-edge interactions. We propose a novel deep reinforcement learning-based foveation model, DRIFT, that sequentially generates and recognizes mixed-acuity images. Training of DRIFT requires only image-level category labels and encourages fixations to contain task-relevant information, while maintaining data efficiency. Specifically, we train a foveation actor network with a novel Deep Deterministic Policy Gradient by Conditioned Critic and Coaching (DDPGC3) algorithm. In addition, we propose to shape the reward to provide informative feedback after each fixation to better guide RL training. We demonstrate the effectiveness of DRIFT on this task by evaluating on five fine-grained classification benchmark datasets, and show that the proposed approach achieves state-of-the-art performance with over 3X reduction in transmitted pixels.

1. Introduction

While edge camera IoT devices, which connect the physical world to the cloud, are revolutionizing data gathering in many consumer and business applications, their limited resources place a significant burden on their service life and operation, warranting cost-aware inference methods. For low-powered IoT devices, transmit energy dominates all other forms of battery usage [11, 17], and so, we are justified in considering the number of edge-cloud interactions as a surrogate for battery usage. Our goal is to minimize pixels transmitted while simultaneously ensuring accuracy that is comparable to a fully trained state-of-art deep neural network¹, which has access to the entire image. Our prototypical setup is an edge-device that sparingly transmits im-

¹The minimum amount of pixels for a standard Inception-V3 input size is about 299×299 to achieve good accuracy. The pictures captured by modern high resolution cameras are often larger than this size.

age regions to a cloud server equipped with abundant computational resources (e.g. a Inception-V3 network [37]), to interpret received inputs².

Fine-grained classification poses fundamental challenges in the IoT setup and highlights a fundamental dilemma between accuracy and cost: *On the one hand an instance cannot be classified accurately, unless the cloud sees the most discriminative parts. On the other hand an edge device can neither transmit the entire image due to bandwidth/power constraints, nor can it locally identify those parts due to lack of computational resources.*

In this context we are compelled to adopt a novel interactive edge-cloud model. Our approach is a novel deep reinforcement learning-based foveation model, DRIFT, that sequentially generates and recognizes mixed-acuity images. Instead of transmitting the full details all at once, the edge device first transmits a preliminary coarse thumbnail, e.g. 30×30 , to the server, which then actively but sparingly interacts with the edge device to seek image regions of value (see Fig.1), based only on the past received inputs. To summarize, our method (1) operates on mixed-acuity inputs whose resolution varies across the image, and more importantly, (2) actively and sequentially determines which image regions should be perceived with more visual details based on a low-acuity input, without access to the class label, a process we humans perform naturally³. While existing trained deep network architectures are an option, they are fundamentally ill-suited for this distributed task, and their variants, namely, attention-based methods require the entire high-acuity image to be available up-front.

Our DRIFT model consists of three neural networks: (1) a backbone CNN to extract visual features from low/mixed-acuity input images; (2) a foveation actor network to generate a sequence of fixation actions; and (3) an image classification network to predict the final class label. We pro-

²This edge-cloud setup is a conventional model in IoTs (see [31]).

³When we humans view a novel scene we do not perceive its full complexity at once, but rather, we foveate [26]. In doing so, our brain processes information from high-acuity foveal regions and the coarser-resolution periphery. The ability to process mixed-acuity inputs and actively ‘choose’ where to fixate significantly increases our efficiency, particularly because we do not need to observe every detail before recognizing a scene.

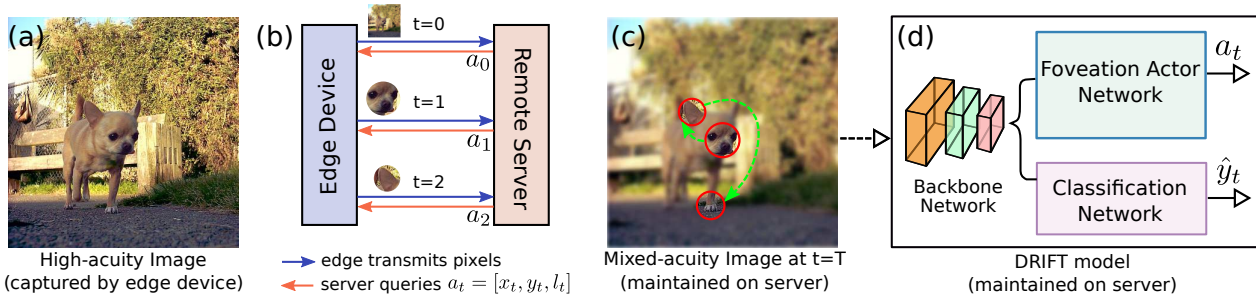


Figure 1. (a) The original high-acuity image captured by a low power edge device; (b) The IoT scenario: At $t = 0$ the edge device only transmits a thumbnail with extremely low resolution, e.g. 30×30 , to the server. The server then calculates the next fixation point (location and size) to query more high-acuity pixels from the edge device; (c) The mixed-acuity image which includes all the received pixels so far at $t = T$; (d) The proposed DRIFT model, described in Sec.3.

pose a novel reinforcement learning (RL) reward to guide the training so that the model fixates on regions that lead to high accuracy, while limiting the total transmitted high-acuity pixels. Given a low-acuity image transmitted from the remote edge device, a cloud server with DRIFT is able to predict locations of the most discriminative visual cues and actively query and incorporate such visual details by further interacting with the edge device.

