

Tidying Deep Saliency Prediction Architectures

Navyasri Reddy^{*1}, Samyak Jain^{*1}, Pradeep Yarlagadda^{*1} and Vineet Gandhi¹

Abstract—Learning computational models for visual attention (saliency estimation) is an effort to inch machines/robots closer to human visual cognitive abilities. Data-driven efforts have dominated the landscape since the introduction of deep neural network architectures. In deep learning research, the choices in architecture design are often empirical and frequently lead to more complex models than necessary. The complexity, in turn, hinders the application requirements. In this paper, we identify four key components of saliency models, i.e., input features, multi-level integration, readout architecture, and loss functions. We review the existing state of the art models on these four components and propose novel and simpler alternatives. As a result, we propose two novel end-to-end architectures called SimpleNet and MDNSal, which are neater, minimal, more interpretable and achieve state of the art performance on public saliency benchmarks. SimpleNet is an optimized encoder-decoder architecture and brings notable performance gains on the SALICON dataset (the largest saliency benchmark). MDNSal is a parametric model that directly predicts parameters of a GMM distribution and is aimed to bring more interpretability to the prediction maps. The proposed saliency models can be inferred at 25fps, making them suitable for real-time applications. Code and pre-trained models are available at <https://github.com/samyak0210/saliency>.

I. INTRODUCTION

Predicting the salient regions in a scene is a fundamental ability, which empowers primates to rapidly analyze/interpret the complex surroundings by locating and devoting the focus only on sub-regions of interest [1]. The work by [2] triggered early interest in the computational modeling of visual saliency from images, i.e., identifying areas that are salient in a scene. Since then, a large variety of saliency detection models have been proposed and find usages in a wide range of applications involving machine vision. Many recent works show that availability of saliency maps enhance cognitive abilities of robots and helps improving performance in variety of tasks including human-robot interaction [3]; identification, and recognition of objects [4]; scene classification [5]; detecting and tracking regions of interest [6]; proposal refinement [7] and visual search in unknown environments (allowing search of regions with higher importance first) [8]. Our work is application agnostic and focuses on improving the general saliency prediction and can cater to large variety of applications in robotic vision. Some example results from our SimpleNet model are illustrated in Figure 1.

Formally, computational saliency models predict the probability distribution of the location of the eye fixations over the image, i.e., the saliency map. Where human observers look in images is often used as a ground truth estimate of

image saliency. The predictions are evaluated using a variety of metrics, which are broadly classified as location-based or distribution-based [9]. The location-based metrics measure the accuracy of saliency models at predicting discrete fixation locations. Distribution based metrics compute the difference/similarity between predicted and ground truth distributions (assuming that the ground truth fixation locations are sampled from an underlying probability distribution).

The last few years have seen tremendous advancements in the field, mainly due to the application of Deep Neural Network architectures for the task and the availability of large scale datasets [10]. Recent works have analyzed the saliency estimation models over different evaluation metrics to add interpretability to saliency scores [9]. Interestingly, the interpretability of saliency model architectures has not been systematically explored. To this end, we propose a componential analysis which can be used to compare a model from another; reduce redundancies in the model without compromising the performance and can help customizations based on application requirements.

We identify four key components in saliency models. First is the input features i.e., to directly send the image to the saliency models or employ transfer learning using pre-trained networks. The second component is the multi-level integration. It is understood that multiscale features (at different spatial and semantic hierarchy) capture a broad spectrum of stimuli, and the combination improves model performance. This aspect concerns how the hierarchy is imbibed in the model. The third aspect is the readout architecture, which concerns the form of output i.e., to directly predict a saliency map or to predict parameters of an assumed underlying distribution. The fourth aspect is the loss function. Different works use a different combination of loss functions; however, most of these choices are only justified empirically. We explore ways to validate these combinations more formally. Overall, our work makes the following contributions:

- We separate components of saliency models and discuss the progress on each of them in reference to the literature. Such analysis can help better interpret the models i.e., assess component-wise weaknesses, strengths, and novelty. The analysis allows to optimize saliency models by trying alternates for a particular component while freezing the rest of them.
- We propose an encoder-decoder based saliency detection model called SimpleNet. The main novelty of SimpleNet is a UNet like multi-level integration [11]. SimpleNet is fully convolutional; end to end trainable; has lower complexity than counterparts and allows real-time inference. It gives consistent performance over

^{*} Equal Contribution

¹ CVIT, KCIS, International Institute of Information Technology Hyderabad, India navyasri.reddy@research.iiit.ac.in

