



HAL
open science

Deep learning investigation for chess player attention prediction using eye-tracking and game data

Justin Le Louedec, Thomas Guntz, James L. Crowley, Dominique Vaufreydaz

► To cite this version:

Justin Le Louedec, Thomas Guntz, James L. Crowley, Dominique Vaufreydaz. Deep learning investigation for chess player attention prediction using eye-tracking and game data. ETRA 2019 - ACM Symposium On Eye Tracking Research & Applications, Jun 2019, Denver, United States. pp.1-15, 10.1145/3314111.3319827 . hal-02100389

HAL Id: hal-02100389

<https://inria.hal.science/hal-02100389>

Submitted on 15 Apr 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Deep learning investigation for chess player attention prediction using eye-tracking and game data

Justin Le Louedec, Thomas Guntz, James L. Crowley, Dominique Vaufreydaz

Univ. Grenoble Alpes, CNRS, Inria, Grenoble INP, LIG, 38000 Grenoble, France

Author Version

Abstract

This article reports on an investigation of the use of convolutional neural networks to predict the visual attention of chess players. The visual attention model described in this article has been created to generate saliency maps that capture hierarchical and spatial features of chessboard, in order to predict the probability fixation for individual pixels. Using a skip-layer architecture of an autoencoder, with a unified decoder, we are able to use multiscale features to predict saliency of part of the board at different scales, showing multiple relations between pieces. We have used scan path and fixation data from players engaged in solving chess problems, to compute 6600 saliency maps associated to the corresponding chess piece configurations. This corpus is completed with synthetically generated data from actual games gathered from an online chess platform. Experiments realized using both scan-paths from chess players and the CAT2000 saliency dataset of natural images, highlights several results. Deep features, pretrained on natural images, were found to be helpful in training visual attention prediction for chess. The proposed neural network architecture is able to generate meaningful saliency maps on unseen chess configurations with good scores on standard metrics. This work provides a baseline for future work on visual attention prediction in similar contexts.

Keywords: Deep neural network, Computer vision, Visual attention, Chess

1 Introduction

Humans possess a remarkable ability to discriminate and select relevant information from an environment without attending to the entire scene. This ability for selective attention enables very fast interpretation and rapid interaction. Models for this process are of growing interest for both computer vision and neurosciences. Modelling visual attention not only provides an understanding of how the human visual system processes information but also has practical applications for constructing artificial systems for object recognition, event detection [Gu et al., 2015; González-Díaz et al., 2013; Buso et al., 2015], action recognition [Safaei and Foroosh, 2017; Daria, 2016], object segmentation [Duan et al., 2014; Xu et al., 2017], as well as computer graphics [Lavoué et al., 2018; Bülbül, 2012].

Most work on Visual attention prediction has attempted to quantitatively or qualitatively predict gaze location. These locations are often displayed in the form of saliency maps, i.e. a map that displays the probability that each pixel may be fixated (receive visual attention). For example, the right side of figure 1 shows a saliency map computed from eye-tracking data. In this example, one can see that whiter parts of the saliency maps are more frequently fixated than others. Aggregated eye-tracking information from many users can be used to estimate the probability that a pixel will be fixated. However, accurate prediction requires an extensive amounts of eye-tracking data.

Current techniques for saliency or visual attention prediction fall into two main categories: Bottom-Up and top-Down approaches. Bottom-Up approaches are common in computer vision where, local image appearance is assumed to stimulate fixation, and the goal is to provide a probability that each pixel will be fixated without prior knowledge of the task or task context. Top down approaches are generally task-driven and require an understanding of the image, scene, context, or task. Visual

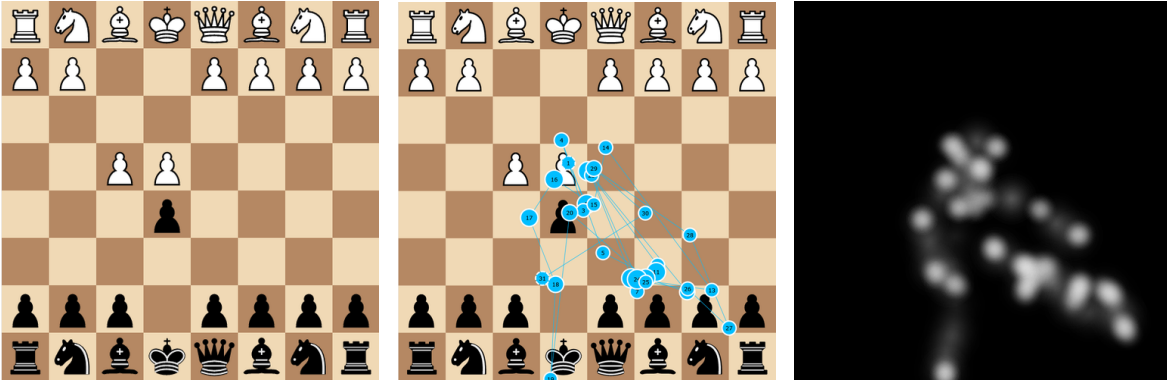


Figure 1: Saliency map from a chess game. At left, input chess board image. At centre, eye tracking results (points represent eye fixations, lines are scan path between fixations). At right, saliency map computed from eye tracking. For each pixel, a probability between 0 (black) and 1 (white) is computed.

attention in Chess can be seen as a combination of these two approaches as salient areas may be defined by both visual pattern recognition (bottom-up analysis) and an understanding of the game (top-down prediction) [Charness et al., 2001; Reingold and Charness, 2005; Sheridan and Reingold, 2017]. This suggests the extraction of local features and interpretation from images that are very different from the natural images most commonly used in visual attention prediction. In particular, the saliency of chess pieces depends on the current game configuration, expressed as configurations of pieces (Figure 1).

In this paper, we address the problem of visual attention prediction for chess players using a convolutional neural network (CNN). While CNN’s have been used for attention prediction in non-natural image such as comics [Bannier et al., 2018], we believe that this is first investigation for the use in predicting attention in chess. Our investigation builds on previous works on visual attention [Huang et al., 2015; Kruthiventi et al., 2015; Wang and Shen, 2017] and Chess expertise [Charness et al., 2001; Reingold and Charness, 2005; Sheridan and Reingold, 2017], to create a network architecture capable of predicting visual attention of chess players. We examine several training strategies, and evaluate interest of different training corpora including a state-of-the-art saliency corpus, a bottom-up corpus from actual eye-tracking data from chess players and a top-down corpus from online games built in the context of the project.

