

Variational Saccading: Efficient Inference for Large Resolution Images

Jason Ramapuram¹²
Jason.Ramapuram@etu.unige.ch

Maurits Diephuis¹
Maurits.Diephuis@unige.ch

Frantzeska Lavda¹²
Frantzeska.Lavda@etu.unige.ch

Russ Webb³
rwebb@apple.com

Alexandros Kalousis¹²
alexandros.kalousis@hesge.ch

¹ University of Geneva,
Switzerland

² Geneva School of Business
Administration, HES-SO,
Switzerland

³ Apple Inc.,
1 Infinite Loop,
Cupertino, CA

Abstract

Image classification with deep neural networks is typically restricted to images of small dimensionality such as $\mathbb{R}^{224 \times 224}$ in Resnet models [14]. This limitation excludes the $\mathbb{R}^{4000 \times 3000}$ dimensional images that are taken by modern smartphone cameras and smart devices. In this work, we aim to mitigate the prohibitive inferential and memory costs of operating in such large dimensional spaces. To sample from the high-resolution original input distribution, we propose using a smaller proxy distribution to learn the co-ordinates that correspond to regions of interest in the high-dimensional space. We introduce a new principled variational lower bound that captures the relationship of the proxy distribution’s posterior and the original image’s co-ordinate space in a way that maximizes the conditional classification likelihood. We empirically demonstrate on one synthetic benchmark and one real world large resolution DSLR camera image dataset that our method produces comparable results with $\sim 10\times$ faster inference and lower memory consumption than a model that utilizes the entire original input distribution. Finally, we experiment with a more complex setting using mini-maps from Starcraft II [6] to infer the number of characters in a complex 3d-rendered scene. Even in such complicated scenes our model provides strong localization: a feature missing from traditional classification models.

1 Introduction

Direct inference over large input spaces allows models to leverage fine grained information that might not be present in their downsampled counterparts. We demonstrate a simple example of such a scenario in Figure 1, where the task is to identify speed limits. The downsampled image does not contain the required information to correctly solve the task; on the other hand direct inference over the original input space is memory and computationally intensive.



Figure 1: (a) Original large resolution image. (b) Downsampled image processable (in a reasonable time-frame) by a typical Resnet model using existing computational resources.

In order to work over such large dimensional input spaces, we take inspiration from the way the human visual cortex handles high dimensional input. Research in neuroscience [23, 48, 54] and attention for eye-gaze [55] have suggested that human beings enact rapid eye movements (or saccades [44]) to different locations within the scene to gather high resolution information from local patches. More recent research [46, 47] has shown that humans and macaque monkeys stochastically sample saccades from their environment and merge them into a continuous representation of perception. These saccades are also not necessarily only of the salient object(s) in the environment, but have a component of randomness attached to them. In this work we try to parallel this stochastic element through the use of a learned sampling distribution, conditioned on auxiliary information provided via a proxy distribution. We explore two different types of proxy distributions in this work: (a) one where the proxy distribution is simply the downsampled version of the original, as in Figure 1 & (b) one where the proxy distribution is of a completely different modality. For the latter, we experiment with a Starcraft II scenario [56], and use the game minimap as the proxy distribution.

2 Related Work

Saliency Methods: The analysis of salient (or interesting) regions in images has been studied extensively in computer vision [21, 27, 29, 30]. Important regions are quantified by simple low-level features such as intensity, color and orientation changes. These methods fail to generalize to complex scenes with non-linear relationships between textures and colors [7]. More recently, deep convolutional networks have been exploited to directly learn saliency at multiple feature levels (eg: [8, 37]) as well as to learn patch level statistics [58]. None of these methods directly learn “where” to look without information about the entire image.

CNN Approaches: Current state of the art CNN models on the other hand separate the *entire* image, into cropped regions [60], employ pyramid decompositions [58] over the *entire* image, or utilize large pooling [9] / striding operands. These methods are challenging because they are either lossy, resulting in poor classification accuracy, or they are too memory and

computationally intensive (see Experiments Section 4) as they run convolutional filters over the entire image.

Region Proposal Methods: Another approach to CNN models are region proposal networks such as R-CNN [18], R-CNN++ [47] and YOLO [46] to name a few. The R-CNN methods generate a set of candidate extraction regions, either by extracting a fixed number of proposals as in the original work [18], or by utilizing a CNN over the *entire image* to directly predict the ROI [47]. They then proceed to enact a form of pooling over these regions, compute features, and project the features to the space of the classification likelihood. In contrast to R-CNN, our method uses an informatively learned posterior to extract the exact number of required proposals, rather than the 2000 proposals as suggested in the original work. R-CNN++ on the other hand doesn't scale with ultra-large dimensional images as direct inference over these images scales with the dimensionality of the images. Furthermore, the memory usage of R-CNN++ increases with the dimensionality of the images whereas it does not for our proposed model. YOLO on the other hand, resizes input images to $\mathbf{R}^{488 \times 488}$ and simultaneously predicts bounding boxes and their associated probabilities. While YOLO produces quick classification results, it trades off accuracy of fine-grained details. By resizing the original image, critical information can be lost (see Figure 1). Our proposed method on the other hand has no trouble with small details since it has the ability to directly control its foveation to sample the full resolution image.