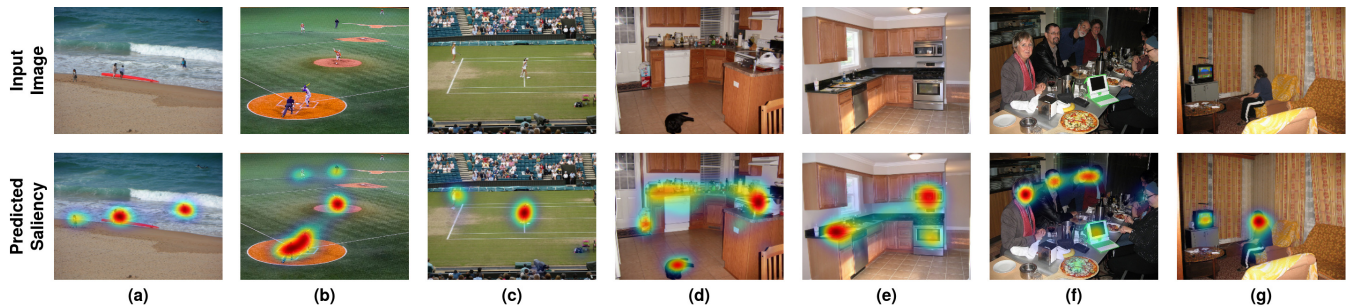


Fig. 1. Example results of our approach on images from Salicon dataset. Saliency models can play key role in application like (a) drone surveillance, (b,c) robotics cameras for sports, (d,e) indoor navigation and (f,g) social interactions.

multiple metrics on SALICON and MIT benchmarks, outperforming state of the art models over five different metrics (with significantly notable improvements on KLdiv metric).

- We propose a parametric model called MDNSal, which predicts parameters of a GMM instead of a pixel-level saliency map. The main novelty of MDNSal is in readout architecture with a modified Negative Log Likelihood (NLL) loss formulation. It achieves near state of the art performance on SALICON and MIT benchmarks.

II. KEY COMPONENTS OF SALIENCY MODELS

A. Input features

Early attempts relied on handcrafted low-level features for saliency prediction. Seminal work by Itti [2] relied on color, intensity, and orientation maps (obtained using Gabor filter). Valenti [12] use isophotes (lines connecting points of equal intensity), color gradients, and curvature features. Zhang [13] computes saliency maps by analyzing the topological structure of Boolean maps generated through random sampling. Bruce [14] use low-level local features (patch level) in combination with information-theoretic ideas. Jude [15] included high-level information by using detectors for faces, people, cars, and horizon. However, most of these methods remain elusive on generic high-level feature representation.

Recent works are dominated by deep learning architectures owing to their strong performance. Most of this success can be attributed to Convolutional architectures [16], [17], [18], [19]. Some works have also explored combining CNN with recurrent architectures [20]. The breakthrough happened via transfer learning of high-level features trained for image classification [21], [22], [23], [24], [25]. The large scale SALICON dataset [10] was pivotal in transfer learning process (allowing efficient fine-tuning). Initial approaches relied on Alexnet or VGG features [16], [18], [26]. Other notable architectures like ResNet, DenseNet and NasNet [20], [19] were then explored. Recent works also explore the combinations of features from multiple pre-trained networks [19]. There is enough evidence to agree that using pre-trained features brings significant gains on the task of saliency prediction. We analyze two important design choices: (a)

which pre-trained network to pick and (b) should pre-trained weights be frozen or fine-tuned.

B. Multi-level integration

It is evident that deep learning models utilizing high-level features significantly outperform the older counterparts, which rely on low-level handcrafted features. However, recent work [27] suggests that explicitly modelling low-level contributions to saliency could be used to improve the robustness of saliency models. They quantitatively show this by changing the input features to low-level intensity contrast features (ICF) and keeping the rest of the architecture the same.

Deep networks have employed two main strategies to resolve this concern. The first is to send different image scales as input in parallel. SALICON [28] model uses two different image scales and the idea was then extended to multiple scales [29], [30]. The multi scale spatial stimuli can be tackled by this approach, but not necessarily the semantics. An alternate approach is to take features from different stages of pre-trained CNN. Different levels of semantics (from low, mid, and high-level features) can be thus directly incorporated into the model to resolve the concerns raised in [27]. The work by [16] takes a weighted sum of features at different levels post resizing, where the weights are trained through the network. ML-Net [18] resizes and concatenates features from different levels of VGG-16 model and passes it through additional convolutions layers to predict the map. The work in [31] individually predicts saliency maps for features from different stages of VGG-19 and then fuses them. In this work, we propose a novel UNet [11] like architecture for incorporating multi-level features. The single-stream network with skip connections speeds up training, and the structure allows for an organic hierarchical refinement from high to low-level features (symmetric expansion over high-level context to enable precise localization).