The article is organized as follow. Section 2 provides background about the problem of salience prediction. Section 3 describes our approach, choices and our model. In section 4, we describe the importance of data and the creation of our datasets. In section 5, we detail performance results of different training strategies using attention prediction metrics and qualitative results from different chess configurations. In a final section, we discuss the results and limitations.

2 Visual attention in Chess

Visual attention has evolved to enable rapid analysis of interesting and relevant regions in visual scene. This frees resources so that higher-level cognitive processes may perform more complex tasks such as scene understanding or decision making from limited data.

In their paper, Posner and Petersen [1990] argue that visual attention is not the function of a single area of the brain but rather the result of a network of different components. He explains that visual recognition depends on successions of neural processes to extract information visual stimuli.

Modern Eye tracking systems make it possible to record the locations in a visual scene that receive fixation as well as the path between fixations (scan path). Fixations are characterized by the time spent on different parts of the scene and can be used as indicators on most viewed and fixated areas on images or scenes. They are often transformed to continuous saliency maps by convolution with a Gaussian function where intensity indicates duration of fixation.

As said, figure 1 shows an example of eye tracking data with a chess configuration in the background. On the center image, points indicate fixations and the lines between points are the scan path of gaze

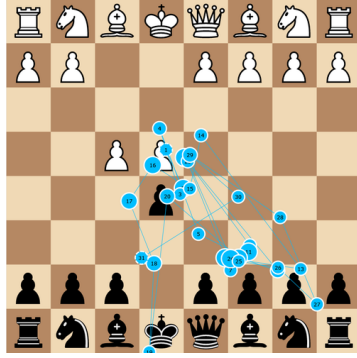


Figure 2: Example of eye tracking data for a user on a chess configuration

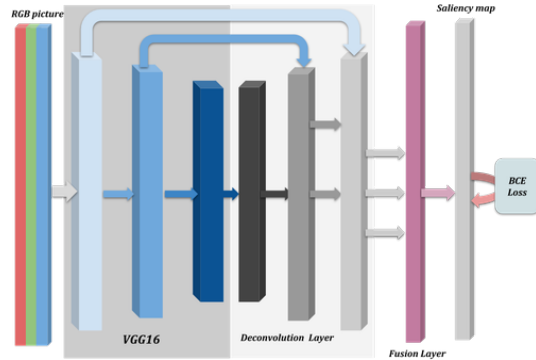


Figure 3: Architecture of our Deep Neural Network

between fixations. This data was acquired for an experimentation on observing and interpreting subjects in problem solving tasks (chess in particular) through multimodal observations (eye tracking, posture, facial emotion, etc.) [Guntz et al., 2018].

Eye tracking maps, captured from human subjects looking at scenes with a small number of salient objects, show that humans tend to focus on a dominant object and ignore background information. Such an approach is not appropriate for Chess as the salient objects are pieces, while their saliency is conditioned by tactical importance of their role in a configuration (Figure 1). In addition, the saliency of individual chess pieces depends on the chess expertise of the player [Reingold and Charness, 2005].

Chess expertise and expertise in general are frequently studied in psychology. Charness et al. [2001] discovered that experts fixate more often on empty squares than intermediate players. Experts also display a much larger number of fixations on relevant pieces than intermediate. It was argued that experts were encoding configurations rather than pieces and produced patterns of saccades (quick movement of eyes between two or more phases of fixation). In their article, Reingold and Charness [2005] demonstrated that expert players use larger visual spans while processing structured chess positions. They also documented that expert players would make fewer fixations and focus on relations between pieces rather than fixate a large number of individual pieces.

In a more recent paper, Sheridan and Reingold [2017] proposed that chess expertise is mainly related to the detection of complex patterns and configurations. When improving their skills, players become capable of detecting increasingly complex patterns. This may explain why experts do not need to fixate pieces individually, and instead detect interesting patterns by looking at their center, using peripheral vision to detect the configuration. This specificity of visual attention of player depending on their skills is a difficulty to address in this research.

3 Deep Neural Network overview

3.1 Architecture Overview

In the last few years, many deep learning architectures have been proposed for prediction of visual attention. Deep learning solutions generally achieve better performance than traditional machine learning techniques. For example, Huang et al. [2015] propose a fine-tuning method with objective functions and integrating information from the image at different scales, achieving top scores for standard metrics such as AUC-Judd or Linear correlation (respectively 0.87 and 0.74, see section 5.1.2 for metric descriptions).

Other models, such as Deepfix [Kruthiventi et al., 2015] or Deep visual attention [Wang and Shen, 2017], have attempted to model visual attention using convolutional neural networks. Deepfix focused on location dependent patterns using "Location biased convolution layer", concatenating Gaussian blobs to the inputs of convolutional layers. Our work is inspired by Deep visual attention, which used

an autoencoder network to generate saliency maps from images using skip layers and a split decoder. Both studies provide good performances on free-view scenes of natural images on the MIT300 dataset [Judd et al., 2012] and made use of multi-scale features.

A particular strength of Convolutional Neural Networks (CNN) is to hierarchically capture features such as corners, edges and shared common patterns from raw images and to integrate such information to recognize more complex patterns such as faces, eyes, or the crown of a king. This integration of high and low level features is useful to capture characteristics of a piece.

The interaction of pieces at different positions on the board determines if an area is salient or not. We base our approach on the work of Wang and Shen [2017], a multi-scale deep neural network encoding features from images to decode them as saliency maps. The model is depicted in Figure 3. As Wang and al. model, it is composed of an auto-encoder with skip layer connections. The encoder part of the model is a VGG16 [Simonyan and Zisserman, 2014] pretrained on ImageNet [Deng et al., 2009], while the decoder part is a succession of deconvolution layers with RELUs, in charge of upsampling feature maps at different dimension/scale levels.