Sequential Attention: Sequential attention models have been extensively explored through the literature, from utilizing Boltzman Machines [9, 13, 36], enacting step-by-step CNN learning rules [45], to learning scanning policies [10, 8] as well as leveraging regression based targets [25]. Our model takes inspiration from the recent Attend-Infer-Repeat (AIR) [15] and its extensions (SQAIR) [55], D.R.A.W [22], and Recurrent Attention Models (RAM) [0, 41]. While RAM based models allow for inference over large input images, they utilize a score function estimator [20] coupled with control variates [19]. Our algorithm on the other hand utilizes pathwise estimators [34, 59] which have been shown to have lower variance [63] in practice. In contrast to AIR and general attention based solutions, we do not use the entire image to build our attention map. Our model can infer where to attend using a summary with different semantics or encoding than the original image distribution (see Experiment 4.3). In addition, as opposed to adding a classifier in an ad-hoc manner as in AIR and SQAIR, we derive a new principled lower bound on the conditional classification likelihood that allows us to relate the posterior of the proxy-distribution to the co-ordinate space of the original input. This direct use of supervised information in an end-to-end manner allows our model to converge very rapidly (100-300 epochs) vs AIR which takes 50,000-200,000 epochs [15] to successfully converge.

Interpretability: With the surge of deep learning, understanding the model decision making process has become more important. While prior work took a post-mortem approach on trained models by computing gradients of the conditional likelihood with respect to the input image [59, 42, 61, 63], recent work such as Capsule Networks [49], InfoGAN [11], and numerous others [11, 49, 61, 63] directly attempt to learn models that are interpretable¹. Our model attempts to follow the latter of the two paradigms by extracting crops of regions that directly maximize the conditional classification likelihood. In contrast to the existing methods mentioned above we do not parse the entire input image to provide interpretability.

¹See [64] for a more thorough treatment of interpretability.

3 Variational Objective

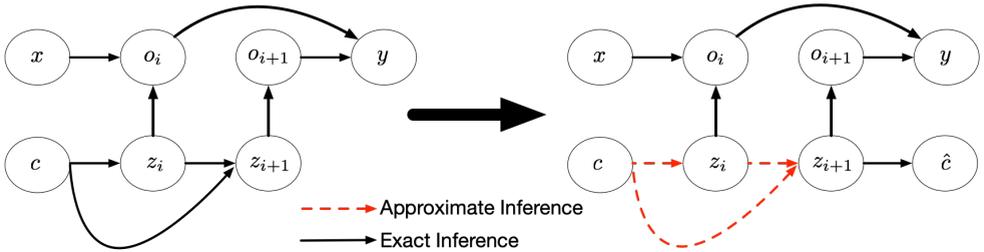


Figure 2: *Left*: Original (intractable) objective (Equation 1). *Right*: Newly derived tractable objective (Equation 5) that leverages a VRNN to approximate posteriors (red-dashed lines).

Given an image $\mathbf{x} \in \mathbb{R}^{K \times K}$, a corresponding proxy image $\mathbf{c} \in \mathbb{R}^{J \times J}$, $J \ll K$, and a corresponding class label $\mathbf{y} \in \mathbb{R}$, our objective is defined as maximizing $\log p_{\theta}(\mathbf{y}|\mathbf{x})$ for θ . We are only interested in the case where $p(\mathbf{y}|\mathbf{c}) \neq p(\mathbf{y}|\mathbf{x})$, i.e. the proxy distribution is not able to solve the classification task of interest. Assuming that \mathbf{c} provides no new information for the classification objective, $p_{\theta}(\mathbf{y}|\mathbf{x}) = p_{\theta}(\mathbf{y}|\mathbf{x}, \mathbf{c})$, we can reformulate our objective as:

$$\log p_{\theta}(\mathbf{y}|\mathbf{x}) = \log p_{\theta}(\mathbf{y}|\mathbf{x}, \mathbf{c}) = \log \int \int \frac{p_{\theta}(\mathbf{y}, \mathbf{o}_{\leq T}, \mathbf{z}_{\leq T}, \mathbf{c}, \mathbf{x})}{p(\mathbf{c}, \mathbf{x})} d\mathbf{z}_{\leq T} d\mathbf{o}_{\leq T} \quad (1)$$

We have introduced (and marginalized out) two sets of T latent variables in Equation (1): $\{\mathbf{z}\}_{i=1}^T, \mathbf{z}_i \in \mathbb{R}^3$ and $\{\mathbf{o}\}_{i=1}^T, \mathbf{o}_i \in \mathbb{R}^{L \times L}$, $L \ll J$. These correspond to the posteriors $\mathbf{z}_i \sim p_{\phi}(\mathbf{z}_i|\mathbf{c}, \mathbf{z}_{<i})$, induced by \mathbf{c} and a set of dirac distributions, $\mathbf{o}_i \sim \delta[f_{ST}(\mathbf{x}, \mathbf{z}_i)]$, centered at a differentiable function, f_{ST} , implemented using Spatial Transformer networks (ST) [R1]². This differentiable function produces crops, \mathbf{o}_i , of our large original input, \mathbf{x} , using a posterior sample from $p_{\phi}(\mathbf{z}_i|\mathbf{c}, \mathbf{z}_{<i})$. In general the true posterior, $p_{\phi}(\mathbf{z}_i|\mathbf{c}, \mathbf{z}_{<i})$, is intractable or difficult to approximate [R2]. To resolve this we posit a set of variational approximations [R3], $q_{\phi}(\mathbf{z}_i|\mathbf{c}, \mathbf{z}_{<i}) \approx p_{\phi}(\mathbf{z}_i|\mathbf{c}, \mathbf{z}_{<i})$, and introduce them via a multiply-by-one constant into the expanded joint distribution implied by the graphical model in Figure 2³:

$$\log p_{\theta}(\mathbf{y}|\mathbf{x}) = \log \int \int p_{\theta_y}(\mathbf{y}|\mathbf{o}_{\leq T}) p_{\theta_o}(\mathbf{o}_{\leq T}|\mathbf{z}_{\leq T}, \mathbf{x}) p_{\phi}(\mathbf{z}_{\leq T}|\mathbf{z}_{<T}, \mathbf{c}) \frac{q_{\phi}(\mathbf{z}_{\leq T}|\mathbf{z}_{<T}, \mathbf{c})}{q_{\phi}(\mathbf{z}_{\leq T}|\mathbf{z}_{<T}, \mathbf{c})} d\mathbf{z}_{\leq T} d\mathbf{o}_{\leq T} \quad (2)$$