C. Readout Architectures

The commonly used readout architecture consists of few convolutional layers post the encoder followed by 1×1 convolutions to control the size of the output saliency map. It is also common to learn an additional prior [16], [18], [27]. The prior is often aimed to compensate for




Fig. 2. SimpleNet Architecture

the central fixation bias. Work by [20] employs an lstm based readout architecture and learns a set of 2D Gaussian priors parameterized by their mean and variance (instead of a single one). In this paper, we employ minimal readout architecture with only convolutional layers. We show that combined with an UNet like multi-resolution encoder; such architecture can outperform state of the art models which rely on priors, complex architectures [20] or multi-network feature combinations [19].

Interestingly, most of the state of the models directly predict an image as output (the saliency map). Parametric models for computing image saliency have not been explored. Although parametric models come with a bound on the complexity of the model (even if the amount of data is unbounded), they come with several advantages; especially, they are easier to understand and interpret (Where saliency models should look next? [32]). Furthermore, they allow better integration with downstream applications. The importance of predicting distributions has been nicely motivated in [33]. To this end, we propose a novel readout architecture, which directly predicts parameters of a 2D GMM (mean, variance, and mixture weights). The proposed readout architecture can be plugged at the end of any given architecture to output a parametric distribution. We show that, although bounded, the parametric models can achieve near state of the art performance.

D. Loss functions

Mean Squared Error (MSE) between predicted and ground truth has been employed as loss function [17]. ML-Net introduced a normalized version of MSE [18]. Most of the recent efforts directly use one of commonly used evaluation metrics or a combination of them as a loss function. The most commonly used loss is computing KL-divergence (KLdiv) between the estimated and ground-truth saliency maps [28]. Some papers use a variation of it like negative log likelihood [16] or cross entropy [31] instead. Recent works use KLdiv in combination with other metrics like Pearson's Correlation Coefficient (CC), Normalized Scanpath Saliency(NSS), and Similarity. These combinations bring clear improvement in performance [19], [20]. However, the combinations are often decided empirically. In this work, we

provide formal insights to choose a minimal and comprehensive loss function.

III. PROPOSED ARCHITECTURES

We propose two end-to-end architectures SimpleNet and MDNSal. SimpleNet is an encoder-decoder architecture that predicts the pixel-wise saliency values, while MDNSal is a parametric model that predicts parameters of a GMM distribution. We now describe each of the models in detail.

A. SimpleNet

The overall architecture of the SimpleNet is shown in Figure 2. It is a fully convolutional, single-stream encoder-decoder architecture, which is end to end trainable. The name SimpleNet is derived from the design goal to keep each component simple and minimal, resulting in an architecture which easy to train, comprehend and reproduce, without compromising the performance.

a) Input features: SimpleNet directly takes input from the pre-trained architectures designed for image classification. We explore four different architectures VGG-16, ResNet-50, DenseNet-161, and PNASNet-5 and compare their performance. The feature extraction layers (shown as encoder blocks in Figure 2) are initialized with the pre-trained weights and later fine-tuned for saliency prediction.

b) Multi-level integration: SimpleNet employs a UNet like architecture that symmetrically expands the input features starting at the final layer of the encoder (input features). The symmetric expansion enables precise localization. Every step of the expansion consists of an upsampling of the feature map, a concatenation with the corresponding scale feature map from the encoder. The number of channels are then reduced using 3×3 convolutions followed by ReLU.

c) Readout architecture: The readout architecture consists of two 3×3 convolutional layers; the first is followed by ReLU, and the second uses a sigmoid to output the saliency map.

d) Loss function: The loss function compares the output saliency map with the ground truth. We use a combination of Kullback-Leibler Divergence(KLdiv) and Pearson Cross Correlation (CC) metrics as a loss function. KLdiv is an information-theoretic measure of the difference between