There is a need to express the similarity in the relation between two pieces being close to each other or far away across the board. This means that feature maps at the same dimension must be deconvolved the same way by the decoding part of the network. To do that, we chose to combine the different deconvolution layers of Wang and al. network into only one decoding network. Deconvolution functions in deep learning are in fact transposed convolutions or fractionally strided convolutions which are mislabeled as deconvolution. In fact, these are simply spatial convolutions using a specific padding to up-sample in a learnable way the feature map. The use of deconvolutions rather than fixed kernel upsampling functions comes from learning to discriminate salient or not salient features. Doing so, the network becomes capable of extracting features from an image and deciding which ones are related to salient areas of an image. As an added benefit, this approach provides a higher computation efficiency by progressively reducing the feature space dimensions.

The decoding part of the deep network is a succession of deconvolution step on feature maps of different dimensions to up-sample the signal to the size of the input image. Three feature maps are taken at various stages of the encoding network and are decoded using the decoding network, giving 3 corresponding saliency maps. Those saliency maps are then combined and passed through a convolutional layer called a "fusion layer", which is reduced to form a final saliency map. This architecture makes it possible to learn the transformation of feature maps to salient areas and to combine different scale saliency maps into a single global map describing the input image as displayed in figures 4 to 8.

3.2 Training loss

To train the network, one needs to define a loss function between two saliency maps. The proposed model is close to a standard auto-encoder architecture which basically learns to encode some data and to decode it identically to the input. A direct approach to compare the output to the input would be to compute a distance metric that could be minimized during the training process. The basic, simple and intuitive L1 loss ($|gt - f(x)|$ with gt as the ground-truth and $f(x)$ the output of the auto-encoder) provides a possible solution. It would also be possible to use the L2 distance or Mean Square Error (MSE) expressed as $(gt - f(x))^2/N$ with N being the batch size.

In our case, the network does not learn to generate a copy of the input in the form of a chessboard with pieces on it, but the visual attention prediction from this chessboard. The visual prediction is a probability distribution ranges from 0 to 1 (0 not salient at all and 1 very salient). The *Kullback-Leibler* (KL) divergence [Kullback and Leibler, 1951] expresses the degree of divergence of two probability distributions. Minimizing the KL divergence is equivalent to reducing the distance between them. However, the *Kullback-Leibler* divergence is not useful as a metric, as it does not respect the triangle inequality: it is not symmetrical, it does not have an upper bound, and often $KL(P|Q) \neq KL(Q|P)$ as discussed in [Le Meur and Baccino, 2013; Bylinskii et al., 2016].

Binary Cross Entropy ($-gt.log(f(x)) - (1 - gt).log(1 - f(x))$) is a good candidate in the case of saliency map generation. It can be seen as relatively close to the *Kullback-Leibler* divergence. Indeed, cross entropy is often used to measure the dissimilarity between two distributions. Minimizing cross

entropy makes it possible to generate data as close as possible to the ground-truth, while avoiding *Kullback-Leibler* divergence’s problems. Thus, Binary cross entropy is used as a loss function in the experiments described below 5.

4 Training corpora

4.1 Bottom-up datasets from eye tracking data (*ETDS* and *ETDSa*)

4.1.1 Initial dataset

The initial dataset comes from a multimodal study with people solving chess problems [Guntz et al., 2018]. For these experiments, they recorded the eye-gaze of 30 players (10-51 years, 2 female, age: $M = 30.0$; $SD = 12.9$) from various levels (from beginners to experts, with a majority of intermediate players), trying to solve 13 chess tasks of increasing difficulty in limited time:

- 2 *openings* tasks: *Kings Gambit* and a variation of the *Caro-Kann* defense.
- 11 *N-Check-Mate* tasks (putting the enemy king in a position of checkmate in N moves, N ranges from 1 to 6) tasks.

In this corpus, 9 players were active expert players with *Elo* ratings¹ ranged from 1759 to 2150 ($M = 1937.3$; $SD = 123.2$). For the intermediate players, the *Elo* ratings ranged from 1100 to 1513 ($M = 1380.3$; $SD = 100.5$). 6 players were casual players who were not currently playing in club. The average recording time per participant is 13:35 minutes ($MIN = 4:54$, $MAX = 23:54$, $SD = 5:02$).

The saliency maps of the openings are not very informative. Common openings are repetitive and well known by players. They tend to follow more their intuition at this point of the game rather than reasoning. Thus, eye tracking data does not provide satisfactory information for our prediction task and we decided to focus our study on the 11 *N-Check-Mate* tasks that require more cognitive processes to be solved.

Visual attention differs among chess players and depends on their level of expertise [Charness et al., 2001]. These differences are reflected in saliency maps. The primary experiments using these data exhibited that the saliency maps are too sparse to let a neural network converge. As it is commonly done in saliency map prediction from natural images [Banner et al., 2018], we decided to group each task into one saliency map, averaging all saliency maps from all players computed from eye tracking data. This computation provides a mean representation closer to most players’ visual attention. Doing so, the current corpus ends-up with 11 configurations and the same amount of saliency maps. With such limited data, we needed ways to augment it while keeping data interesting and relevant. This corpus will be referenced throughout the rest of the article as the *ETDS* dataset.

4.1.2 Data augmentation

Data augmentation is about creating variations of the dataset, to multiply the examples the network will learn on and help it generalizing its knowledge to more and more complex examples. Usually, this includes many transformations such as rotating, translating, flipping, scaling or cropping a part of the image, or even adding some noise to the input in some adversarial training paradigms. Several augmentation techniques are relevant in our context:

- The first augmentation is to cut the board into smaller pieces, similarly for the corresponding saliency map. So we let the network focus on specific parts of the board, learning features at different scales. We used a sliding window of 3x3, 4x4 and 5x5 squares on the chess board, creating various sub-views of the configuration. This augmentation multiplies the amount of data by 150. The goal here is to make the network learn smaller and more precise features at the scale of few cells.

¹ The *Elo* system is a method to calculate rating for players based on tournament performance. Ratings vary between 0 and approximately 2850. https://en.wikipedia.org/wiki/Elo_rating_system (last seen 09/2018)

- The second data augmentation consists in inverting the color of each piece (black to white and *vice versa*). Because nearly every configuration is possible while playing white and black, inverting colors adds diversity to the configurations the network encounter while training. The amount of training data is multiplied by 2.
- The last augmentation changes the orientation of the board to the point of view of the other player. The underlying hypothesis is that, at any time, a chess player exploring his adversary possibilities looks at the board the same way the initial player does. Again, this augmentation increases the available training data by a 2 factor.