Since the space for a fixation action (location and size) is large, discretizing/enumerating this space would be intractable in practice. Therefore, we propose solving in continuous space and train a foveation policy with a novel Deep Deterministic Policy Gradient by Conditioned Critic with Coaching (DDPGC3) algorithm. Compared to the original DDPG algorithm [20], several modifications are made: First, DDPG trains a critic to approximate an action-value function [43] to evaluate the learned policy, and uses the evaluation results to guide reinforcement learning. While this action-value function is globally shared among all state-action pairs in [20], we found this global function is difficult to approximate in our foveation problem. Instead, we propose training the critic to approximate a conditioned state-value function that is uniquely defined on every RL episode and more easy to approximate. Second, the actor network parameters in [20] are updated completely based on the critic's evaluation. However, at the early training stage a deficient critic can easily misguide the updates. Consequently, we propose updating the actor network by coaching [13], i.e. by both the critic's policy evaluation as well as the actions generated by a heuristic oracle. We observe that our improvements on DDPG stabilizes the training procedure.

Contributions: (1) An active image acuity exploration model, DRIFT, is proposed, for IoT applications which require efficient data transmission. DRIFT is able to sequentially infer fixation points from low-acuity images and classify based on mixed-acuity inputs; (2) A novel reward function is introduced so that the proposed DRIFT model can be trained with weak image-level class labels instead of more fine-grained labels on locations and sizes; (3) We pro-

pose training a foveation actor network via (I) a conditioned critic that approximates a unique state-value function conditioned on every input image, and (II) a coaching mechanism that combines the critic's evaluation with a heuristic oracle; (4) Experiments on five fine-grained classification datasets show that DRIFT achieves competitive performance with substantially fewer pixels compared to standard deep CNN models. (5) Furthermore, since DRIFT discovers discriminative visual features, it can also be used to generate hard attention which boosts standard classification performance for existing deep CNN models. Finally, although we demonstrate the proposed model's effectiveness in classification, DRIFT is a general active image acuity exploration solution which can be applied to other domains.

2. Related Work

While there are a number of works that deal with edge-computing, wireless sensor networks (see [46, 17, 31]) and resource constrained learning (see [29, 52, 4, 12, 38, 41]), the focus of these works tend to be one-shot, and interaction between cloud and edge device is not considered. DRIFT is the first to propose a sequential decision making process, whereby a cloud with computational resource interacts with a bandwidth limited edge device. This interaction is necessitated by fine-grained classification problem.

Recently, various attention models [45, 10, 18, 34, 24] have been proposed to enable CNNs to attend to specific image regions for multiple vision tasks. Nonetheless, these methods are optimized for foveated interaction leading to inefficient data transmission. In particular, they still operate on single high-acuity level inputs, and cannot sequentially infer attentions from low/mixed-acuity inputs. Consequently, all image details have to be revealed as inputs, a priori, instead of accumulation of a sequence of fixations.

Our research is related to foveation. Deng *et al.* [9] found that humans were able to correctly recognize an object by only revealing few high-acuity regions (referred to as fixations here) on a heavily blurred image. They thus propose crowd-sourcing to collect such location annotations

and training detectors on these discriminative features to boost classification accuracy. Matzen *et al.* [25] extended this concept and proposed an automatic but brute-force approach by initializing hundreds of random fixations per image, and iteratively optimizing to adjust each fixation based on the classification scores of their corresponding foveated images. These brute-force approaches require too many fixation regions to be transmitted per image, and are thus incompatible with IoT setup.

Different from [9, 25], which take low-acuity inputs, Almeida *et al.* [2] and Recasens *et al.* [33] proposed generating attention maps from images with standard input sizes, to either down-sample backgrounds [2] or up-sample foregrounds [33]. The approach of generating attention maps falls into a broader family of attention models, which has been predominantly applied to image classification [27, 3, 49, 10, 50], segmentation [18], visual question answering [34, 19], detection [48], image captioning [45], and so forth. As is the case for deep CNNs, these attention models require the full high-acuity images to be transmitted, and their performance degrades significantly on low-acuity inputs (see Sec. 4). In contrast our paper focuses on automatically inferring fixations from extremely low-acuity inputs (e.g. 30×30). We take a sequential and additive approach: The proposed DRIFT model is able to accumulate knowledge, recursively refine its fixations, and finally produce fixation locations that are optimized for classification accuracy as well as data efficiency. DRIFT learns to avoid exhaustive search, and thus is superior to the brute-force approach in [25].

In our RL formulation, fixations are modeled as a sequence of actions generated by a foveation actor model, which is similar in spirit to a few RL-based object detection works, e.g. [24, 32, 15, 6, 7]. However, our work is significantly different in that: (1) the proposed DRIFT learns fixation actions without any supervision on object locations, therefore it is more scalable to large scale data; (2) our action space is infinite, whereas in [32, 15, 6, 7] the actions can only be chosen within a restricted list, which limits diversity of model outputs; and (3) our low-acuity input uses much less information compared to the full-resolution input images in these detection methods.

3. Methodology

3.1. Foveation for Internet-of-Things

While our method is general, for concreteness we consider foveation within the context of image classification. We assume there are two types of representations for any scene/object: a low-acuity image I_{low} with limited visual details, and a high-acuity image I_{high} . For example, I_{high} could have standard Inception-V3 input size 299×299 , and I_{low} is a down-sampled version with size 30×30 . A low-battery edge device could directly transmit all the details on

I_{high} to a cloud server which leads to high classification accuracy, but to transmit I_{high} is expensive.