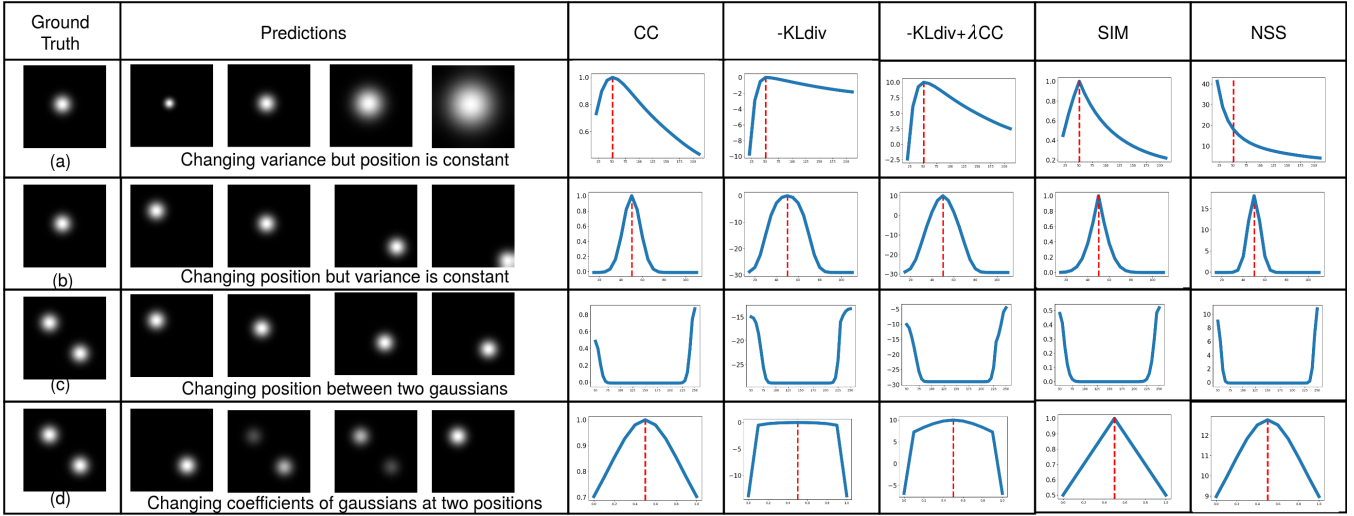


Fig. 3. We synthetically vary saliency predictions w.r.t the ground truth in order to quantify effects on the loss functions. Each row corresponds to varying a single parameter value of the prediction: (a) Variance, (b) location on a single mode, (c) location between two modes, (d) relative weights between two modes. The x and y-axis spans the parameter range and values of loss functions respectively. The dotted red line corresponds to the ground truth.

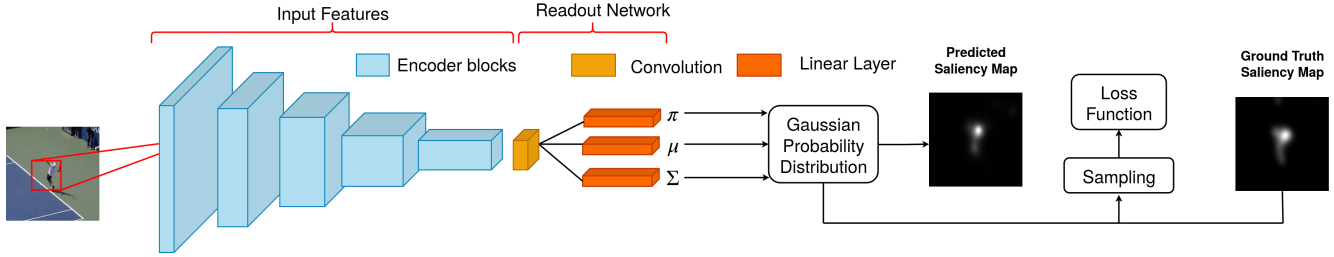


Fig. 4. MDNSal Architecture

two probability distributions:

$$KLdiv(P, Q) = \sum_i Q_i \log\left(\epsilon + \frac{Q_i}{P_i + \epsilon}\right), \quad (1)$$

where P, Q are predicted and ground truth maps respectively and ϵ is a regularization term.

CC is a statistical method used generally in the sciences for measuring how correlated or dependent two variables are

$$CC(P, Q) = \frac{\sigma(P, Q)}{\sigma(P) \times \sigma(Q)}. \quad (2)$$

The combination of KLdiv and CC is motivated by the analysis presented in Figure 3 (the analysis is inspired by [9]). KLdiv is highly sensitive to false negatives and results in steeper costs (consider case (a), (b) in Figure 3). Steeper costs lead to larger gradients, which are crucial in initial training (which motivates the use of KLdiv or its variants as a backbone of loss function). CC is symmetric in terms of false positives and false negatives and gives typical behavior in each scenario of Figure 3. The combination provides appropriate behavior in each of the studied scenarios while maintaining steeper costs (scenarios (a), (d) (Figure 3)). We also explore the use of Normalized Scanpath Saliency (NSS) as a loss function in the ablation studies.

B. MDNSal

The overall architecture of the MDNSal is shown in Figure 4. The network is inspired by the literature on Mixture Density Networks [34]. MDNSal is a parametric model which gives a compressed representation, allowing faster processing at the application level.

a) Input features: Similar to SimpleNet; we apply transfer learning from pre-trained image classification networks. The fine-tuning of the pre-trained features is extremely crucial in MDNSal and leads to significant performance improvements.

b) Multi-level integration: MDNSal only uses the features from the last convolutional layer of the pre-trained networks and is devoid of multi-level integration. Since the outputs are parameters instead of a per-pixel map, the multi-level features do not play a significant role.

c) Readout architecture: The readout architecture consists of a convolutional layer to reduce the number of channels followed by a ReLU. The output is then passed to three parallel fully connected layers predicting mixture weight (π), mean (μ), and the covariance matrix (Σ) for each Gaussian. For C mixtures, the sizes of the output layers are $C, C \times 2$ and $C \times 2$ for π, μ and Σ respectively. We predict only the diagonal elements of the covariance matrix Σ .