While remaining small in regards of usual state-of-the-art corpora, the final *ETDSa* (for augmented *ETDS*) contains 6600 examples.

4.2 Top-Down dataset generated from games (*GDS*)

As stated in previous section, Top-Down approaches in visual attention prediction focus on relevant information in an image depending on the task to accomplish. In chess, the task is well defined as it is to move, for a specific configuration, the piece that is most satisfying a condition which could be to win the game, to put the king in check, etc. Because we are working with visual attention and saliency maps, we do not want to train a network to play chess (unlike Alphazero [Silver et al., 2017]) but to generate saliency maps depending on the pieces chess players consider to play at each round, helping the network to learn relevant pieces for a given configuration.

From the beginning of the project, we used Lichess², an online and offline chess platform. It allows to create personalized games and to play against other players or a computer. With thousands of registered players and millions of games played every month, this is a gold mine to get training data. They grant access to a monthly database in the PGN format (Portable Game Notation)³. This database contains all games played during the past month (more than 22 Million for March 2018). Because we are working on visual attention with images, we had to transform these PGN games into images. We developed a tool which reads games and, for each turn, creates the image board as seen from the white and black perspectives.

At this point, we want to generate saliency maps of chess players' visual attention. We need to find a way to generate relevant and accurate saliency maps. We based the processing on three basic hypotheses. The first one being that a chess player is at least looking at two things: the cell where lies the piece he intends to move and its destination. As we know from the PGN data which piece is going to move, we highlighted in the saliency maps these two areas on the board. The second hypothesis is that chess players are quickly following the path taken by the moved piece to check collision and accuracy of its destination. Even if this is not as salient as starting and ending fixations, it could be considered as a succession of eye saccades. We model it in the saliency map as a darker shade of gray showing its smaller relevancy. The last hypothesis is about moves changing the enemy king into check state, making it salient for the attacking player. Those hypotheses permit to create a dataset way larger than the bottom-up dataset (*ETDS*). With around two thousand games and an average of fifty round each, we created a bit more than one hundred and twenty thousand new images for training.

Even though this new dataset is substantially big and rich enough, some biases are present. The first one is that even with the large amount of created configurations, the vast majority of them were openings. Variety of these openings being limited, there are some redundancies. The second bias is about the subjectivity of the generated saliency maps. Our choices about salient areas were driven by our knowledge of the chess game, observations from our data and examples in the literature [Charness et al., 2001; Reingold and Charness, 2005; Sheridan and Reingold, 2017]. The last bias concerns the type of moved pieces. Pawns being the vast majority of them, they are often way more salient because they are moved more often. Queen, rooks, knights and bishops follow pawns in the salient order. The king being not very often moved nor the center of attention as it is not put in check that often, had a limited saliency even for configurations where it was in danger. Also, because of their key role in

²<https://en.lichess.org/> (last seen 12/2018).

³https://en.wikipedia.org/wiki/Portable_Game_Notation (last seen (12/2018)).

System	Pre-training		(Fine-)training	
	CAT2000	<i>GDS</i>	<i>ETDS</i>	<i>ETDSa</i>
V1		✓	✓	
V2	✓		✓	
V3		✓		✓
V4	✓			✓
V5				✓

Table 1: Training conditions of our system.

the game strategy, players do not need to check the kings’ position as often as other pieces (bishops, pawns etc..).

5 Experiments and results

This section presents results obtained using different training strategies. We are unaware of other research on visual attention prediction for chess (or other board games) players in the machine learning literature. Thus, it is not possible to provide comparative experimental performance evaluation. The results presented in this section can provide a baseline for future work on visual attention prediction in similar contexts.

5.1 Quantitative results

5.1.1 Evaluation conditions

To conduct the quantitative evaluation of our approach, we trained several variations of our system. The neural network architecture remains the same (see section 3), but the training conditions vary in terms of pretraining and fine training corpora. At our disposal, we have our 3 datasets *ETDS*, *ETDSa* and *GDS* (see section 4). We adjoin the CAT2000 dataset [Borji and Itti, 2015]. This dataset contains natural images with their saliency maps computed and averaged from eye movements of 120 people. The underlying idea of using it is to extract general saliency deep features before fine-tuning them on our specific corpora. The different versions of our training strategies are depicted in table 1.

The first 2 systems are trained on the *ETDS* dataset. The first one (V1) is pretrained on *GDS* and the second one (V2) on CAT2000. The next 2 systems (V3 and V4) are pretrained the same way but are fine-tuned using the augmented *ETDS* dataset. As a reference, the last system is trained on *ETDS* without pretraining. All systems are trained during 2000 epochs with a learning rate of 0.01. This learning rate has been determined experimentally during the development phase. It showed the best performance, even compared with a cyclical learning rate approach [Smith, 2017].

Concerning the usual training/validation/test sets paradigm, the size of the *ETDS* dataset restricts the available options. First, we cannot create a validation set without drastically reducing the size of the training set. As we do not want to train and to test on similar saliency maps (i.e. saliency maps on the same chess problem from chess players with about the same expertise), we split training and testing sets using tasks rather than using players. We kept 9 configurations as training data and used the 2 lasts, the more complex and difficult ones, as test configurations to assess the performance of our proposal. This splitting strategy ensures that the neural networks are not trained on the test data.

5.1.2 Evaluation metrics

To measure the quality of the generated saliency maps, we consider them as probability distributions. Information about available metrics can be found in Bylinskii and al. [Bylinskii et al., 2016] and Le Meur and Baccino [Le Meur and Baccino, 2013] papers, where the authors explain the interpretation that can be done from those metrics. In this section we will present the metrics used in this article.

Linear Correlation Coefficient (LCC). The linear correlation is a statistical method to determine how much two probability distributions are related one to the other. It is computed as the ratio of

Model/ Metrics	LCC (std)	SIM (std)	NSS (std)	AUC-Judd (std)	AUC-Borji (std)
V1	0.60±0.10	0.68±0.05	0.21±0.11	0.57±0.07	0.67±0.05
V2	0.66±0.06	0.70±0.03	0.23±0.13	0.61±0.13	0.70±0.08
V3	0.64±0.08	0.69±0.04	0.21±0.10	0.62±0.13	0.71±0.09
V4	0.57±0.10	0.65±0.08	0.24±0.09	0.69±0.09	0.74±0.10
V5	0.54±0.03	0.64±0.02	0.20±0.10	0.66±0.09	0.74±0.06

Table 2: Performances of the different flavors of our neural network, using a 6-fold cross validation on the 11 configurations.

the covariance of the two distributions and the product of their standard deviations. The value given by the linear correlation ranges between -1 and +1, the closer to +1 or -1 it gets, the more related the two distributions are. Because of its symmetric computation, linear correlation does not make the difference between false positives and false negatives.