Our **foveation for IoT pipeline** is then defined as: (1) An edge device transmit I_{low} as an initial input to a cloud server. (2) A foveation model on the server infers a fixation point, which defines the coordinates and radius of a small circular image region. The coordinates and radius are then sent back to the edge device. (3) The edge device only transmits new high-acuity contents on I_{high} as specified by the fixation point. (4) The server incorporates the newly received high-acuity pixels to I_{low} , updates its posterior, and requests new coordinates if not confident. Intuitively, a good foveation model should reach a balance between *accuracy* and *transmission efficiency*, i.e. it should fixate on the most discriminative image regions so that good classification results can be achieved, while keeping the overall transmitted high-acuity pixels at a very low amount.

We view this foveation pipeline as an instance of Markov Decision Problem (MDP), and adopt reinforcement learning (RL) for training. We employ RL to train the foveation model because the optimal foveation policy should not be learned with any explicit supervision other than the objective of optimizing transmission efficiency and classification accuracy. It is thus difficult to define such a loss via standard supervised learning, but in RL, this training objective can be easily reflected by a reward function.

3.2. Markov Decision Process Formulation

We consider foveation for IoT as a sequential decision making problem, where the foveation model interacts with a dynamic environment E at discrete timesteps. At each timestep t , the foveation model receives an observation state s_t , takes an action a_t , and receives a scalar reward $r_t = r(s_t, a_t)$. This MDP process can be formally modeled by: action space \mathcal{A} , state space \mathcal{S} , transition dynamics from s_t to s_{t+1} after receiving a_t , and reward function $r(s_t, a_t)$. The foveation model implements a policy function π , which maps states to a distribution over actions: $\mathcal{S} \rightarrow \mathcal{P}(\mathcal{A})$. The return at timestep t is defined as the sum of discounted future rewards $R_t = \sum_{i=t}^{\infty} \gamma^{(i-t)} r(s_i, a_i)$ with a discount factor $\gamma \in [0, 1]$. Note that the return R_t depends on the actions taken, and thus depends on the policy π . The goal of RL is to find the best policy which maximizes the expected return $\mathbb{E}_{r_t, s_t \sim E, a_t \sim \pi} [\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t)]$. Next we explain in detail each component of our MDP.

Episode: To mimic the IoT scenario, we take images of the standard input size (e.g. 299×299 for Inception-V3) as I_{high} , and down-sample it as I_{low} (e.g. 30×30). I_{low} is thus left with very limited visual details (see Fig.5). At $t = 0$, the edge device transmits I_{low} . On receiving I_{low} , the server-side environment E interpolates it back to input size and uses it as the initial input for the foveation

by an actor network parameterized by w , which outputs a continuous deterministic policy $a = \pi_w(s)$. To optimize the policy, it takes a typical policy evaluation and improvement scheme. Policy evaluation uses state-value function $Q(s_t, a_t)$ to evaluate the current policy's expected return, where $Q(s_t, a_t) = \mathbb{E}_{r_t \geq t, s_t \geq t \sim E, a_t \geq t \sim \pi}(R_t | s_t, a_t)$. Here the state-value function is approximated by a critic network parameterized by θ , denoted as Q_θ , which only serves to train the actor network and is discarded during testing. Policy improvement uses the critic's estimation to improve the current policy model so that better $Q(s_t, a_t)$ can be reached.

Formally, the critic network is trained to optimize the temporal-difference (TD) term of the Bellman equation:

$$J_\theta = \min_{\theta} \mathbb{E}_{s_t, a_t, r_t \sim \beta} [(Q_\theta(s_t, a_t) - q_t)^2], \quad (5)$$

where $q_t = r_t + \gamma Q_{\theta'}(s_{t+1}, a_{t+1})$, β is the distribution of off-policy $(s_t, a_t, r_t, s_{t+1}, a_{t+1})$ samples stored in a replay buffer, $Q_{\theta'}$ is a separate target network used to generate TD targets q_t . The weights of the target network are updated by having them slowly track the learned networks: $\theta' = \tau\theta + (1 - \tau)\theta'$ with $\tau \ll 1$. Both the replay buffer and target network are originally introduced in [28] to decorelate training samples and stabilize the training process.

In [20] the objective for training the actor network is simply to maximize the critic's estimation:

$$J_w = \max_w \mathbb{E}_{s_t, a_t, r_t \sim \beta} [Q_\theta(s_t, \pi_w(s_t))]. \quad (6)$$

Conditioned Critic with Coaching We observe that DDPG fails to train a good foveation actor. Our analysis follows:

First, the global state-value function $Q(s_t, a_t)$ is too difficult to be approximated by the critic network $Q_\theta(s_t, a_t)$. Intuitively, given (s_t, a_t) , in the original formulation (Eq. 5), the critic network is expected to estimate R_t , which reflects the reward r_t . r_t depends on the ground-truth label y and g 's prediction \hat{y} , while \hat{y} further depends on the high-acuity region specified by a_t . Since the critic does not have access to any of $\{I_{high}, g, y, r_t\}$ by definition, the estimation of R_t is difficult. Observing this issue, we propose training a **conditioned critic** which approximates a unique state-value function defined for each episode k , $Q(s_t, a_t | C^k)$, where $C^k = [f(I_{high}^k), y^k]$ is the condition, I_{high}^k and y^k referring to the high-acuity image and ground-truth label for the k -th episode. By substituting the conditioned critic in Eq. 5; letting $q_t = r_t + \gamma Q_{\theta'}(s_{t+1}, a_{t+1} | C^k)$, and ψ as the distribution of episodes, we have a new objective for critic training:

$$J_\theta = \min_{\theta} \mathbb{E}_{s_t, a_t, r_t \sim \beta, C^k \sim \psi} [(Q_\theta(s_t, a_t | C^k) - q_t)^2], \quad (7)$$

Second, in Eq. 6 the actor network's optimization is solely based on the critic's estimation. Since the critic network parameter θ is randomly initialized, the critic's estimation is initially a random guess. This significantly slows down the training process and impedes convergence of actor training.