TABLE I
VALIDATION RESULTS ON ALL THREE STUDIED DATASETS

	SimpleNet					MDNSal				
	KLdiv	CC	AUC	NSS	SIM	KLdiv	CC	AUC	NSS	SIM
SALICON	0.193	0.907	0.871	1.926	0.797	0.217	0.899	0.868	1.893	0.797
MIT1003	0.558	0.786	0.907	2.870	0.626	0.634	0.779	0.904	2.814	0.624
CAT2000	0.256	0.895	0.883	2.400	0.758	0.293	0.889	0.878	2.329	0.751

TABLE II
SIMPLENET’S VALIDATION RESULTS ON SALICON WITH VARIOUS ENCODERS

Models	CC	NSS	KLdiv	AUC	SIM
VGG-16	0.871	1.863	0.238	0.864	0.772
ResNet-50	0.895	1.881	0.211	0.868	0.786
DenseNet-161	0.902	1.930	0.210	0.87	0.795
PNASNet-5	0.907	1.926	0.193	0.871	0.797

TABLE III
SIMPLENET’S VALIDATION RESULTS ON SALICON WITH VARIOUS LOSS FUNCTIONS

Loss Functions	CC	KLdiv	NSS	AUC	SIM
KL	0.904	0.223	1.935	0.870	0.797
KL + CC	0.906	0.192	1.925	0.871	0.798
KL + CC + NSS	0.900	0.204	1.998	0.872	0.794

d) *Loss function:* We define Negative log-likelihood (NLL) loss function to train the parameters of the Gaussians as follows:

$$NLL(P, Q) = - \sum_i q_i \log(p_i + \epsilon). \quad (3)$$

Where i represents exhaustive sampling across spatial dimensions of the image, p_i is the likelihood of the sampled point to fit the distribution with C Gaussians, q_i is the corresponding ground truth value and ϵ is a small constant. p_i is further defined as follows:

$$p(x; \pi, \mu, \Sigma) = \sum_{c=1}^C \pi_c \frac{1}{\sqrt{(2\pi_c)^2 |\Sigma_c|}} e^{-\frac{1}{2}(x-\mu_c)^T \Sigma_c^{-1} (x-\mu_c)}. \quad (4)$$

Similar to SimpleNet we use a combination of NLL and CC for training MDNSal.

IV. EXPERIMENTS AND RESULTS

A. Datasets

a) *SALICON*:: We use SALICON dataset for training our models. It consists of 10,000 training, 5000 validation, and 5000 test images. It is the largest dataset for saliency prediction and was labeled based on mouse-tracking (shown to be equivalent to the eye-fixations recorded with an eye-tracker). Our experiments are based on the newest release, SALICON 2017, from the LSUN challenge.

b) *MIT300*:: MIT300 test dataset consists of 300 natural images with eye-tracking data of 39 observers. The labels of MIT300 are non-public. We use MIT1003 consisting of 1003 images to fine-tune the initial model trained on

TABLE IV
MDNSAL’S VALIDATION RESULTS ON SALICON WITH VARIOUS NUMBER OF GAUSSIANS

No of Gaussians(C)	CC	KLdiv	NSS	AUC	SIM
8	0.882	0.256	1.849	0.864	0.778
16	0.892	0.240	1.881	0.867	0.787
24	0.895	0.233	1.887	0.867	0.789
32	0.899	0.224	1.892	0.868	0.797
48	0.896	0.231	1.889	0.868	0.790
64	0.896	0.230	1.892	0.868	0.790

SALICON dataset. We have done 10 fold cross-validation by splitting the images into 903 train, and 100 validation and have chosen the best model for test submission.

c) *CAT2000*:: CAT2000 consists of 4000 images(2000 train and 2000 test) taken from 20 different categories like Action, Art, Cartoon etc., with eye-tracking data from 24 observers. Similar to MIT1003, we split the images into 1800 train and 200 validation and perform fine-tuning on our model.

B. Experimental Setup

We resize the input images into 256x256 resolution for both the models. We train SimpleNet for 10 epochs with learning rate starting from 1e-4 and reducing it after 5 epochs. MDNSal is trained for 50 epochs with learning rate 1e-4. We use 32 Gaussians ($C = 32$) in MDNSal. Backpropagation was performed using ADAM optimizer in both the networks. The model trained on SALICON was fine-tuned using MIT1003 and CAT2000. We submitted the test results for SALICON to LSUN17¹ and MIT300 test results to MIT Saliency Benchmark². We only present validation results on the CAT2000 dataset.