By applying Jensen’s inequality and re-framing the marginalization operand as an expectation we can rewrite Equation 2 from above as a lower bound of $\log p_{\theta}(\mathbf{y}|\mathbf{x})$:

$$\log p_{\theta}(\mathbf{y}|\mathbf{x}) \geq \int \left[\mathbb{E}_{q_{\phi}} \left(\log \left[p_{\theta_y}(\mathbf{y}|\mathbf{o}_{\leq T}) p_{\theta_o}(\mathbf{o}_{\leq T}|\mathbf{z}_{\leq T}, \mathbf{x}) \right] \right) - D_{KL}[q_{\phi}(\mathbf{z}_{\leq T}|\mathbf{z}_{<T}, \mathbf{c}) || p_{\phi}(\mathbf{z}_{\leq T}|\mathbf{z}_{<T}, \mathbf{c})] \right] d\mathbf{o}_{\leq T} \quad (3)$$

²See Appendix Sections 7.2 and 7.3 for more information about Spatial Transformers and a possible method for minimizing their memory usage.

³Full derivation in Appendix Section 7.1.

We also observe that the KL divergence between the true set of posteriors $p_{\phi}(\mathbf{z}_{\leq T}|\mathbf{z}_{<T}, \mathbf{c})$ and the approximate posteriors $q_{\phi}(\mathbf{z}_{\leq T}|\mathbf{z}_{<T}, \mathbf{c})$ can be re-written in terms of the VRNN [16] Evidence Lower Bound (ELBO) [12, 24] and the marginal data distribution $p(\mathbf{c})$:

$$\begin{aligned} & -D_{KL}[q_{\phi}(\mathbf{z}_{\leq T}|\mathbf{c}, \mathbf{z}_{<T})||p_{\phi}(\mathbf{z}_{\leq T}|\mathbf{c}, \mathbf{z}_{<T})] \\ & = \\ & \mathbb{E}_{q_{\phi}}\left(\sum_{i=1}^T \log p_{\theta}(\hat{\mathbf{c}}|\mathbf{z}_{\leq i}) - D_{KL}(q_{\phi}(\mathbf{z}_i|\mathbf{c}, \mathbf{z}_{< i})||p_{\theta}(\mathbf{z}_i|\mathbf{c}, \mathbf{z}_{< i}))\right) - \log p(\mathbf{c}) \end{aligned} \quad (4)$$

$p_{\theta}(\mathbf{z}_i|\mathbf{c}, \mathbf{z}_{< i})$ in Equation 4 refers to the learned prior introduced by the VRNN, while $\hat{\mathbf{c}} \sim \log p_{\theta}(\hat{\mathbf{c}}|\mathbf{z}_{\leq i})$ is the VRNN reconstruction. All temporal dependencies $< T$ are incorporated by passing the hidden state of an RNN across functions (see [16] for a more thorough treatment). Coupling the VRNN ELBO with the fact that $-\log p(\mathbf{c})$ is always a positive constant, we can preserve the bound from Equation 3 and update our reframed objective as:

$$\begin{aligned} \log p_{\theta}(\mathbf{y}|\mathbf{x}) & \gtrsim \left[\mathbb{E}_{q_{\phi}}\left(\log p_{\theta, \mathbf{y}}(\mathbf{y}|f_{\theta_{\text{conv}}}(\underbrace{f_{ST}(\mathbf{z}_1, \mathbf{h}_1, \mathbf{x})}_{\mathbf{o}_1}, \dots, \underbrace{f_{ST}(\mathbf{z}_T, \mathbf{h}_T, \mathbf{x})}_{\mathbf{o}_T}))\right) \right. \\ & \quad + \\ & \quad \left. \underbrace{\mathbb{E}_{q_{\phi}}\left(\sum_{i=1}^T \log p_{\theta}(\hat{\mathbf{c}}|\mathbf{z}_{\leq i}) - D_{KL}(q_{\phi}(\mathbf{z}_i|\mathbf{c}, \mathbf{z}_{< i})||p(\mathbf{z}_i|\mathbf{c}, \mathbf{z}_{< i}))\right)}_{\text{VRNN ELBO}} \right] \end{aligned} \quad (5)$$

This leads us to our final optimization objective, shown above in Equation (5), which utilizes an empirical estimate of the expectation and marginalization operands and the substitutions of the functional approximations of the dirac distributions, $\mathbf{o}_i \sim \delta[f_{ST}(\mathbf{x}, \mathbf{z}_i)]$, to provide a novel lower bound on $\log p_{\theta}(\mathbf{y}|\mathbf{x})$. This lower bound allows us to classify a set of crops, $\{\mathbf{o}_i\}_{i=1}^T$, of the original distribution, $p(\mathbf{x})$, utilizing location information inferred by the posterior, $q_{\phi}(\mathbf{z}_{\leq T}|\mathbf{c}, \mathbf{z}_{< T})$, of the proxy distribution, $p(\mathbf{c})$.

3.1 Interpretation

Current state of the art research in neuroscience for attention [16, 26] suggest that humans sample saccades approximately every 250ms and integrate them into a continuous representation of perception. We parallel this within our model by utilizing a discrete $\mathbf{o}_i \sim \delta[f(\mathbf{x}, \mathbf{z}_i)]$ for sampling saccades and continuous latent representations (isotropic-gaussian posteriors) for the concept of perception. An additional requirement is the ability to transfer this continuous latent representation across saccades. This is enabled through the use of the VRNN, which relays information between posteriors through its RNN hidden state.