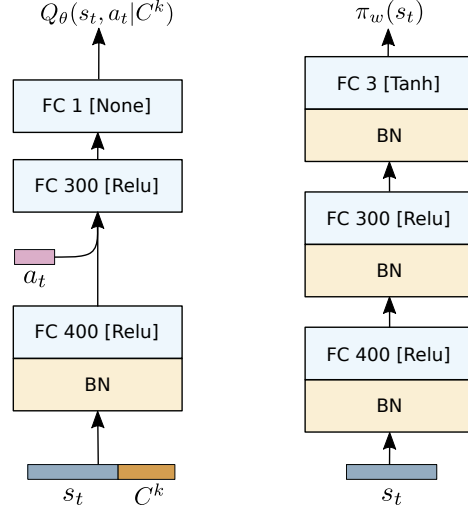


Figure 3. Our critic (left) and actor (right) network architecture. FC: fully-connected layer. BN: batch normalization layer. Numbers: the amount of neurons. Brackets: activation functions.

To solve this problem, we leverage the idea of coaching [13] and introduce a low-cost heuristic oracle⁵, as in [51], which provides a policy better than random guessing, and can be also used to guide early stage actor training. The actor training by **coaching** objective is defined as:

$$J_w = \max_w \mathbb{E}_{s_t, a_t, r_t, a'_t \sim \beta, C^k \sim \psi} [(1 - \epsilon)Q_\theta(s_t, \pi_w(s_t) | C^k) - \epsilon|\pi_w(s_t) - a'_t|^2], \quad (8)$$

where a'_t is the action taken by the heuristic oracle given s_t , and ϵ is an exponentially decreasing factor with respect to the training progress. We refer to this actor-critic RL training strategy with Eq. 7-8 as DDPG by Conditioned Critic with Coaching (DDPGC3). Using the final feature map prior to spatial pooling in f , we perform a 1×1 convolution with the ground-truth class's classifier to get a response map m (for Inception-V3: feature map is shaped $8 \times 8 \times 2048$ and m is 8×8). We then sample a location (x', y') based on m 's values, randomly sample a radius $l' \in [-1, 1]$, and use (x', y', l') to construct a'_t . Even though our naive a'_t only provides a coarse clue on the classifier's response over the low-acuity observation I_t , it still helps to speed up and stabilize the actor training by significantly saving efforts spent on random explorations caused by the deficient critic during early training. The oracle, which has access to the GT label, is only used during training and discarded when testing.

3.5. Implementation, Training and Deployment

Implementation We implemented our model using TensorFlow [1]. For the backbone f and classification network g , we adopted the Inception-V3 architecture [37], i.e. f outputs a 2048-d feature vector, and g is a fully-connected layer

⁵We use the term heuristic since a true oracle is impossible to realize without exhaustively searching over the large action space.

datasets (%)	CUB		Cars		Dogs		Aircrafts		Food101	
	acc	pix	acc	pix	acc	pix	acc	pix	acc	pix
Random	30.1	15.0	44.1	15.0	33.2	15.0	38.4	15.0	40.8	15.0
Center	59.3	15.0	55.3	15.0	58.4	15.0	69.6	15.0	42.3	15.0
Saliency	37.5	15.5	29.7	14.8	39.8	16.1	33.6	15.4	28.1	12.2
Attention	44.4	15.6	51.5	14.9	46.4	14.8	63.6	14.9	35.5	14.8
BubbleNet	65.5	-	-	-	-	-	-	-	56.1	-
DRIFT	74.4	10.1	82.8	11.5	71.6	14.1	86.7	14.4	75.5	11.4
I_{low}	13.9	1.0	7.8	1.0	17.1	1.0	6.6	1.0	8.9	1.0
I_{high}	81.6	100.0	91.2	100.0	81.8	100.0	87.2	100.0	85.0	100.0
DRIFT _E	80.1	32.6	88.5	33.6	78.0	35.6	88.0	35.8	81.9	33.6

Table 1. IoT classification setting. DRIFT outperforms other foveation methods, while requiring substantially fewer pixels to be transmitted.

datasets (%)	CUB	Cars	Dogs	Air	Food
Bilinear [21]	84.1	91.3	-	84.1	-
RA-CNN [10]	85.3	92.5	87.3	88.2	-
FCAN [22]	84.3	91.5	88.9	-	86.3
GP [44]	85.8	92.8	-	89.8	85.7
MAMC [36]	86.2	92.8	84.8	-	-
DFL-CNN [42]	87.4	93.1	-	91.7	-
Inception-V3	81.6	91.2	81.8	87.2	85.0
DRIFT _I	83.7	92.2	82.9	90.7	86.6
ResNet-50	83.2	92.2	85.7	89.8	85.8
DRIFT _R	86.2	93.6	87.3	91.7	88.6