C. Ablation Analysis

We examine the effects of (a) changing the input feature, (b) using different combinations of the loss function, and (c) the significance of hierarchy. The analysis is made using SimpleNet model on SALICON validation set. Table II illustrates the results by varying the pre-trained network for the input feature component. PNASNet-5 achieves the best overall results and is used as the backbone for all the following experiments. Ablation results by adding CC and NSS to the KLdiv loss are presented in Table III. Adding CC improves the performance over just using KLdiv loss (on both KL and CC metrics). Higher performance on the NSS metric can be achieved by adding an NSS term to the

¹<http://saliicon.net/challenge-2017/>

²<http://saliency.mit.edu>

TABLE V
OUR PROPOSED MODELS PERFORMANCE ON THE SALICON TEST.

	KLdiv↓	CC↑	AUC↑	NSS↑	SIM↑	IG↑	sAUC↑
EMLNET [19]	0.520	0.886	0.866	2.050	0.780	0.736	0.746
SAM-Resnet [20]	0.610	0.899	0.865	1.990	0.793	0.538	0.741
MSI-Net [35]	0.307	0.889	0.865	1.931	0.784	0.793	0.736
GazeGAN [36]	0.376	0.879	0.864	1.899	0.773	0.720	0.736
MDNSal (Ours)	0.221	0.899	0.865	1.935	0.790	0.863	0.736
SimpleNet (Ours)	0.201	0.907	0.869	1.960	0.793	0.880	0.743

TABLE VI
OUR PROPOSED MODELS PERFORMANCE ON THE MIT SALIENCY BENCHMARK.

	KLdiv↓	CC↑	AUC↑	NSS↑	SIM↑	sAUC↑	EMD↓
EMLNET [19]	0.84	0.79	0.88	2.47	0.68	0.70	1.84
DeepGaze2 [27]	0.96	0.52	0.88	1.29	0.46	0.72	3.98
SALICON [28]	0.54	0.74	0.87	2.12	0.60	0.74	2.62
DPNSal [37]	0.91	0.82	0.87	2.41	0.69	0.74	2.05
DenseSal [38]	0.48	0.79	0.87	2.25	0.67	0.72	1.99
DVA [31]	0.64	0.68	0.85	1.98	0.58	0.71	3.06
MDNSal (Ours)	0.47	0.78	0.86	2.25	0.67	0.71	1.96
SimpleNet (Ours)	0.42	0.79	0.87	2.30	0.67	0.71	2.06

loss; however, it brings minor reductions in the KLdiv and CC metric. To keep things minimal, we use KLdiv+CC loss for later experiments.

We also explore the significance of multi-level integration by learning SimpleNet by just using the last conv layer of PNASNet-5. The CC drops to 0.89 from 0.907, and KLdiv increases to 0.22 from 0.193, indicating the importance of multi-level integration. Finally, the results of the validation set on all three datasets are presented in Table I.

We perform another set of experiments on MDNSal. The first ablation experiment is aimed to understand the impact of changing the number of Gaussians(C). The results are presented in Table IV and it is observed that using 32 Gaussians gives best performance and thus we use the same number of mixtures in the later experiments. As next experiment, we relax the constraint of only predicting the diagonal values of the covariance matrix. We predict the full covariance matrix with positive-definite constraints, which is necessary to compute the loss. To enforce this constraint we adopt the method by [39], where we predict lower triangular matrix(L) and get covariance matrix using $A = LL^T$ and $\Sigma^{-1} = A$. Using full covariance matrix however did not give any visible improvements (CC remained same to 0.899 on SALICON validation set). Hence, we use the diagonal approximation in all the remaining experiments on MDNSal.

D. Comparison with state of the art

We quantitatively compare our models with state of the art on SALICON and MIT300 test sets. Table V shows the results on the SALICON dataset in terms of KLdiv, CC, AUC, NSS, SIM, IG, and sAUC metrics. SimpleNet gives consistent results on all seven metrics. It achieves the best performance on five different metrics and outperforms state of the art by a large margin on KLdiv and IG. We are third-best in NSS metric; however, if crucial, that can be compensated by adding an NSS loss, as indicated in the ablation studies. Although parametric, surprisingly, MDNSal

also gives competent performance across various metrics (only second to SimpleNet on four metrics). SimpleNet and MDNSal also achieve state of the art performance on MIT300 test dataset, as shown in Table VI. SimpleNet gives the best results on KLdiv and is competent in all other metrics. MDNSal gives a similar performance with the second-best results on KLdiv and EMD.

The work by [9] recommends CC as one of the ideal metrics to report, as it makes limited assumptions about input format and treats both false positives and negatives symmetrically. They further suggest KL and IG as good choices concerning benchmark intended to evaluate saliency maps as probability distributions. Both our models give a leading performance on KLdiv and CC on both MIT300 and SALICON test sets, which makes them an ideal choice for the task of saliency prediction.