In addition, [16, 26] show that attention does not always focus on the most salient object in an image, but at times randomly attends to other parts of the scene. This behavior can be interpreted as a form of exploration as done in reinforcement learning. In our work, since the sampling distribution $q_{\phi}(\mathbf{z}_i|\mathbf{c}, \mathbf{z}_{< i})$ is stochastic, it provides a natural way to explore the space of the original input distribution $p(\mathbf{x})$, without the need for specific exploration methods such as ϵ -greedy [52] or weight noise [12]. We validate this in Experiment 4.1.1.1.

4 Experiments

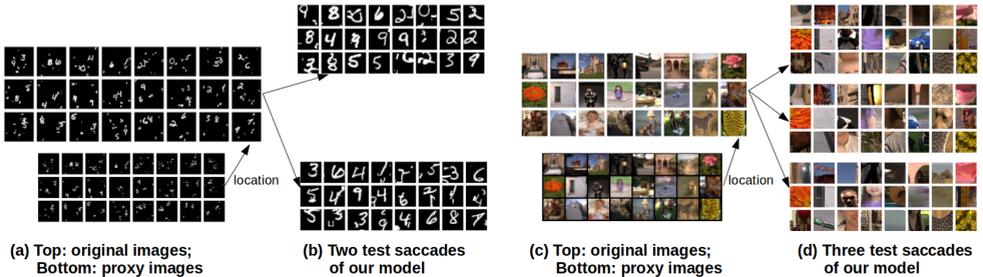


Figure 3: (a,b): Two-Digit-Identification ClutteredMNIST $\in \mathbb{R}^{2528 \times 2528}$; (c,d): MIT-5k $\in \mathbb{R}^{2528 \times 2528}$. Listed are test and proxy images (a,c) and their extracted test saccades (b,d).

| Image Size: $\mathbb{R}^{2528 \times 2528}$ | #params | gpu memory (160 batch-size) | time / epoch | accuracy MIT-Adobe-5k | accuracy Two Digit MNIST Identification |
|---|---------|---|---|--------------------------|--|
| resnet18 - 144 crops | 11.4M | 79G (naive) 6.5G (checkpoint) ⁴ | 1052.43s (naive) 1454.05s (checkpoint) | 63.6% +/- 0.03 | 97.3 +/- 0.006 |
| variational saccading | 7.4M | 4.1G | 120 s | 62.7% +/- 0.03 | 95.23 +/- 0.03 |

| Image Size: $\mathbb{R}^{100 \times 100}$ | #params | gpu memory (160 batch-size) | time / epoch | accuracy Two Digit MNIST Sum | accuracy Two Digit MNIST Identification |
|---|---------|--------------------------------|--------------|---------------------------------|--|
| resnet18 - full image | 11M | 6.6G | 59.27s | 99.86 +/- 0.01 | 97.4 +/- 0.003 |
| RAM [▣] | - | - | - | 91% | 93% |
| DRAM [▣] | - | - | - | 97.5% | 95% |
| variational saccading | 7.4M | 2.8G | 37s | 97.2 +/- 0.04 | 95.42 +/- 0.002 |

Table 1: Our model infers $\sim 9\text{-}10\times$ faster and utilizes less GPU memory than the baselines in high dimensions. *Top*: Large resolution trials on MIT-Adobe-5k and Two-Digit MNIST Identification. *Bottom*: Small resolution baseline trials used to situate work against RAM [▣] and DRAM[▣].

We evaluate our algorithm on three classification datasets where we analyze different induced behaviors of our model. We utilize Two-Digit MNIST for our first experiment in order to situate our model against baselines; we then proceed to learn a classification model for the large MIT-Adobe 5k dataset and finally attempt to learn a model that can accurately count marines in a complicated, dynamic, large resolution Starcraft II [56] map. The first two experiments utilize downsampled original images, $\mathbf{c} \in \mathbb{R}^{32 \times 32}$, as the proxy distribution, while the third uses the game-minimap from Starcraft II, $\mathbf{c} \in \mathbb{R}^{64 \times 64}$, as an auxiliary source of information. We demonstrate that our model has comparable accuracy to the best baseline models in the first two experiments, but we infer $\sim 9\text{-}10\times$ faster and utilize far less GPU memory than a naive approach. We provide visualizations of the model’s saccades in Figure 3; this aids in interpreting what region of the original input image aids the model in maximizing the desired classification likelihood. We utilize resnet18 as our naive baseline and did not observe any performance uplifts from using larger models for our three experiments.

⁴Checkpointing caches the forward pass operation as described in [▣]. The naive approach parallelizes across 8 GPUs and splits each of the 144 crops across the GPUs.

We implement the VRNN using a fully convolutional architecture where conv-transpose layers are used for upsampling from the vectorized latent space. The crop classifier is implemented by a standard fully-convolutional network, followed by a spatial pooling operation on the results of the convolution over the crops, \mathbf{o}_i . Adam [53] was used as an optimizer, combined with ReLU activations; batch-norm [28] was used for dense layers and group-norm [52] for convolutional layers. For more details about specific architectural choices see our code⁵ and Appendix Section 7.6 in the supplementary material.

4.1 Two-Digit MNIST: Identification & Summing

Two-Digit-Cluttered MNIST is a benchmark dataset used in RAM [41], DRAM [2] and as a generative target in AIR [15] and SQAIR [55]⁶. The objective with the initial set of experiments is to identify the digits present in the image (ignoring the distracting clutter), localize them, and predict a multi-class label using the localized targets. This form of learning, where localization information is not directly provided, is known as weakly supervised learning [1, 43, 44]. In the first setting we compare our model to RAM [41], DRAM [2] and a baseline resnet18 [24] model that operates over the entire image and directly provides classification outputs. As in RAM and DRAM, we also examine a case where the learning objective is to sum two digits placed in an image (without clutter). In order to provide a fair comparison we evaluate our model in the original dimension ($\mathbb{R}^{100 \times 100}$) suggested by the authors [2, 41]. We observe (Table 1 bottom) that our method improves upon RAM and DRAM and gets close to the baseline resnet18 results.