Table 2. Standard classification setting. DRIFT consistently improves baselines’ performance.⁷

followed by a softmax. The architectures for our actor network π_w and critic network Q_θ are illustrated in Fig.3. For a given backbone with a default input size, e.g. 299×299 for Inception-V3, we define I_{high} to be the standard input image, and generate I_{low} by down-sampling I_{high} to 30×30 (only retaining 1% pixels). We set the smallest and largest fixation radius b_1, b_2 to 15 and 75, episode length T to 5, data-efficiency reward trade-off λ to 5.0, threshold thr to 25% of I_{high} pixels, discount factor γ to 0.9, target network update rate τ to $1e^{-4}$. As in [20], we add Ornstein-Uhlenbeck noise [39] to our actor policy for exploration.

Training We first pretrain $g \circ f$ on I_{high} with a standard classification loss. Then we train π_w and Q_θ by the proposed DDPGC3 algorithm (Sec. 3.4) for 60 epochs, with a SGD optimizer, a batch size of 32, and a fixed learning rate of $1e^{-4}$. For the beginning 50 epochs we freeze f and g , and then in the remaining epochs f and g are updated by a standard classification loss with the foveated images I_T as input. The size of experience replay buffer β for RL was 50,000. The decreasing factor ϵ for coaching in Eq. 8 is set to 0.7 initially and decays 0.96 every 1000 training updates. During training, $(s_t, a_t, a'_t, r_t, s_{t+1}, a_{t+1})$ samples are first pushed into the replay buffer, and then randomly sampled to update the actor and critic.

Deployment The critic network is deleted after training. The part $[f, \pi_w, g]$ is maintained at a cloud server. Once a low-acuity image is received, $\pi_w \circ f$ networks can be used to generate sequential fixation points, and $g \circ f$ to classify the resulting foveated images. Thus, both π_w and g take features generated by a shared backbone f . In real-world deployments where *server computation efficiency* is required, we see that we place no limitation on π_w to use cheap features while g may use expensive features so that a balance between data and computation efficiency might be reached.

4. Experiments

Our goal is the IoT setup consisting of poorly resourced edge camera device communicating data to a cloud-server

to perform fine-grained classification. To this end we show that DRIFT achieves state-of-art performance with significantly fewer high-acuity pixels relative to a fully trained DNN model that has access to full high-resolution images.

Experiments were conducted on five fine-grained classification datasets: CUB-200-2011 [40], Stanford Cars [16], Dogs [8], Aircrafts [23], and Food-101 [5]. We chose these datasets since the distinctions among categories are subtle and highly local, which requires a foveation model to fixate on the most discriminative regions to classify an image. The detailed statistics of these datasets are summarized in Table 3. Only the image-level category labels were used for training, while extra annotations such as bounding boxes and parts were NOT used.

datasets	CUB	Cars	Dogs	Air	Food
# Category	200	196	120	100	101
# Train	5,994	8,144	12,000	6,667	75,750
# Test	5,794	8,041	8,580	3,333	25,250

Table 3. The statistics of fine-grained datasets.

4.1. Internet-of-Things Setting

Setting We first compare the proposed DRIFT under the IoT scenario as described in Sec.3.1. Specifically, with the low-acuity images as initial inputs, we use DRIFT to generate fixation points under different foveation strategies, acquire more high-acuity pixels from the edge device, and finally use the trained $f \circ g$ to classify the foveated images. For this IoT setting, two criteria are considered: (1) the classification accuracy, and (2) the percentage of high-acuity pixels transmitted. A good foveation model should achieve high classification accuracy while requiring fewer high-acuity pixels.



Figure 4. Comparison on mixed-acuity foveated images.

⁷In both Table 1 and 2, blank marks (-) indicate unavailable results.




Figure 5. Qualitative result of the proposed DRIFT model. Each cell contains 4 images, from left to right: the (rescaled) low-acuity image I_{low} (input), the high-acuity image I_{high} , DRIFT’s foveated image I_T , and the zoomed high-acuity image by the tightest bounding box (shown in green) around the fixation points. On I_T , the fixation actions are also shown in red circles.

Baselines Seven IoT foveation strategies are compared: (1) *Random*: Fixate at random locations (uniform distribution); (2) *Center*: Fixate at the image center; (3) *Saliency*: Given an input image, we first obtain a class prediction \hat{y} , generate a class-response saliency map for \hat{y} following [51], and then sample a fixation location based on the map values. The procedure is repeated for T steps; (4) *Attention*: We trained a multi-attention model, MA-CNN [50], with T parts. Given an input image I_{low} , it generates T attention maps, from which T fixation locations are sampled. (5) *BubbleNet*: BubbleNet [25] initializes 128 fixation locations per image, iteratively optimizes each fixation and selects the best ones based on prediction entropy. We report its published results with the same Inception architecture; (6) *DRIFT*: Use the proposed model to generate sequential fixations; (7) *DRIFT_E*: In this strategy, if the prediction’s entropy on our I_T is higher than a threshold, we explore all high-acuity pixels in I_{high} instead. The threshold is controlled so that only 25% of the test images is used at full I_{high} . For (1-2), we control the fixation radius so that 15% of high-acuity pixels are explored for easy comparisons. For (3-4), we randomly sample the fixation radius.⁸

Results The results are shown in Table 1. We also provide the direct classification results using I_{low} and I_{high} as context. First, observe that DRIFT consistently outperforms the other five foveation strategies. While transmitting a similar number or fewer high-acuity pixels, DRIFT achieves much higher accuracy with our fixation approach. For example, on Aircrafts [23] we achieve 86.7% accuracy by only transmitting 14.4% of the high-acuity pixels, only 0.5% lower than the result with full high-acuity images (87.2%). Moreover, with DRIFT_E we are able to obtain an even higher accuracy (88.0%). This indicates that DRIFT’s fixations indeed contain the most discriminative regions, which can be validated by Fig.5. Taking a low-acuity image with limited information as input, it successfully fixates on objects of interest (e.g. the black dog), or the discriminative visual parts

of an object (e.g. the headlight and grille of the BMW).