We qualitatively compare results of the proposed models with other state of the art methods. The results on couple of images from MIT300 test dataset are shown in Figure 5. The ground truth images are chosen from the carefully curated set in [32]. Our model performs well both in terms on coverage (predicting all the salient regions), accurate localization and the relative order of importance.

V. CONCLUSION

In this paper, we identify four key components of the saliency detection architectures and analyze how the previous literature has approached the individual components. The analysis helps to explore agreements, redundancies, gaps, or need for optimization over these components. Using that as a basis, we propose two novel architectures called SimpleNet and MDNSal. SimpleNet improves upon the encoder-decoder architectures, and MDNSal opens up a new paradigm of parametric modeling. Both models are devoid of complexities like prior maps, multiple input streams, or recurrent units and still achieve the state of the art performance on public saliency benchmarks. Our work suggests that the way

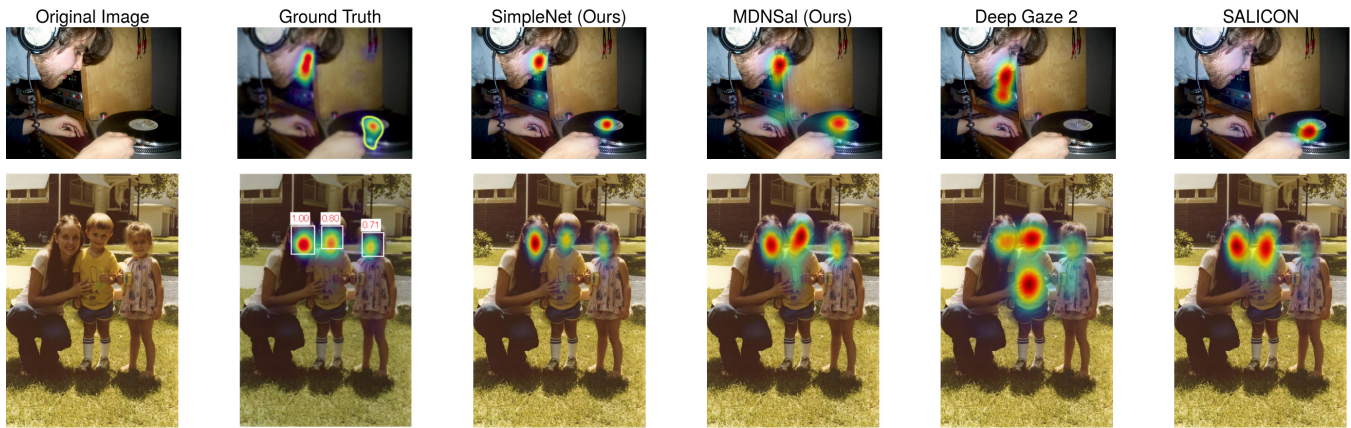


Fig. 5. Examples of predicted saliency maps. Both images are taken from MIT300 test set and ground truth images are taken from [32]. We compare results of the proposed SimpleNet and MDNSal models with other state of the art approaches. First row (action): SimpleNet and MDNSal accurately predicts both the person’s face and where he is looking (indicated by yellow boundary). In contrast other models miss out on either the action or the face. Second row (Faces with relative importance): our model gives accurate localization on all three faces and preserves relative importance.

forward is not necessarily to design more complex architectures but a modular analysis to optimize each component and possibly explore novel (and simpler) alternatives.