We extend the Two-Digit-Cluttered MNIST identification experiment from above to a new experiment where we classify large dimensional images, $\mathbf{x} \sim \mathbb{R}^{2528 \times 2528}$. As in the previous experiment we evaluate our model against a baseline resnet18 model. Resnet models are tailored to operate over $\mathbb{R}^{224 \times 224}$ images; in order to use large images, we divide an original $\mathbb{R}^{2528 \times 2528}$ image into $\mathbb{R}^{144 \times 224 \times 224}$ individual crops and feed each crop into the model. We then sum the logit outputs of the model and run the pooled result through a dense layer. This allows the model to make a single classification decision for the entire image using all 144 crops: $\mathbf{y} = f_{\theta_d}(\sum_{i=1}^{144} g_{\theta_c}(\mathbf{x}_i))$, $\mathbf{x}_i \in \mathbb{R}^{224 \times 224}$. f_{θ_d} represents a multi-layer dense network and g_{θ_c} is a multi-layer convolutional neural network that operates on individual crops \mathbf{x}_i . While it is also possible to concatenate each logit vector $g_{\theta_c}(\mathbf{x}) = [g_{\theta_c}(\mathbf{x}_i), g_{\theta_c}(\mathbf{x}_{i-1}), \dots, g_{\theta_c}(\mathbf{x}_0)]$, and project it through the dense network $f_{\theta_d}(g_{\theta_c}(\mathbf{x}))$, the tasks we operate over do not necessitate relational information [50] and pooled results directly aid the classification objective. We visualize saccades (Figure 3), the model accuracy, training-time per epoch and GPU memory (Table 1) and observe that our model performs similarly (in terms of accuracy) in higher dimensions, while inferring $\sim 10\mathbf{x}$ faster and using **only 5%** of the total GPU memory in contrast to a traditional resnet18 model.

⁵https://github.com/jramapuram/variational_saccading

⁶The authors do not use the cluttered version of the two-digit dataset for the AIR variants.

4.1.1 Ablation Studies

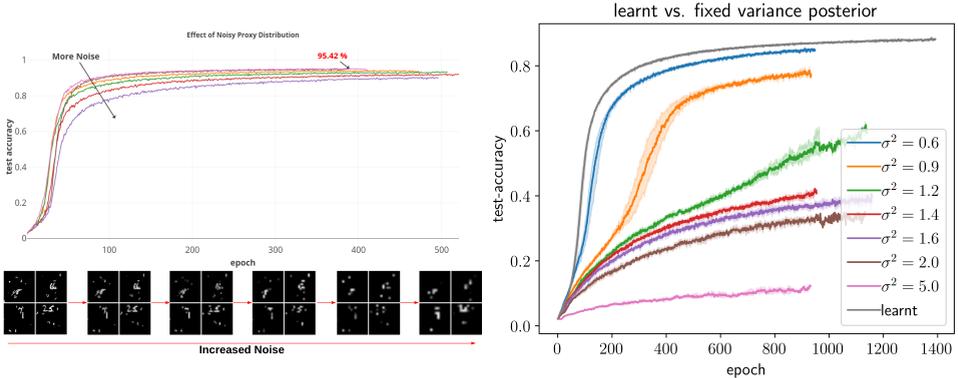


Figure 4: *Top-left*: Effect of noisy proxy distribution on test accuracy. *Bottom-left*: Left to right correspond to noisier versions of the same proxy distribution used in above graph. *Right*: Test accuracy for Two-Digit ClutteredMNIST identification under a range of fixed variances, σ^2 , of the isotropic-gaussian posterior $q_{\phi}(\mathbf{z}_i | \mathbf{z}_{<i}, \mathbf{c})$.

Robustness to Noisy Proxy Distribution: Since the proxy distribution is critical to our formulation, we conduct an ablation study using the two-digit cluttered identification problem from experiment 4.1. We vary the amount of noise in the proxy distribution as shown in the bottom of Figure 4-left. The test curves shown on the top of the same figure demonstrates that our method is robust to noisy proxy distributions. In general, we found that our method worked even in situations where the proxy distribution only contained a few points, allowing us to infer positional information to index the original distribution $p(\mathbf{x})$.

Quantifying Learned Exploration: In order to validate the hypothesis that the learned variance, σ^2 , of our isotropic-gaussian posterior $q_{\phi}(\mathbf{z}_i | \mathbf{z}_{<i}, \mathbf{c})$ is useful in the learning process, we repeat the identification experiment from Experiment 4.1 using the noisiest proxy-distribution from the previous ablation study (right most example in Figure 4-bottom-left). We compare our model (*learnt*) against the same model with varying fixed variance: $\sigma^2 = [0.6\text{I}, 0.9\text{I}, 1.2\text{I}, 1.4\text{I}, 1.6\text{I}, 2.0\text{I}, 5.0\text{I}]$. We repeat each experiment five times and plot the mean and variance of the test accuracy in Figure 4-right. We observe a **clear** advantage in terms of convergence time and accuracy for the *learnt* model.

4.2 MIT-Adobe 5k

MIT-Adobe 5k [9] is a high resolution DSLR camera dataset consisting of six classes: {abstract, animals, man-made, nature, None, people}. While the dimensionality of each image is large, the dataset has a total of 5000 samples. This upper-bounds the performance of deep models with millions of parameters (without the use of pre-training / fine-tuning and other unsupervised techniques). We examine this scenario because it presents a common use case of learning in a low-sample regime.