Second, results in Table 1 indicate that approaches like Saliency [51] and Attention [50] fail to infer good fixations from low-acuity input (Fig.4), faring worse than Center fixations. This is because (a) they are not designed to operate on low/mixed-acuity inputs, and (b) they cannot accumulate prior fixations to inform future actions (DRIFT handles this via its proposed state representations and RL training guided by the dense rewards) Third, observe that on all datasets except for Food-101, Center fixation performs much better than Random fixation. This is because these datasets are artificially constructed by human with a center bias. For images in real-world deployments where the center prior no longer holds, we can expect a larger performance gap between DRIFT and Center fixations.

4.2. Standard Setting

Setting and Baselines Having validated DRIFT’s ability in optimizing data transmission for IoT, we ask if DRIFT’s fixation strategy is also beneficial to standard classification tasks. Specifically, as shown in Fig.5, we can fit a bounding box around DRIFT’s fixations. The box is similar to a hard attention, with the only difference that it is generated via a sequential foveation procedure from a low-acuity input image. In this setting, we simply treat DRIFT as a hard attention model, and verify whether it can boost classification results for any baseline classification model under the standard fine-grained classification setting.

Specifically, we use DRIFT’s hard attentions to zoom into the original images (Fig.5). Given a baseline model, we simply fuse its predictions on the original image and the zoomed image by DRIFT’s attention. We tested two baseline models: Inception-V3 [37] and ResNet-50 [14]. For all the five datasets, they are pre-trained on ImageNet [8], and further trained for 30 epochs with a RMSProp optimizer and a batch size of 32. The learning rate is initialized as 0.01 and decays 0.96 every 4 epochs. The input sizes are 299 and 448 for Inception-V3 and ResNet-50, respectively. We use DRIFT_I and DRIFT_R to represent the corresponding

⁸The proposed DRIFT model requires even less high-acuity pixels, so comparison is fair (see Table.1).




Figure 6. Visual patterns in fixation actions. Each cell contains four example fixation patches which belong to the same cluster.

datasets	CUB	Cars	Dogs	Air
Random	8.0	34.8	25.6	11.9
Center	36.2	89.1	61.2	32.9
DRIFT	44.3	91.5	63.2	50.9

Table 4. Localization results in hit rate.

classification results using our hard attentions together with Inception-V3 and ResNet-50 baseline models.

Results Table 2 shows our results. We observe a clear positive effect of DRIFT’s attention selection on classification accuracy. In particular, on average DRIFT_I is 1.9% higher than Inception-V3 in absolute accuracy, while DRIFT_R is 2.1% higher than ResNet-50, and has already achieved better or comparable performance to existing state-of-the-arts. This again demonstrates DRIFT’s ability to fixate on discriminative regions and filter out background clutter. Note that to compare purely in accuracy between Table.1 and 2 is not meaningful since DRIFT and DRIFT_E in Table.1 use far fewer pixels resulting in only around 10% and 30% pixel transmissions respectively while achieving accuracy very close to state-of-the-art as reported in Table.2.

4.3. Discussion and Analysis

Where does DRIFT fixate? First, inspired by [47] we use hit rate to evaluate the localization performance. Specifically, taking the boxes generated by DRIFT as in Sec. 4.2, we count a box as a hit when its intersection with the ground-truth box⁹ is greater than 90% of its own area, otherwise as a miss, and then measure $\frac{\#hits}{\#hits + \#misses}$. The localization performance is shown in Table 4. We also show results of a randomly-generated box and a center-located box of 1/2 image size. DRIFT’s localization performance is consistently superior. Evidently, DRIFT’s fixations are strongly correlated to object locations, even though trained without location labels.

Second, we aim to discover and visualize common patterns in DRIFT’s fixations to better understand its learned foveation policy. Specifically, we collect the local image patches specified by every fixation action, and perform a k-means clustering ($k = 50$) over their visual features. The clusters with top popularity are shown in Fig. 6. It is evident that DRIFT performs implicit part detection during fix-

acc (%)	DDPG	+ con. critic	+ coaching	DRIFT
CUB	51.0	57.1	67.0	74.4
Cars	48.3	61.4	76.6	82.8
Dogs	53.8	58.5	66.6	71.6

Table 5. Ablative analysis with results on foveated images.

ation. This experiment also shows the potential applications of DRIFT in visual discovery.