REFERENCES

- [1] C. Koch and S. Ullman, “Shifts in selective visual attention: towards the underlying neural circuitry,” in *Matters of intelligence*. Springer, 1987, pp. 115–141.
- [2] L. Itti, C. Koch, and E. Niebur, “A model of saliency-based visual attention for rapid scene analysis,” *IEEE TPAMI*, no. 11, pp. 1254–1259, 1998.
- [3] M. Staudte and M. W. Crocker, “Visual attention in spoken human-robot interaction,” in *2009 4th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2009, pp. 77–84.
- [4] B. Schauerte, J. Richarz, and G. A. Fink, “Saliency-based identification and recognition of pointed-at objects,” in *IROS*, 2010.
- [5] A. Borji and L. Itti, “Scene classification with a sparse set of salient regions,” in *ICRA*, 2011.
- [6] S. Frintrop and M. Kessel, “Most salient region tracking,” in *ICRA*, 2009.
- [7] L. Chen, P. Huang, and Z. Zhao, “Saliency based proposal refinement in robotic vision,” in *2017 IEEE International Conference on Real-time Computing and Robotics (RCAR)*. IEEE, 2017, pp. 85–90.
- [8] A. Rasouli, P. Lanillos, G. Cheng, and J. K. Tsotsos, “Attention-based active visual search for mobile robots,” *Autonomous Robots*, pp. 1–16, 2019.
- [9] Z. Bylinskii, T. Judd, A. Oliva, A. Torralba, and F. Durand, “What do different evaluation metrics tell us about saliency models?” *IEEE TPAMI*, vol. 41, no. 3, pp. 740–757, 2018.
- [10] M. Jiang, S. Huang, J. Duan, and Q. Zhao, “Salicon: Saliency in context,” in *CVPR*, 2015.
- [11] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *MICCAI*, 2015.
- [12] Z. Bylinskii, T. Judd, A. Borji, L. Itti, F. Durand, A. Oliva, and A. Torralba, “Mit saliency benchmark,” 2015.
- [13] J. Zhang and S. Sclaroff, “Saliency detection: A boolean map approach,” in *ICCV*, 2013.
- [14] N. D. Bruce and J. K. Tsotsos, “Saliency, attention, and visual search: An information theoretic approach,” *Journal of vision*, vol. 9, no. 3, pp. 5–5, 2009.
- [15] T. Judd, K. Ehinger, F. Durand, and A. Torralba, “Learning to predict where humans look,” in *ICCV*, 2009.
- [16] M. Kümmerer, L. Theis, and M. Bethge, “Deep gaze i: Boosting saliency prediction with feature maps trained on imagenet,” *arXiv:1411.1045*, 2014.
- [17] S. S. Kruthiventi, K. Ayush, and R. V. Babu, “Deepfix: A fully convolutional neural network for predicting human eye fixations,” *IEEE TIP*, vol. 26, no. 9, pp. 4446–4456, 2017.
- [18] M. Cornia, L. Baraldi, G. Serra, and R. Cucchiara, “A deep multi-level network for saliency prediction,” in *ICPR*, 2016.
- [19] S. Jia and N. D. Bruce, “Eml-net: An expandable multi-layer network for saliency prediction,” *arXiv:1805.01047*, 2018.
- [20] M. Cornia, L. Baraldi, G. Serra, and R. Cucchiara, “Predicting human eye fixations via an lstm-based saliency attentive model,” *IEEE TIP*, vol. 27, no. 10, pp. 5142–5154, 2018.
- [21] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” in *NIPS*, 2012.
- [22] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” *arXiv:1409.1556*, 2014.
- [23] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *CVPR*, 2016.
- [24] B. Zoph, V. Vasudevan, J. Shlens, and Q. V. Le, “Learning transferable architectures for scalable image recognition,” in *CVPR*, 2018.
- [25] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, “Densely connected convolutional networks,” in *CVPR*, 2017.
- [26] N. Liu and J. Han, “A deep spatial contextual long-term recurrent convolutional network for saliency detection,” *IEEE TIP*, vol. 27, no. 7, pp. 3264–3274, 2018.
- [27] M. Kummerer, T. S. Wallis, L. A. Gatys, and M. Bethge, “Understanding low-and high-level contributions to fixation prediction,” in *ICCV*, 2017.
- [28] X. Huang, C. Shen, X. Boix, and Q. Zhao, “Salicon: Reducing the semantic gap in saliency prediction by adapting deep neural networks,” in *ICCV*, 2015.
- [29] N. Liu, J. Han, T. Liu, and X. Li, “Learning to predict eye fixations via multiresolution convolutional neural networks,” *IEEE transactions on neural networks and learning systems*, vol. 29, no. 2, 2016.
- [30] G. Li and Y. Yu, “Visual saliency based on multiscale deep features,” in *CVPR*, 2015.
- [31] W. Wang and J. Shen, “Deep visual attention prediction,” 2017.
- [32] Z. Bylinskii, A. Recasens, A. Borji, A. Oliva, A. Torralba, and F. Durand, “Where should saliency models look next?” in *ECCV*, 2016.
- [33] M. Kummerer, T. S. Wallis, and M. Bethge, “Saliency benchmarking made easy: Separating models, maps and metrics,” in *ECCV*, 2018.
- [34] C. M. Bishop, “Mixture density networks,” 1994.
- [35] A. Kroner, M. Senden, K. Driessens, and R. Goebel, “Contextual encoder-decoder network for visual saliency prediction,” *arXiv:1902.06634*, 2019.
- [36] Z. Che, A. Borji, G. Zhai, X. Min, G. Guo, and P. Le Callet, “How is gaze influenced by image transformations? dataset and model,” *IEEE TIP*, 2019.
- [37] T. Oyama and T. Yamanaka, “Influence of image classification accuracy on saliency map estimation,” 2018.
- [38] —, “Fully convolutional densenet for saliency-map prediction,” in *ACPR*, 2017.
- [39] G. Dorta, S. Vicente, L. Agapito, N. D. Campbell, and I. Simpson, “Structured uncertainty prediction networks,” in *CVPR*, 2018.