We downsample the large original images to $\mathbf{x} \in \mathbb{R}^{3 \times 2528 \times 2528}$ to evaluate against a baseline resnet [24] model. The baseline model operates over 144 crops per image as in the

previous experiment. Test saccades (non-cherry picked) of our model are visualized in Figure 3; the saccades allow us to gain an introspective view into the model decision making process. Some of the interesting examples are that of the ‘people’ class: in the example with the child (third to the right in the bottom row of Figure 3-c), the model saccades to the adult as well as the child in the image. Other notable examples are leveraging the spotted texture of cheetah fur and the snout of the dog. As observable from Table 1-top, our model has comparable accuracy to the baseline resnet model, but infers $\sim 9\text{-}10\times$ faster and uses far less GPU memory than a naive approach.

4.3 Starcraft II - Count the Marines

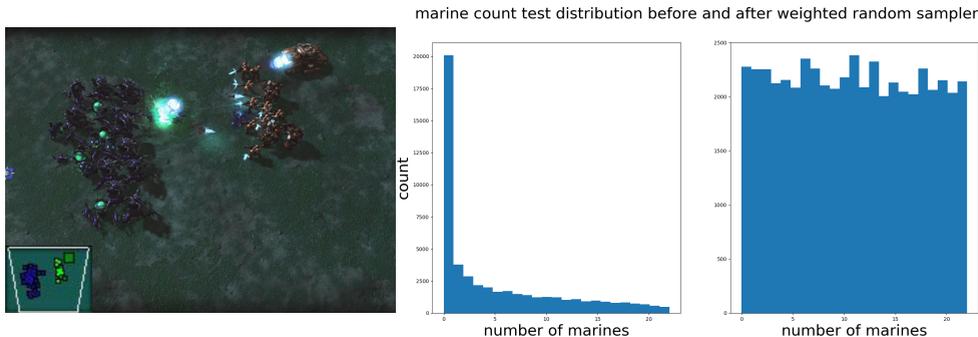


Figure 5: *Left*: full resolution image with minimap (proxy distribution, $p(\mathbf{c})$) in the bottom left corner. *Right*: initial test output distribution and normalized test distribution using weighted random sampler with replacement.

In this experiment we generate a dataset of 250,000 samples of a randomly initialized Starcraft II [56] scenario. The maps generated use the full resolution ($\mathbf{x} \in \mathbb{R}^{1024 \times 1024}$) images which have complex structure such as bloom, background texture and anti-aliased sprites (see Figure 5 left). Units from the Zerg faction attack Terran units until all units of one faction are eliminated. Both factions generate a random number of three distinct types of units with a maximum of 22 per unit-type. Our objective in this scenario is to predict the number of Terran marines in the map at any given time given training data with the correct counts; we use the minimap (Figure 5-left bottom corner) as our proxy distribution, $p(\mathbf{c})$, to infer positional information in the full-resolution image $p(\mathbf{x})$. Even though the minimap has an estimate of the total count the model cannot directly use this information to make a prediction due to the fact that *there are multiple unit types with the same minimap grid size*. This forces the model to have to look at the true map in order to infer the number of marines present. In addition, the minimap is only used to infer locations in the full-resolution images and not directly for classification.

The generated dataset was initially biased to a count of zero marines; in order to learn a model that did not instantly have $> 50\%$ accuracy we used a weighted random sampler (with replacement) on both the train and test sets. We visualize the effect of this sampler on the entire test distribution in the histograms shown in Figure 5-right.

As with the previous experiment we compare our model to a baseline resnet18 that utilized the entire image to predict the Terran marine count. In this experiment we observe that our model (50%) does not perform as well as the baseline resnet18 (65%) in terms of classi-

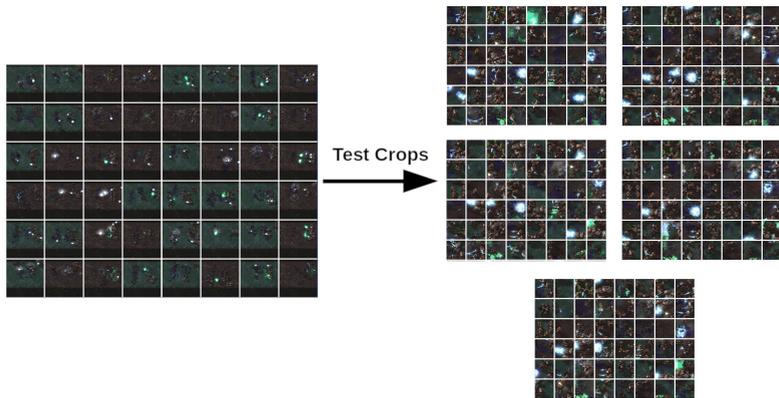


Figure 6: *Left*: Starcraft II full resolution test images $\in \mathbb{R}^{1024 \times 1024}$. *Right*: Five Test saccades (per input image) extracted by our model; almost all contain localized Terran marines.

fication accuracy, however we observe reasonable localizations in Figure 6. We hypothesize that this is due to the fact that we use a completely different modality for the proxy distribution. This makes it challenging for the model to learn how to saccade around the space as it needs to first learn the relationship of the simplistic color RGB mini-map and the complicated, texture filled, full-resolution map. We hypothesize that the reason for the performance degradation is due to double-counting of marines. Our work presents a strong first attempt at probabilistic large image classification using location information derived from a proxy image. The method also provides localization for free as a byproduct of the process.

5 Conclusion

We demonstrate a novel algorithm capable of working with ultra-large resolution images for classification and derive a new principled variational lower bound that captures the relationship of a proxy distribution’s posterior and the original image’s co-ordinate space. We empirically demonstrate that our model works with low memory and inference costs on ultra-large images in three datasets.

