# **End-to-end Convolutional Network for Saliency Prediction**

Junting Pan and Xavier Giró-i-Nieto Universitat Politecnica de Catalunya (UPC) Barcelona, Catalonia/Spain

junting.pan@alu-etsetb.upc.edu, xavier.giro@upc.edu

## Abstract

The prediction of saliency areas in images has been traditionally addressed with hand crafted features based on neuroscience principles. This paper however addreses the problem with a completely data-driven approach by training a convolutional network. The learning process is formulated as a minimization of a loss function that measures the Euclidean distance of the predicted saliency map with the provided ground truth. The recent publication of large datasets of saliency prediction has provided enough data to train a not very deep architecture which is both fast and accurate. The convolutional network in this paper, named JuntingNet, won the LSUN 2015 challenge on saliency prediction with a superior performance in all considered metrics.

### 1. Introduction

This work presents an end-to-end convolutional network (convnet) for saliency prediction. Our objective is to compute saliency maps that represent the probability of visual attention. This problem has been traditionally addressed with hand-crafted features inspired by neurology studies. In our case we have adopted a completely data-driven approach, training a model with a large amount of annotated data.

Convnet is a popular architecture in the field of deep learning and has been widely explored for visual pattern recognition, ranging from a global scale image classification to a more local object detection or semantic segmentation. The hierarchy of layers of convnets are also inspired by biological models and actually recent works have pointed at a relation between the activity of certain areas in the brain with hierarchy of layers in the convnets [1]. Provided with enough training data, convnets show impressive results, often outperforming other hand-crafted methods . In many popular works, the output of the convnet is a discrete label associated to a certain semantic class. The saliency prediction problem, though, addresses the problem



Stimuli Image (iSUN)





Eye Fixation Data(iSUN)



Stimuli Image (SALICON)

Mouse Click Data (SALICON)

Figure 1. Images (right) and saliency maps (left) from the iSUN and SALICON datasets.

of a continuous range of values that estimate the probability of a human fixation on a pixel. These values present a spatial coherence and smooth transition that this work addresses by using the convnet as a regression solver, instead of a classifier.

The training of a convolutional network requires a large amount of annotated data that provides a rich description of the problem. Our work has benefited from the recent publication of two datasets: iSun [21] and SALICON [12]. These datasets propose two different approaches for saliency prediction. While iSun was generated with an eye-tracker to annotate the gaze fixations, the SALICON dataset was built by asking humans to click on the most salient points on the image. The different nature of the saliency maps of the two datasets can be seen in Figure 1. The large size of these datasets has provided for the first time the possibility of training a convnet.

Our main contribution has been the design of an endto-end convnet for saliency prediction, the first one from this type, up to the authors knowledge. The network, called *JuntingNet*, has proved its superior performance in the Large-scale Scene UNderstanding (LSUN) challenge 2015 [23]. The developed model has been publicly available at http://bit.ly/juntingnet.

This paper is structured as follows. Section 2 presents the previous works using convolutional networks for saliency prediction. Our system is presented in Section 3 and its results on the LSUN challenge reported in Section 4. The conclusions and future directions are contained in Section 5.

### 2. Related work

JuntingNet presents the next natural step to two main trends in deep learning: using convnets for saliency prediction and training these networks by formulating and end-toend problem. This section refers to some related work in these two fields.

### 2.1. Deep learning for saliency prediction

An early attempt of predicting saliency model with a convnet was the *ensembles of Deep Networks (eDN)* [19], which proposed an optimal blend of features from three different convnet layers who were finally combined with a simple linear classifier trained with positive (salient) or negative (non-salient) local regions. This approach inspired *DeepGaze* [15], which only combined features from different layers but, in this case, from a much deeper network. In particular, *DeepGaze* used the existing *AlexNet* convnet [14], which had been trained for an object classification task, not for saliency prediction. *JuntingNet* adopts a not very deep architecture as *eDN*, but it is end-to-end trained as a regression problem, avoiding the reuse of precomputed parameters from another task.