How much gain does ‘C3’ provide? While keeping all other settings fixed, we re-trained the foveation actor network π_w with three different strategies: DDPG, DDPG + Conditioned Critic, and DDPG + Coaching. Table 5 shows the classification results on their foveated images (detailed test setting in Sec.4.1). The original actor-critic training scheme as in DDPG [20] fails in our foveation problem, due to the reasons analyzed in Sec.3.4, i.e. a global state-value function difficult to approximate by the critic, and less informative guides provided by a randomly initialized critic. By conditioning the critic on every training episode, on average the accuracy is improved by 8.0% over the three datasets. Moreover, by coaching the actor using the policy sampled from a heuristic oracle that reduces exploration efforts, on average a 19.0% performance gain is obtained. Finally, the full DRIFT model, trained with the proposed DDPGC3 algorithm, brings a 25.2% average improvement in absolute accuracy; clearly DDPGC3 trains a better foveation policy.

5. Conclusion

We considered IoT scenarios where the cost of transmitting high-acuity images from an edge device to the cloud exceeds the transmission/power budget. Our solution is DRIFT, a novel deep-RL approach to generate sequential fixations with a foveated field-of-view. DRIFT avoids discretizing the state-action space, which would be prohibitively expensive, and instead solves a continuous-control problem. As part of our solution we introduce a novel use of a *conditioned critic* and a *coaching* strategy; we also provide an example of shaping the reward function to accelerate convergence. Experiments show high accuracy and data-efficiency of our approach on challenging classification tasks. Finally, although we demonstrated the proposed model’s effectiveness in classification, DRIFT is a general active image exploration solution which can be applied to other domains.

⁹CUB, Cars, Dogs and Aircrafts provide ground-truth bounding boxes.

References

- [1] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensorflow: a system for large-scale machine learning. In *OSDI*, volume 16, pages 265–283, 2016. 5
- [2] Ana Filipa Almeida, Rui Figueiredo, Alexandre Bernardino, and José Santos-Victor. Deep networks for human visual attention: A hybrid model using foveal vision. In *Iberian Robotics conference*, pages 117–128. Springer, 2017. 3
- [3] Jimmy Ba, Volodymyr Mnih, and Koray Kavukcuoglu. Multiple object recognition with visual attention. *arXiv preprint arXiv:1412.7755*, 2014. 3
- [4] Tolga Bolukbasi, Kai-Wei Chang, Joseph Wang, and Venkatesh Saligrama. Resource constrained structured prediction. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017. 2
- [5] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101 – mining discriminative components with random forests. In *European Conference on Computer Vision*, 2014. 6
- [6] Míriam Bellver Bueno, Xavier Giró-i Nieto, Ferran Marqués, and Jordi Torres. Hierarchical object detection with deep reinforcement learning. *Deep Learning for Image Processing Applications*, 31:164, 2017. 3
- [7] Juan C Caicedo and Svetlana Lazebnik. Active object localization with deep reinforcement learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2488–2496, 2015. 3
- [8] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*, 2009. 6, 7
- [9] Jia Deng, Jonathan Krause, and Li Fei-Fei. Fine-grained crowdsourcing for fine-grained recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 580–587, 2013. 2, 3
- [10] Jianlong Fu, Heliang Zheng, and Tao Mei. Look closer to see better: Recurrent attention convolutional neural network for fine-grained image recognition. In *CVPR*, volume 2, page 3, 2017. 2, 3, 6
- [11] Malka N Halgamuge, Moshe Zukerman, Kotagiri Ramamohanarao, and Hai Le Vu. An estimation of sensor energy consumption. *Progress In Electromagnetics Research B*, 2009. 1
- [12] Manjesh Hanawal, Csaba Szepesvari, and Venkatesh Saligrama. Unsupervised Sequential Sensor Acquisition. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, pages 803–811. PMLR, 2017. 2
- [13] He He, Jason Eisner, and Hal Daume. Imitation learning by coaching. In *Advances in Neural Information Processing Systems*, pages 3149–3157, 2012. 2, 5
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 7
- [15] Zequn Jie, Xiaodan Liang, Jiashi Feng, Xiaojie Jin, Wen Lu, and Shuicheng Yan. Tree-structured reinforcement learning for sequential object localization. In *Advances in Neural Information Processing Systems*, pages 127–135, 2016. 3
- [16] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *4th International IEEE Workshop on 3D Representation and Recognition (3dRR-13)*, Sydney, Australia, 2013. 6
- [17] Benoît Latré, Bart Braem, Ingrid Moerman, Chris Blondia, and Piet Demeester. A survey on wireless body area networks. *Wireless Networks.*, 2011. 1, 2
- [18] Kunpeng Li, Ziyang Wu, Kuan-Chuan Peng, Jan Ernst, and Yun Fu. Tell me where to look: Guided attention inference network. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 2, 3
- [19] Junwei Liang, Lu Jiang, Liangliang Cao, Li-Jia Li, and Alexander Hauptmann. Focal visual-text attention for visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6135–6143, 2018. 3
- [20] Timothy Paul Lillicrap, Jonathan James Hunt, Alexander Pritzel, Nicolas Manfred Otto Heess, Tom Erez, Yuval Tassa, David Silver, and Daniel Pieter Wierstra. Continuous control with deep reinforcement learning, Jan. 26 2017. US Patent App. 15/217,758. 2, 4, 5, 6, 8
- [21] Tsung-Yu Lin, Aruni RoyChowdhury, and Subhransu Maji. Bilinear cnn models for fine-grained visual recognition. In *Proceedings of the IEEE international conference on computer vision*, pages 1449–1457, 2015. 6
- [22] Xiao Liu, Tian Xia, Jiang Wang, Yi Yang, Feng Zhou, and Yuanqing Lin. Fully convolutional attention networks for fine-grained recognition. *arXiv preprint arXiv:1603.06765*, 2016. 6
- [23] S. Maji, J. Kannala, E. Rahtu, M. Blaschko, and A. Vedaldi. Fine-grained visual classification of aircraft. Technical report, 2013. 6, 7
- [24] Stefan Mathe, Aleksis Pirinen, and Cristian Sminchisescu. Reinforcement learning for visual object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2894–2902, 2016. 2, 3
- [25] Kevin Matzen and Noah Snavely. BubbleNet: Foveated imaging for visual discovery. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1931–1939, 2015. 3, 7
- [26] George W McConkie and Keith Rayner. The span of the effective stimulus during a fixation in reading. *Perception & Psychophysics*, 17(6):578–586, 1975. 1
- [27] Volodymyr Mnih, Nicolas Heess, Alex Graves, et al. Recurrent models of visual attention. In *Advances in neural information processing systems*, pages 2204–2212, 2014. 3
- [28] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529, 2015. 5
- [29] Feng Nan and Venkatesh Saligrama. Adaptive classification for prediction under a budget. In I. Guyon, U. V. Luxburg,