6 Acknowledgments

We would like to thank Amazon Web Services (AWS) and the Swiss National Supercomputing Centre (CSCS) for their generous GPU grants, without which the experiments in this paper would not be possible. We would also like to thank our hard-working systems administrator, Yann Sagon, from the University of Geneva, who keeps our internal cluster operational.

References

- [1] Bogdan Alexe, Nicolas Heess, Yee W Teh, and Vittorio Ferrari. Searching for objects driven by context. In *Advances in Neural Information Processing Systems*, pages 881–889, 2012.

- [2] Jimmy Ba, Volodymyr Mnih, and Koray Kavukcuoglu. Multiple object recognition with visual attention. 2015.
- [3] Loris Bazzani, Hugo Larochelle, Vittorio Murino, Jo-anne Ting, and Nando D Freitas. Learning attentional policies for tracking and recognition in video with deep networks. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pages 937–944, 2011.
- [4] Y-Lan Boureau, Jean Ponce, and Yann LeCun. A theoretical analysis of feature pooling in visual recognition. In *Proceedings of the 27th international conference on machine learning (ICML-10)*, pages 111–118, 2010.
- [5] Nicholas J Butko and Javier R Movellan. Optimal scanning for faster object detection. In *Computer vision and pattern recognition, 2009. cvpr 2009. ieee conference on*, pages 2751–2758. IEEE, 2009.
- [6] Vladimir Bychkovsky, Sylvain Paris, Eric Chan, and Frédo Durand. Learning photographic global tonal adjustment with a database of input / output image pairs. In *The Twenty-Fourth IEEE Conference on Computer Vision and Pattern Recognition*, 2011.
- [7] Zoya Bylinskii, Adrià Recasens, Ali Borji, Aude Oliva, Antonio Torralba, and Frédo Durand. Where should saliency models look next? In *European Conference on Computer Vision*, pages 809–824. Springer, 2016.
- [8] Feilong Cao, Yuehua Liu, and Dianhui Wang. Efficient saliency detection using convolutional neural networks with feature selection. *Information Sciences*, 456:34–49, 2018.
- [9] Gustavo Carneiro, Antoni B Chan, Pedro J Moreno, and Nuno Vasconcelos. Supervised learning of semantic classes for image annotation and retrieval. *IEEE transactions on pattern analysis and machine intelligence*, 29(3):394–410, 2007.
- [10] Tianqi Chen, Bing Xu, Chiyuan Zhang, and Carlos Guestrin. Training deep nets with sublinear memory cost. *arXiv preprint arXiv:1604.06174*, 2016.
- [11] Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in neural information processing systems*, pages 2172–2180, 2016.
- [12] Junyoung Chung, Kyle Kastner, Laurent Dinh, Kratarth Goel, Aaron C Courville, and Yoshua Bengio. A recurrent latent variable model for sequential data. In *Advances in neural information processing systems*, pages 2980–2988, 2015.
- [13] Misha Denil, Loris Bazzani, Hugo Larochelle, and Nando de Freitas. Learning where to attend with deep architectures for image tracking. *Neural computation*, 24(8):2151–2184, 2012.
- [14] Heiner Deubel and Werner X Schneider. Saccade target selection and object recognition: Evidence for a common attentional mechanism. *Vision research*, 36(12):1827–1837, 1996.

- [15] SM Ali Eslami, Nicolas Heess, Theophane Weber, Yuval Tassa, David Szepesvari, Geoffrey E Hinton, et al. Attend, infer, repeat: Fast scene understanding with generative models. In *Advances in Neural Information Processing Systems*, pages 3225–3233, 2016.
- [16] Ian C Fiebelkorn, Mark A Pinsk, and Sabine Kastner. A dynamic interplay within the frontoparietal network underlies rhythmic spatial attention. *Neuron*, 99(4):842–853, 2018.
- [17] Meire Fortunato, Mohammad Gheshlaghi Azar, Bilal Piot, Jacob Menick, Matteo Hessel, Ian Osband, Alex Graves, Volodymyr Mnih, Remi Munos, Demis Hassabis, et al. Noisy networks for exploration. *ICLR*, 2018.
- [18] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 580–587, 2014.
- [19] Paul Glasserman. *Monte Carlo methods in financial engineering*, volume 53. Springer Science & Business Media, 2013.
- [20] Peter W Glynn. Likelihood ratio gradient estimation for stochastic systems. *Communications of the ACM*, 33(10):75–84, 1990.
- [21] Stas Goferman, Lihi Zelnik-Manor, and Ayellet Tal. Context-aware saliency detection. *IEEE transactions on pattern analysis and machine intelligence*, 34(10):1915–1926, 2012.
- [22] Karol Gregor, Ivo Danihelka, Alex Graves, Danilo Rezende, and Daan Wierstra. Draw: A recurrent neural network for image generation. In *International Conference on Machine Learning*, pages 1462–1471, 2015.
- [23] Mary Hayhoe and Dana Ballard. Eye movements in natural behavior. *Trends in cognitive sciences*, 9(4):188–194, 2005.
- [24] 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.
- [25] David Held, Sebastian Thrun, and Silvio Savarese. Learning to track at 100 fps with deep regression networks. In *European Conference on Computer Vision*, pages 749–765. Springer, 2016.
- [26] Randolph F Helfrich, Ian C Fiebelkorn, Sara M Szczepanski, Jack J Lin, Josef Parvizi, Robert T Knight, and Sabine Kastner. Neural mechanisms of sustained attention are rhythmic. *Neuron*, 99(4):854–865, 2018.
- [27] Xiaodi Hou and Liqing Zhang. Saliency detection: A spectral residual approach. In *Computer Vision and Pattern Recognition, 2007. CVPR'07. IEEE Conference on*, pages 1–8. IEEE, 2007.
- [28] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.