#### 2.2. End to end semantic segmentation

Fully Convolutional Networks (FCNs) [17] addressed the semantic segmentation task which predicting the semantic label of every individual pixel in the image. This approach dramatically improved previous results on the challenging PASCAL VOC segmentation benchmark [6]. The idea of an end-to-end solution for a 2D problem as as semantic segmentation was refined by DeepLab-CRF [5], where the spatial consistency of the predicted labels is checked with a Conditional Random Field (CRF), similarly to the hierarchical consistency enforced in [7]. In our work, we adopt the end-to-end solution for a regression problem instead of a classification one, and we also introduce a postfiltering stage, which consists of a Gaussian filtering that smoothes the resulting saliency map.

### 3. JuntingNet

This paper presents *JuntingNet*, an end-to-end convnet for saliency prediction. The parameters of our network are learned by minimizing an Euclidean loss function defined directly on the ground truth saliency maps.

#### 3.1. Architecture

The detailed architecture of JuntingNet is illustrated in Figure 2. The network contains five learned layers: three convolutional layers and two fully connected layers, which can also be interpreted as 1x1 convolutions.

The proposed architecture is not very deep if compared to other networks in the state of the art. Popular architectures trained on the 1, 200, 000 images of the ILSRVC 2012 challenge proposed from 7 [14] to 22 layers [18]. *JuntingNet* is defined by only 5 layers which are trained separately on two training datasets collections of diverse sizes: 6,000 for iSun and 10,000 for SALICON. This adopted shallow depth tries to prevent the overfitting problem, which is a great risk for models with a large amount of parameters, such as convnets.

The detailed description of the convnet stages is the following:

- 1. The input volume has size of [96x96x3] (RGB image), a size smaller than the [227x227x3] proposed in AlexNet. Similarly to the shallow depth, this design parameter is motivated to reduce the possibilities of overfitting.
- 2. The receptive field of the first 2D convolution is of size [5x5], and its outputs define a convolutional layer with 32 neurons. This layer is followed by a ReLU activation layer which applies an element wise non-linearity. Later, a max pooling layer progressively reduces the spatial size of the input image. Despite the loss of visual resolution at the output, this reduction also reduces the amount of model parameters and prevents overfitting. The max-pooling layer selects the maximum value of every [2x2] region, taking strides of two pixels.
- 3. The output of the previous stage has a size of [46x46x32]. The receptive field of this second stage is [3x3]. Again, this is followed by a RELU layer and a max-pooling layer of size [2x2].
- 4. Finally, the last convolutional layer is fed with an input of size [22x22x64]. The receptive of this layer is also of [3x3] and it has 64 neurons. A ReLU and max pooling layers are stacked too.
- 5. A first fully connected layer receives the output of the third convolutional layer with a dimension of [10x10x64]. It contains a total of 4,608 neurons.



Figure 2. Convnet architecture for JuntingNet.

- 6. The second fully connected layer consist of a maxout layer with 2,304 neurons. The maxout operation [8] computes the pairs of the previous layers output.
- 7. Finally, the output of the last maxout layer is the saliency prediction array. The array is reshaped to have 2D dimensions and resized to the stimuli image size. Finally, a 2D Gaussian filter with a standard deviation of 3.0 is applied.

#### **3.2.** Training parameters

The limited amount of training data for our architecture made overfitting a significant challenge, so we used different techniques to minimize its effects. Firstly, we apply norm constraint regularization for the maxout layers [8]. Secondly, we use data augmentation technique by mirroring all images. We also tested a dropout layer [10] after the first fully connected layer, with a dropout ratio of 0.5 (50% of probability to set a neurons output value to zero). However, this did not make much of a difference, so it is not included to the final model.

The weights in all layers are initialized from a normal Gaussian distribution with zero mean and a standard deviation of 0.01, with biases initialized to 0.1. Ground truth values that we used for training are saliency maps with normalized values between 0 and 1.

For validation control purposes, we split the training partitions of iSUN and SALICON datasets into 80% for training and the rest for real time validation. The network was trained with stochastic gradient descent (SGD) and Nesterov momentum SGD optimization method that helps the loss function to converge faster. The learning rate



Figure 3. Learning curves for iSUN models.



Figure 4. Learning curves for SALICON models.

was changing over time; it started with a higher learning rate 0.03 and decreased during the course of training until 0.0001. We set 1,000 epochs to train a separate network for each dataset. Figures 3 and 4 present the learning curves for the iSUN and SALICON models, respectively.