Similarity (SIM). The similarity, also referred as histogram intersection, is the sum of minimums of two normalized saliency maps. A null value indicates that there is no overlap while 1 means that the saliency maps are the same.

Normalized Scanpath Saliency (NSS). NSS was introduced by Peters et al. [2005] as a correspondence measure between a normalized saliency map and a ground truth at fixation points. It is sensitive to false positives, relative differences in saliency across the image and general monotonic transformations. A mean subtraction step is added to gain robustness against any kind of linear transformations. With a value above 0, this metric expresses a correspondence above chance between the two saliency maps, below 0 an anti-correspondence.

The last metric employed in this paper is the *Area Under the Curve (AUC)*. In visual attention modeling, the main goal is to predict if a pixel is going to be stared at or not, which is basically a binary classification. The ROC (Receiver Operating Characteristic) is in signal detection theory a measure of the trade-off between true and false positives while varying a threshold [Fawcett, 2006]. For saliency evaluation, we consider the saliency map as a binary classifier, varying the threshold to transform probabilities to {0,1} values. The AUC metric is then the area under the ROC curve. Some variants exist in literature: *AUC-Judd* [Judd et al., 2009, 2012; Bylinskii et al., 2016] and *AUC-Borji* [Borji et al., 2013]. The difference resides in the way the *AUC-Borji* metric computes the false positive rate as the ratio of values above a threshold sampled at random pixels (as many as fixations, sampled uniformly over the whole image) divided by the number of saliency map pixels at the same threshold.

5.1.3 Performance assessment

The quantitative results of the different training strategies of the proposed neural architecture are presented in table 2. Looking at this table, one can first draw some general comments. All systems are better than chance ($NSS > 0$) and similarity between the saliency maps and the ground-truth is always over 0.54. As the correlation ranges from 0.54 (V5, only trained on the non augmented *ETDS* corpus) to 0.66 (V2), the LCC metric confirms that our approach is promising.

As a second evaluation point, one can compare the effectiveness of the different (pre)training strategies. Without pretraining and augmentation, V5 is the worst model except for AUC metrics. V2 and V4 are the best models looking globally at all metrics. Nevertheless, V4, pretrained with CAT2000 and then *ETDSa*, presents better results on AUC metrics, meaning that it has a better discrimination power between fixated/unfixated pixels. As expected, pretraining our model on the CAT2000 dataset appears to give better results (V4) even without data augmentation (V2). Concerning correlation metrics (LCC, SIM, NSS), V4 is not as good as without augmentation (V2) but offers the better performances concerning AUC metrics.

These results show the effectiveness of our investigation. Systems pretrained with CAT2000 have better general scores. This important result validates that deep features computed on natural image saliency maps are interesting to process synthetic images like our chess board.

5.2 Qualitative results

The final evaluation consists in qualitative results, i.e. human evaluation of generated saliency maps. This evaluation will first focus on actual chess problems, then on pure synthetic examples, comparing the best systems (V3 and V4) of each training strategy in the qualitative evaluation.

5.2.1 Comparative results

Saliency map prediction from V3 and V4 models given two unknown chess configurations are depicted on figure 4 (first row). On the second row, the figure displays ground-truth results computed from eye tracking data of chess players for these configurations. Third and fourth rows present respectively saliency map predictions for systems V3 and V4. After analysis, several differences can be observed between the two training strategies and the ground-truth. On one hand, for configuration 1, the network trained on CAT2000 correctly focuses on both kings, rooks and the bishop as actual chess players did. On the other hand, the other network does not show as much interest for the rook and bishop on the right side of the board. For the second configuration, both networks correctly focus on the white king, the rook and the white queen. The black queen and black rook which are the keys to the check mate are salient also for both networks.

5.2.2 Prediction on synthetic examples

In this section, we will qualitatively assess performances of our networks on several unseen synthetic configurations. Configurations from figures 5 to 8 were created using the board editor from the Lichess platform.

The figure 5 depicts a simple configuration. Here the black king is put in check by the white queen. We would expect fixations on the two pieces and on the vertical line between them showing their relation. V3 is close to this expected result but with some salient areas on the left part of the board. V4 shows less saliency between the two pieces but some on the right part of the board. A more complicated configurations result is drawn on Figure 6, showing the white king in check. The only possibility is to move the king. The expected salient areas would be the white rooks, white king and white queen and the black knight if we consider the king in check. If we do not know the check rules, the white queen diagonal and rook’s vertical lines would be salient as they can attack the enemy king this way. Here V3 considers all pieces and the position from where the queen would put the enemy king in check. But rook’s verticals are not considered. V4 is considering moving the right rook upward and the black knight putting the king in check with more saliency. Figure 7 is different from previous examples considering the number of pieces. The white bishop in B5 is in danger because of the black pawn in C6. The logic way to solve this problem is to move the bishop either back or to the left. V3 considers the area surrounding the king/queen of the enemy and including the white bishop (the king was in check the round before). V4 considers the bottom right region (castling area) to be more salient but do not include the white bishop. The last configuration (Figure 8) shows a white knight having two choices. It can either take the pawn and be taken by the queen or take the queen and take the pawn next round. We would expect the area around the three pieces to be salient as well as the space between them. V3 saliency map corresponds to these assumptions choosing the queen to be more salient with the knight. V4 chooses also the queen and the knight but is less focused. Unexpectedly, it adds noisy salient areas everywhere on the board.

These qualitative results offer a glance at what deep neural networks could bring to visual attention prediction for chess players. As in quantitative results (see previous section), training with the CAT2000 or on synthetic *GDS* dataset provides different results. This suggests again that deep features learned on natural images are relevant in our saliency computation for chess board images. It also highlights that the generated *GDS* corpus is of interest to learn relationships between pieces.