- S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 4727–4737. Curran Associates, Inc., 2017. 2
- [30] Deepak Pathak, Pulkit Agrawal, Alexei A Efros, and Trevor Darrell. Curiosity-driven exploration by self-supervised prediction. In *International Conference on Machine Learning (ICML)*, volume 2017, 2017. 4
- [31] Charith Perera, Chi Harold Liu, and Srimal Jayawardena. The emerging internet of things marketplace from an industrial perspective: A survey. *IEEE Transactions on Emerging Topics in Computing*, 3(4):585–598, 2015. 1, 2
- [32] Aleksis Pirinen and Cristian Sminchisescu. Deep reinforcement learning of region proposal networks for object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6945–6954, 2018. 3
- [33] Adria Recasens, Petr Kellnhofer, Simon Stent, Wojciech Matusik, and Antonio Torralba. Learning to zoom: a saliency-based sampling layer for neural networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 51–66, 2018. 3
- [34] Kevin J Shih, Saurabh Singh, and Derek Hoiem. Where to look: Focus regions for visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4613–4621, 2016. 2, 3
- [35] David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin Riedmiller. Deterministic policy gradient algorithms. In *ICML*, 2014. 4
- [36] Ming Sun, Yuchen Yuan, Feng Zhou, and Errui Ding. Multi-attention multi-class constraint for fine-grained image recognition. In *The European Conference on Computer Vision (ECCV)*, September 2018. 6
- [37] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016. 1, 5, 7
- [38] Kirill Trapeznikov and Venkatesh Saligrama. Supervised sequential classification under budget constraints. In *Artificial Intelligence and Statistics*, pages 581–589, 2013. 2
- [39] George E Uhlenbeck and Leonard S Ornstein. On the theory of the brownian motion. *Physical review*, 36(5):823, 1930. 6
- [40] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011. 6
- [41] Joseph Wang, Kirill Trapeznikov, and Venkatesh Saligrama. An lp for sequential learning under budgets. In *Artificial Intelligence and Statistics*, pages 987–995, 2014. 2
- [42] Yaming Wang, Vlad I Morariu, and Larry S Davis. Learning a discriminative filter bank within a cnn for fine-grained recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4148–4157, 2018. 6
- [43] Christopher JCH Watkins and Peter Dayan. Q-learning. *Machine learning*, 8(3-4):279–292, 1992. 2
- [44] Xing Wei, Yue Zhang, Yihong Gong, Jiawei Zhang, and Nanning Zheng. Grassmann pooling as compact homogeneous bilinear pooling for fine-grained visual classification. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 355–370, 2018. 6
- [45] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*, pages 2048–2057, 2015. 2, 3
- [46] A. Y. Yang, S. Iyengar, P. Kuryloski, and R. Jafari. Distributed segmentation and classification of human actions using a wearable motion sensor network. In *2008 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, pages 1–8, June 2008. 2
- [47] Jianming Zhang, Sarah Adel Bargal, Zhe Lin, Jonathan Brandt, Xiaohui Shen, and Stan Sclaroff. Top-down neural attention by excitation backprop. *International Journal of Computer Vision*, 126(10):1084–1102, 2018. 8
- [48] Xiaoning Zhang, Tiantian Wang, Jinqing Qi, Huchuan Lu, and Gang Wang. Progressive attention guided recurrent network for salient object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 714–722, 2018. 3
- [49] Bo Zhao, Xiao Wu, Jiashi Feng, Qiang Peng, and Shuicheng Yan. Diversified visual attention networks for fine-grained object classification. *IEEE Transactions on Multimedia*, 19(6):1245–1256, 2017. 3
- [50] Heliang Zheng, Jianlong Fu, Tao Mei, and Jiebo Luo. Learning multi-attention convolutional neural network for fine-grained image recognition. In *Int. Conf. on Computer Vision*, volume 6, 2017. 3, 7
- [51] Yanzhao Zhou, Yi Zhu, Qixiang Ye, Qiang Qiu, and Jianbin Jiao. Weakly supervised instance segmentation using class peak response. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 5, 7
- [52] Pengkai Zhu, Durmus Alp Emre Acar, Nan Feng, Prateek Jain, and Venkatesh Saligrama. Cost aware inference for iot devices. In *Proceedings of Machine Learning Research*, pages 2770–2779. PMLR, 2019. 2