- [29] Laurent Itti and Christof Koch. A saliency-based search mechanism for overt and covert shifts of visual attention. *Vision research*, 40(10-12):1489–1506, 2000.
- [30] Laurent Itti, Christof Koch, and Ernst Niebur. A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on pattern analysis and machine intelligence*, 20(11):1254–1259, 1998.
- [31] Max Jaderberg, Karen Simonyan, Andrew Zisserman, et al. Spatial transformer networks. In *Advances in neural information processing systems*, pages 2017–2025, 2015.
- [32] Diederik P Kingma. "*Variational Inference & Deep Learning: A New Synthesis*". PhD thesis, 2017.
- [33] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*, 2014.
- [34] Durk P. Kingma and Max Welling. Auto-encoding variational bayes. *ICLR*, 2014.
- [35] Adam Kosiorok, Alex Bewley, and Ingmar Posner. Hierarchical attentive recurrent tracking. In *Advances in Neural Information Processing Systems*, pages 3053–3061, 2017.
- [36] Hugo Larochelle and Geoffrey E Hinton. Learning to combine foveal glimpses with a third-order boltzmann machine. In *Advances in neural information processing systems*, pages 1243–1251, 2010.
- [37] Guanbin Li and Yizhou Yu. Visual saliency based on multiscale deep features. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5455–5463, 2015.
- [38] Emmanuel Maggiori, Yuliya Tarabalka, Guillaume Charpiat, and Pierre Alliez. High-resolution image classification with convolutional networks. In *Geoscience and Remote Sensing Symposium (IGARSS), 2017 IEEE International*, pages 5157–5160. IEEE, 2017.
- [39] Aravindh Mahendran and Andrea Vedaldi. Understanding deep image representations by inverting them. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5188–5196, 2015.
- [40] Kirk Martinez and John Cupitt. Vips-a highly tuned image processing software architecture. In *Image Processing, 2005. ICIP 2005. IEEE International Conference on*, volume 2, pages II–574. IEEE, 2005.
- [41] Volodymyr Mnih, Nicolas Heess, Alex Graves, et al. Recurrent models of visual attention. In *Advances in neural information processing systems*, pages 2204–2212, 2014.
- [42] Chris Olah, Alexander Mordvintsev, and Ludwig Schubert. Feature visualization. *Distill*, 2(11):e7, 2017.
- [43] Maxime Oquab, Léon Bottou, Ivan Laptev, and Josef Sivic. Is object localization for free?-weakly-supervised learning with convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 685–694, 2015.

- [44] Jason Ramapuram and Russ Webb. A new benchmark and progress toward improved weakly supervised learning. *BMVC*, 2018.
- [45] Marc’Aurelio Ranzato. On learning where to look. *arXiv preprint arXiv:1405.5488*, 2014.
- [46] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016.
- [47] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015.
- [48] Ronald A Rensink. The dynamic representation of scenes. *Visual cognition*, 7(1-3): 17–42, 2000.
- [49] Sara Sabour, Nicholas Frosst, and Geoffrey E Hinton. Dynamic routing between capsules. In *Advances in Neural Information Processing Systems*, pages 3856–3866, 2017.
- [50] Adam Santoro, David Raposo, David G Barrett, Mateusz Malinowski, Razvan Pascanu, Peter Battaglia, and Tim Lillicrap. A simple neural network module for relational reasoning. In *Advances in neural information processing systems*, pages 4967–4976, 2017.
- [51] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*, 2013.
- [52] Richard S Sutton and Andrew G Barto. *Introduction to reinforcement learning*, volume 135. MIT press Cambridge, 1998.
- [53] Seiya Tokui and Issei Sato. Evaluating the variance of likelihood-ratio gradient estimators. In *International Conference on Machine Learning*, pages 3414–3423, 2017.
- [54] Anne M Treisman and Garry Gelade. A feature-integration theory of attention. *Cognitive psychology*, 12(1):97–136, 1980.
- [55] Sabine Kastner Ungerleider and Leslie G. Mechanisms of visual attention in the human cortex. *Annual review of neuroscience*, 23(1):315–341, 2000.
- [56] Oriol Vinyals, Timo Ewalds, Sergey Bartunov, Petko Georgiev, Alexander Sasha Vezhnevets, Michelle Yeo, Alireza Makhzani, Heinrich Küttler, John Agapiou, Julian Schrittwieser, et al. Starcraft ii: A new challenge for reinforcement learning. *arXiv preprint arXiv:1708.04782*, 2017.
- [57] Martin J Wainwright, Michael I Jordan, et al. Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning*, 1(1–2): 1–305, 2008.
- [58] Lijun Wang, Huchuan Lu, Xiang Ruan, and Ming-Hsuan Yang. Deep networks for saliency detection via local estimation and global search. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3183–3192, 2015.

- [59] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.
- [60] Sai Wu, Mengdan Zhang, Gang Chen, and Ke Chen. A new approach to compute cnns for extremely large images. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, pages 39–48. ACM, 2017.
- [61] Tianfu Wu, Xilai Li, Xi Song, Wei Sun, Liang Dong, and Bo Li. Interpretable r-cnn. *arXiv preprint arXiv:1711.05226*, 2017.
- [62] Yuxin Wu and Kaiming He. Group normalization. *arXiv preprint arXiv:1803.08494*, 2018.
- [63] Matthew D Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *European conference on computer vision*, pages 818–833. Springer, 2014.
- [64] Quan-shi Zhang and Song-Chun Zhu. Visual interpretability for deep learning: a survey. *Frontiers of Information Technology & Electronic Engineering*, 19(1):27–39, 2018.
- [65] Quanshi Zhang, Ruiming Cao, Ying Nian Wu, and Song-Chun Zhu. Growing interpretable part graphs on convnets via multi-shot learning. In *AAAI*, pages 2898–2906, 2017.