6 Discussion

As seen in previous section, the different flavors of our neural network show interesting results and generalizations capability on unseen configurations. For specific configurations, one can see that pieces

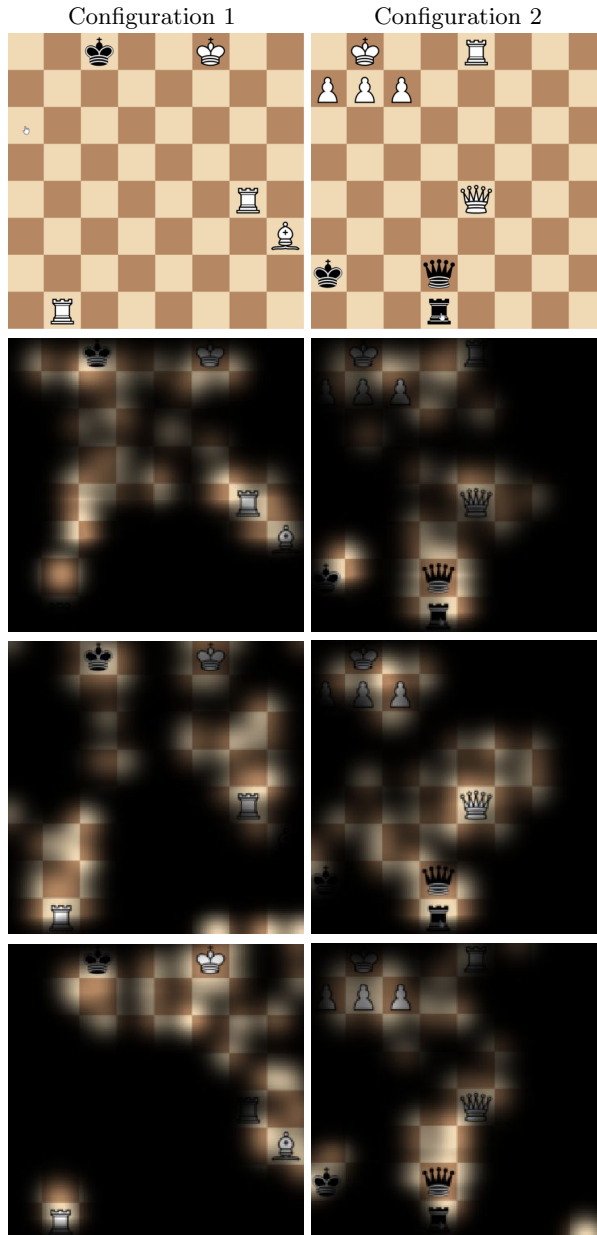


Figure 4: Saliency map generation of two models V3 and V4 on two testing configurations. The first row is the chess configuration, the second row the ground-truth from the users’ eye tracking data. The third row contains the predictions from V3 pretrained on *GDS*, the last row from V4 pretrained on CAT2000.

and their relationship influence the model prediction as it would do for a chess player, sustained by training on the *ETDSa* and *GDS* corpora. Using a well-known visual attention prediction dataset (CAT2000) [Borji and Itti, 2015] to train the decoder part to transform features to salient regions, led to better performances on standard metrics and results closer to the ground-truth. However, it was a bit struggling to generalize to new configurations, displaying some salient areas around the boundaries of the board. Nonetheless, pretraining the decoder part of the network with such dataset seems to be a good way to initialize the network weights. This may be explained by the fact that the network encoder is initialized using VGG weights trained on the ImageNet dataset [Simonyan and Zisserman, 2014].

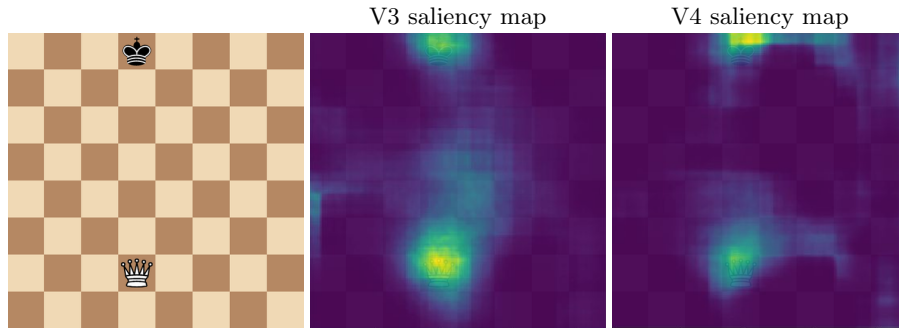


Figure 5: An other simple configuration with only one king and its opponent queen