### 4. Experiments

### 4.1. Datasets

The network was tested in the two datasets proposed in the LSUN challenge [23]:

- **iSUN [21]:** a ground truth of gaze traces on images from the SUN dataset [20]. The collection is partitioned into 6,000 images for training, 926 for validation and 2,000 for test.
- **SALICON [12]:** cursor clicks on the objects of interest from images of the Microsoft COCO dataset [16]. The collection contains 10,000 training images, 5,000 for validation and 5,000 for test.

### 4.2. Results

Our solution is implemented using Python, NumPy and the deep learning library Theano [3, 2]. Processing was performed on an NVidia GPU GTX 980 with 2048 CUDA cores and 4GB of RAM. Our network took between six to seven hours to train for the SALICON dataset, and five to six hours for the iSUN dataset. Every saliency prediction requires 200 ms per image.

We assessed our model on the LSUN saliency prediction challenge 2015 [23]. Table 1 and Table 2 presents our results for iSUN and SALICON datasets. The model was evaluated separately on the testing data of each datasets. The evaluation metrics was adopted of the variety of metrics provided in MIT saliency benchmark [13, 4] defined on both saliency map and fixation points. *JuntingNet* consistenly won the first place of the challenge in all metrics considered in the challenge. A few qualitative results are also provided in Figure 2.

## 5. Conclusions

We designed the first end-to-end ConvNet for saliency prediction, trained only with the datasets of visual saliency provided by the LSUN challenge. With this ConvNet we were able to win the first place in the challenge by large margin. Our results demonstrate that a not very deep ConvNets are capable of achieving good results on a highly challenging task.

Our experiments can be considered as preliminary, as only one configuration and set up was considered. We expect that a more elaborate study of the architecture, use of the dataset and training parameters could still improve the reported performance.

The developed model has been publicly available from http://bit.ly/juntingnet.

### 6. Acnowledgements

We would like to thank the technical support of Albert Gil and Josep Pujal in the setting up of software and hardware necessary to run the experimentation.

The Image Processing Group at the UPC is a SGR14 Consolidated Research Group recognized and sponsored by the Catalan Government (Generalitat de Catalunya) through its AGAUR office.

This work has been developed in the framework of the project BigGraph TEC2013-43935-R, funded by the Spanish Ministerio de Economa y Competitividad and the European Regional Development Fund (ERDF).

We gratefully acknowledge the support of NVIDIA Corporation with the donation of the GeoForce GTX 980 used in this work.

### References

- P. Agrawal, D. Stansbury, J. Malik, and J. L. Gallant. Pixels to voxels: Modeling visual representation in the human brain. *arXiv preprint arXiv:1407.5104*, 2014.
- [2] F. Bastien, P. Lamblin, R. Pascanu, J. Bergstra, I. Goodfellow, A. Bergeron, N. Bouchard, D. Warde-Farley, and Y. Bengio. Theano: new features and speed improvements. *arXiv preprint arXiv:1211.5590*, 2012.
- [3] J. Bergstra, O. Breuleux, F. Bastien, P. Lamblin, R. Pascanu, G. Desjardins, J. Turian, D. Warde-Farley, and Y. Bengio. Theano: a cpu and gpu math expression compiler. In *Proceedings of the Python for scientific computing conference* (*SciPy*), volume 4, page 3. Austin, TX, 2010.
- [4] Z. Bylinskii, T. Judd, A. Borji, L. Itti, F. Durand, A. Oliva, and A. Torralba. Mit saliency benchmark. http://saliency.mit.edu/.
- [5] L. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Semantic image segmentation with deep convolutional nets and fully connected crfs. *CoRR*, abs/1412.7062, 2014.
- [6] M. Everingham, S. A. Eslami, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman. The pascal visual object classes challenge: A retrospective. *International Journal of Computer Vision*, 111(1):98–136, 2014.
- [7] C. Farabet, C. Couprie, L. Najman, and Y. LeCun. Learning hierarchical features for scene labeling. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 35(8):1915– 1929, 2013.
- [8] I. J. Goodfellow, D. Warde-Farley, M. Mirza, A. Courville, and Y. Bengio. Maxout networks. *arXiv preprint* arXiv:1302.4389, 2013.
- [9] J. Harel, C. Koch, and P. Perona. Graph-based visual saliency. In Advances in neural information processing systems, pages 545–552, 2006.
- [10] G. E. Hinton, N. Srivastava, A. Krizhevsky, I. Sutskever, and R. R. Salakhutdinov. Improving neural networks by preventing co-adaptation of feature detectors. *arXiv preprint arXiv:1207.0580*, 2012.