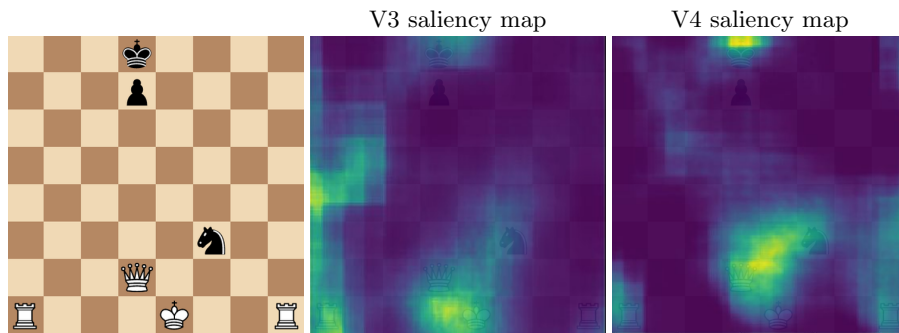


Figure 6: A bit more complicated configuration. The white king is in check

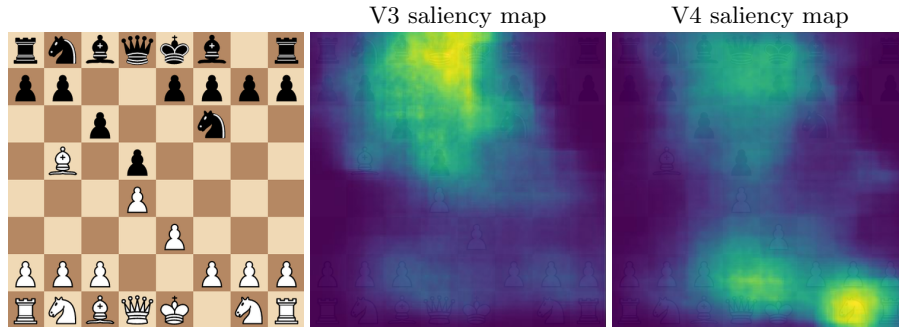


Figure 7: An opening with a bishop in danger

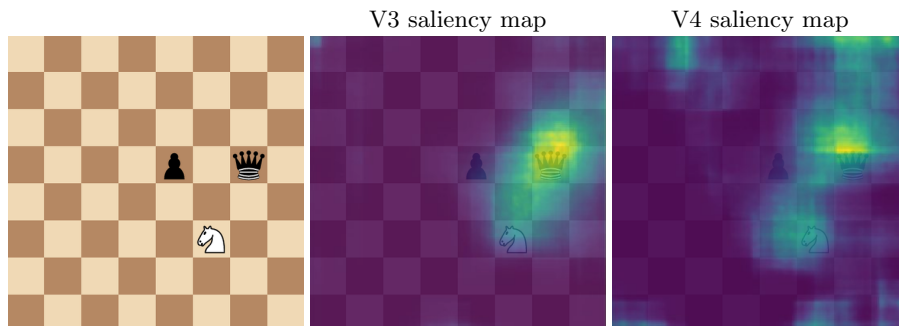


Figure 8: A choice for a knight

The reported results may hide that there are still rooms for improvement, the underlying bottleneck of these researches on chess remaining the amount of available data. The available recordings of eye tracking data in chess context [Guntz et al., 2018] represent a huge amount of work and hours of data. Nevertheless, they contain few configurations and around 30 players recorded in actual chess problem solving tasks. This is not enough for generalization and performance, moreover when using deep learning techniques. This explains the benefit from pre-training on the CAT2000 corpus and from our generated game dataset (*GDS*).

The first way to improve current results would be to record eye tracking data for players playing full games during online tournaments or even when casually training for instance. With such data, the training corpus will contain more configurations and get more eye-tracking data and from a larger panel of players' expertise. With more data, it will be possible to create different deep learning networks based on the players' skill. It has been shown in previous work that for board games such as chess, players do not look at the game the same way depending on their expertise [Charness et al., 2001; Reingold and Charness, 2005; Sheridan and Reingold, 2017]. Doing this separation, one could have different networks pretrained on the whole data, but adjusted on data from different player groups (beginner, intermediates, experts for instance). The resulting neural networks would be able to model visual attention for their respective expertise group.

Using the *GDS* dataset showed to be efficient for producing saliency map when no visual attention prediction dataset is available. Previous depicted results let us assess some points about it:

- There is a content bias. This *GDS* corpus contains more openings, so there are more salient areas on the player side. This could however be partially overthrown after training on the augmented dataset (*ETDSa*).
- The results from the network pretrained on this dataset are showing smaller salient areas focused on pieces. This is a direct consequence of how the dataset has been created, but is interesting as it corresponds more to how beginner/intermediate players would look at the board.
- The network seems to recognize more pieces and their relations when pretrained using this dataset.

The *GDS* corpus can benefit from the immense number of chess games available online. A second way to improve these researches would be to enhance the generation process of this dataset (depicted in section 4). The game dataset generation can be improved by defining a strategy to select interesting parts of games to produce a more balanced corpus, maybe integrating players' statistics to produce saliency maps reflecting what is known about the link between fixations and expertise.

7 Conclusion

Training a CNN for visual attention prediction is deeply task dependent and thus is not a solved problem notably in chess context. The promising quantitative and qualitative result reported here show that using data augmentation and task driven data, visual attention prediction of chess players is possible. The proposed network architecture with an auto-encoder is able to learn and to predict saliency map in the chess context. When focusing on a Bottom-up approach using raw images as input, the model was able to extract features and patterns from limited data. Data augmentation was found to improve performances when pretraining on the synthetic *GDS* dataset, enhancing locations of salient areas rather than extracting correspondence between the saliency maps and the ground-truth, learning smaller and more general patterns. After gathering more chess games to cope with the lack of eye-tracking data, we should be able to use the proposed neural network architecture as input for researches on analyzing and interpreting multimodal signals of players engaged in solving challenging problems.

In this study, we have opened the investigation concerning visual attention prediction for chess players. It may be possible to apply similar approaches to other board games such as GO, where the model used and its training would be the same, with just a variation of the data and game rules. The proposed visual attention model showcases good performances even if there are still rooms for improvement, as working on quality and quantity of our training corpus. The novelty of our work makes it difficult to compare to existing techniques, but current results and ideas could be used as a baseline for future work.

8 Acknowledgments

This research has been funded by the French ANR project CEEGE (ANR-15-CE23-0005), and was made possible by the use of equipment provided by ANR Equipment for Excellence Amiqua4Home (ANR-11-EQPX-0002).

References

- Kévin Bannier, Eakta Jain, and Olivier Le Meur. 2018. Deepcomics: Saliency Estimation for Comics. In *Proceedings of the 2018 ACM Symposium on Eye Tracking Research & Applications (ETRA '18)*. ACM, New York, NY, USA, Article 49, 5 pages. <https://doi.org/10.1145/3204493.3204560>
- Muhammed Abdullah Bülbül. 2012. *Visual attention models and applications to 3D computer graphics*. Ph.D. Dissertation. Bilken University.
- Ali Borji and Laurent Itti. 2015. CAT2000: A Large Scale Fixation Dataset for Boosting Saliency Research. *CoRR* abs/1505.03581 (2015). arXiv:1505.03581 <http://arxiv.org/abs/1505.03581>
- A. Borji, H. R. Tavakoli, D. N. Sihite, and L. Itti. 2013. Analysis of Scores, Datasets, and Models in Visual Saliency Prediction. In *2013 IEEE International Conference on Computer Vision*. 921–928. <https://doi.org/10.1109/ICCV.2013.118>
- Vincent Buso, Jenny Benois-Pineau, and Jean-Philippe Domenger. 2015. Geometrical cues in visual saliency models for active object recognition in egocentric videos. *Multimedia Tools and Applications* 74, 22 (01 11 2015), 10077–10095. <https://doi.org/10.1007/s11042-015-2803-2>
- Zoya Bylinskii, Tilke Judd, Aude Oliva, Antonio Torralba, and Frédo Durand. 2016. What do different evaluation metrics tell us about saliency models? *CoRR* abs/1604.03605 (2016). arXiv:1604.03605 <http://arxiv.org/abs/1604.03605>
- Neil Charness, Eyal M. Reingold, Marc Pomplun, and Dave M. Stampe. 2001. The perceptual aspect of skilled performance in chess: Evidence from eye movements. *Memory & Cognition* 29, 8 (01 12 2001), 1146–1152. <https://doi.org/10.3758/BF03206384>
- Stefic Daria. 2016. *Learning saliency for human action recognition*. Ph.D. Dissertation. Queen Mary University of London.
- J. Deng, W. Dong, R. Socher, L.J. Li, K. Li, and L. Fei-Fei. 2009. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*.
- Lijuan Duan, Jili Gu, Zhen Yang, Jun Miao, Wei Ma, and Chunpeng Wu. 2014. Bio-inspired Visual Attention Model and Saliency Guided Object Segmentation. In *Genetic and Evolutionary Computing*, Jeng-Shyang Pan, Pavel Kromer, and Vaclav Snavel (Eds.). Springer International Publishing, Cham, 291–298.
- Tom Fawcett. 2006. An introduction to ROC analysis. *Pattern Recognition Letters* 27, 8 (2006), 861 – 874. <https://doi.org/10.1016/j.patrec.2005.10.010> ROC Analysis in Pattern Recognition.
- Iván González-Díaz, Hugo Boujut, Vincent Buso, Jenny Benois-Pineau, and Jean-Philippe Domenger. 2013. Saliency-based object recognition in video. *jhal-00799127* (2013).
- Guanghua Gu, Jiaojiao Zhu, and Zexia Liu. 2015. Visual Saliency Detection Based Object Recognition. *Journal of Information Hiding and Multimedia Signal Processing* 6, 6 (11 2015).
- Thomas Guntz, Raffaella Balzarini, Dominique Vaufreydaz, and James L. Crowley. 2018. Multimodal Observation and Classification of People Engaged in Problem Solving: Application to Chess Players. *Multimodal Technologies and Interaction* 2, 2 (March 2018). <https://doi.org/10.3390/mti2020011>

- X. Huang, C. Shen, X. Boix, and Q. Zhao. 2015. SALICON: Reducing the Semantic Gap in Saliency Prediction by Adapting Deep Neural Networks. In *2015 IEEE International Conference on Computer Vision (ICCV)*. 262–270. <https://doi.org/10.1109/ICCV.2015.38>
- Tilke Judd, Frédo Durand, and Antonio Torralba. 2012. A Benchmark of Computational Models of Saliency to Predict Human Fixations. In *MIT Technical Report*.
- T. Judd, K. Ehinger, F. Durand, and A. Torralba. 2009. Learning to predict where humans look. In *2009 IEEE 12th International Conference on Computer Vision*. 2106–2113. <https://doi.org/10.1109/ICCV.2009.5459462>
- Srinivas S. Kruthiventi, Kumar Ayush, and R. Venkatesh Babu. 2015. DeepFix: A Fully Convolutional Neural Network for predicting Human Eye Fixations. *CoRR* abs/1510.02927 (2015). arXiv:1510.02927 <http://arxiv.org/abs/1510.02927>
- S. Kullback and R. A. Leibler. 1951. On Information and Sufficiency. *The Annals of Mathematical Statistics* 22, 1 (1951), 79–86. <http://www.jstor.org/stable/2236703>
- Guillaume Lavoué, Frédéric Cordier, Hyewon Seo, and Chaker Larabi and. 2018. Visual Attention for Rendered 3D Shapes.
- Olivier Le Meur and Thierry Baccino. 2013. Methods for comparing scanpaths and saliency maps: strengths and weaknesses. *Behavior Research Methods* 45, 1 (01 03 2013), 251–266. <https://doi.org/10.3758/s13428-012-0226-9>
- Robert J. Peters, A. M. Iyer, Laurent Itti, and Christof Koch. 2005. Components of bottom-up gaze allocation in natural images. *Vision Research* 45 (2005), 2397–2416.
- M.I. Posner and S.E. Petersen. 1990. The Attention System of the Human Brain. *Annu. Rev. Neurosci* (1990), 13:25–42.
- E. Reingold and N. Charness. 2005. Perception in chess: Evidence from eye movements. (28 07 2005). <http://www.oxfordscholarship.com/view/10.1093/acprof:oso/9780198566816.001.0001/acprof-9780198566816-chapter-14>.
- Marjaneh Safaei and Hassan Foroosh. 2017. Single Image Action Recognition by Predicting Space-Time Saliency. *CoRR* abs/1705.04641 (2017). arXiv:1705.04641 <http://arxiv.org/abs/1705.04641>
- Heather Sheridan and Eyal M. Reingold. 2017. Chess players’ eye movements reveal rapid recognition of complex visual patterns: Evidence from a chess-related visual search task. *Journal of Vision* 17, 3 (2017), 4. <https://doi.org/10.1167/17.3.4> arXiv:/data/journals/jov/936104/i1534-7362-17-3-4.pdf
- David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dhharshan Kumaran, Thore Graepel, Timothy P. Lillicrap, Karen Simonyan, and Demis Hassabis. 2017. Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm. *CoRR* abs/1712.01815 (2017). arXiv:1712.01815 <http://arxiv.org/abs/1712.01815>
- Karen Simonyan and Andrew Zisserman. 2014. Very Deep Convolutional Networks for Large-Scale Image Recognition. *CoRR* abs/1409.1556 (2014). arXiv:1409.1556 <http://arxiv.org/abs/1409.1556>
- Leslie N Smith. 2017. Cyclical learning rates for training neural networks. In *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 464–472.
- Wenguan Wang and Jianbing Shen. 2017. Deep Visual Attention Prediction. *CoRR* abs/1705.02544 (2017). arXiv:1705.02544 <http://arxiv.org/abs/1705.02544>

Yang Xu, Jun Li, Jianbin Chen, Guangtian Shen, and Yangjian Gao. 2017. A novel approach for visual Saliency detection and segmentation based on objectness and top-down attention. In *2017 2nd International Conference on Image, Vision and Computing (ICIVC)*. 361–365. <https://doi.org/10.1109/ICIVC.2017.7984578>