Figure 5. Saliency maps generated by JuntingNet. The first column corresponds to the input image, the second column the prediction from JuntingNet and the third one and the third on to the provided ground truth. First three rows correspond to images from the iSUN dataset, while the last three are from the SALICON dataset.

	Similarity	CC	AUC shuffled	AUC Borji	AUC Judd
Our work	0.6833	0.8230	0.6650	0.8463	0.8693
Xidian	0.5713	0.6167	0.6484	0.7949	0.8207
WHU IIP	0.5593	0.6263	0.6307	0.7960	0.8197
LCYLab	0.5474	0.5699	0.6259	0.7921	0.8133
Rare 2012 Improved	0.5199	0.5199	0.6283	0.7582	0.7846
Baseline: BMS [22]	0.5026	0.3465	0.5885	0.6560	0.6914
Baseline: GBVS [9]	0.4798	0.5087	0.6208	0.7913	0.8115
Baseline: Itti [11]	0.4251	0.3728	0.6024	0.7262	0.7489

Table 1. Results of the LSUN challenge 2015 for saliency prediction with the iSUN dataset.

	Similarity	CC	AUC shuffled	AUC Borji	AUC Judd
Our work	0.5198	0.5957	0.6698	0.8291	0.8364
WHU IIP	0.4908	0.4569	0.6064	0.7759	0.7923
Rare 2012 Improved	0.5017	0.5108	0.6644	0.8047	0.8148
Xidian	0.4617	0.4811	0.6809	0.7990	0.8051
Baseline: BMS [22]	0.4542	0.4268	0.6935	0.7699	0.7899
Baseline: GBVS [9]	0.4460	0.4212	0.6303	0.7816	0.7899
Baseline: Itti [11]	0.3777	0.2046	0.6101	0.6603	0.6669

Table 2. Results of the LSUN challenge 2015 for saliency prediction with the SALICON dataset.

- [11] L. Itti, C. Koch, and E. Niebur. A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (11):1254–1259, 1998.
- [12] M. Jiang, S. Huang, J. Duan, and Q. Zhao. SALICON: Saliency in context. In *Proceedings of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [13] T. Judd, F. Durand, and A. Torralba. A benchmark of computational models of saliency to predict human fixations. In *MIT Technical Report*, 2012.
- [14] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [15] M. Kümmerer, L. Theis, and M. Bethge. Deep gaze i: Boosting saliency prediction with feature maps trained on imagenet. arXiv preprint arXiv:1411.1045, 2014.
- [16] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *Computer Vision–ECCV 2014*, pages 740–755. Springer, 2014.
- [17] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. *CoRR*, abs/1411.4038, 2014.
- [18] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. *arXiv preprint arXiv:1409.4842*, 2014.
- [19] E. Vig, M. Dorr, and D. Cox. Large-scale optimization of hierarchical features for saliency prediction in natural images. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*, pages 2798–2805. IEEE, 2014.

- [20] J. Xiao, J. Hays, K. Ehinger, A. Oliva, A. Torralba, et al. Sun database: Large-scale scene recognition from abbey to zoo. In *Computer vision and pattern recognition (CVPR), 2010 IEEE conference on*, pages 3485–3492. IEEE, 2010.
- [21] P. Xu, K. A. Ehinger, Y. Zhang, A. Finkelstein, S. R. Kulkarni, and J. Xiao. Turkergaze: Crowdsourcing saliency with webcam based eye tracking. *arXiv preprint arXiv:1504.06755*, 2015.
- [22] J. Zhang and S. Sclaroff. Saliency detection: A boolean map approach. In *Computer Vision (ICCV), 2013 IEEE International Conference on*, pages 153–160. IEEE, 2013.
- [23] Y. Zhang, F. Yu, S. Song, P. Xu, A. Seff, and J. Xiao. Largescale scene understanding challenge: Eye tracking saliency estimation.