



Towards Improving the Understanding of Image Semantics by Gaze-based Tag-to-Region Assignments

Tina Walber
Ansgar Scherp
Steffen Staab

Nr. 8/2011

**Arbeitsberichte aus dem
Fachbereich Informatik**

Die Arbeitsberichte aus dem Fachbereich Informatik dienen der Darstellung vorläufiger Ergebnisse, die in der Regel noch für spätere Veröffentlichungen überarbeitet werden. Die Autoren sind deshalb für kritische Hinweise dankbar. Alle Rechte vorbehalten, insbesondere die der Übersetzung, des Nachdruckes, des Vortrags, der Entnahme von Abbildungen und Tabellen – auch bei nur auszugsweiser Verwertung.

The “Arbeitsberichte aus dem Fachbereich Informatik“ comprise preliminary results which will usually be revised for subsequent publication. Critical comments are appreciated by the authors. All rights reserved. No part of this report may be reproduced by any means or translated.

Arbeitsberichte des Fachbereichs Informatik

ISSN (Print): 1864-0346

ISSN (Online): 1864-0850

Herausgeber / Edited by:

Der Dekan:
Prof. Dr. Grimm

Die Professoren des Fachbereichs:

Prof. Dr. Bátori, Prof. Dr. Burkhardt, Prof. Dr. Diller, Prof. Dr. Ebert,
Prof. Dr. Furbach, Prof. Dr. Grimm, Prof. Dr. Hampe, Prof. Dr. Harbusch,
jProf. Dr. Kilian, Prof. Dr. von Korflesch, Prof. Dr. Lämmel, Prof. Dr. Lautenbach,
Prof. Dr. Müller, Prof. Dr. Oppermann, Prof. Dr. Paulus, Prof. Dr. Priese,
Prof. Dr. Rosendahl, jProf. Dr. Scherp, Prof. Dr. Schubert, Prof. Dr. Staab, Prof. Dr.
Steigner, Prof. Dr. Sure, Prof. Dr. Troitzsch, Prof. Dr. Walsh, Prof. Dr. Wimmer, Prof.
Dr. Zöbel

Kontakt Daten der Verfasser

Ansgar Scherp, Tina Walber, Steffen Staab
Institut WeST
Fachbereich Informatik
Universität Koblenz-Landau
Universitätsstraße 1
D-56070 Koblenz
EMail: scherp@uni-koblenz.de, walber@uni-koblenz.de, staab@uni-koblenz.de

Towards Improving the Understanding of Image Semantics by Gaze-based Tag-to-Region Assignments

Tina Walber

WeST, University of Koblenz-Landau,
Germany
walber@uni-koblenz.de

Ansgar Scherp

WeST, University of Koblenz-Landau,
Germany
scherp@uni-koblenz.de

Steffen Staab

WeST, University of Koblenz-Landau, Germany
staab@uni-koblenz.de

May 16, 2011

Abstract

Eye-trackers have been used in the past to identify visual foci in images, find task-related image regions, or localize affective regions in images. However, they have not been used for identifying specific objects in images. In this paper, we investigate whether it is possible to assign image regions showing specific objects with tags describing these objects by analyzing the users' gaze paths. To this end, we have conducted an experiment with 20 subjects viewing 50 image-tag-pairs each. We have compared the tag-to-region assignments for nine existing and four new fixation measures. In addition, we have investigated the impact of extending region boundaries, weighting small image regions, and the number of subjects viewing the images. The paper shows that a tag-to-region assignment with an accuracy of 67% can be achieved by using gaze information. In addition, we show that multiple regions on the same image can be differentiated with an accuracy of 38%.

1 Introduction

To describe the semantics of images on social media platforms such as Flickr¹ users can allocate tags to the images. Nevertheless, tagging describes the semantics of images in a limited way. One step towards improving the understanding of image semantics is to annotate specific image regions instead of the entire image. To conduct such a region-based annotation, different approaches can be applied: The manual annotation of image regions like in the LabelMe data set² is time-consuming. Automatic image segmentation and annotation [13, 1] has many limitations and requirements. Eye-trackers, as new input devices, provide for the users' gaze paths, which are used to advance the specific task-related understanding of images [15, 5, 12]. However, current research does not tackle the identification of objects in images.

In this paper, we investigate the possibilities to improve the understanding of image semantics by using gaze information in the context of tagging. The objective is to assign tags to image regions by analyzing the users' gaze paths. However, establishing such tag-to-region-assignments using gaze information is a complex task. A pre-experiment with 12 subjects tagging images on the photo sharing platform Flickr has shown that the gaze paths are quite different. In order to analyze the gaze paths in a more controlled manner, we have designed an experiment in which 20 subjects have viewed a sequence of 50 tag-image-pairs each. First, a tag is shown to the subjects and subsequently the image. For each tag-image-pair, the subjects had to decide whether or not an object described by the tag is shown on the image. We have recorded the gaze paths and applied nine existing and four new fixation measures to compare their performance on finding tag-to-region assignments. Fixations are short stops in the movements of the eyes, which are briefly focused on a particular point on the screen. These are phases of high visual perception. A fixation measure is a function on the users' gaze path and is calculated for each image region. For the fixation measure with the highest number of correct tag-to-region assignments, we further investigate the impact of extending region boundaries and applying a linear weighting function to support smaller image regions. Finally, the influence between the number of subjects viewing an image and the precision of the tag-to-region assignment is analyzed. As our results show, an assignment of tags to image regions can be achieved at an accuracy of 67%. We have also investigated if different image regions on the same image can be differentiated. In our experiments, we have achieved an accuracy of 38% for distinguishing two different image regions.

The remainder of this paper is organized as follows: The related work on annotating image regions is discussed in the subsequent section. The

¹<http://www.flickr.com/>

²<http://labelme.csail.mit.edu/>

design of our experiment is described in Section 3. The detailed analysis of the users' gaze information for determining a tag-to-region assignment is conducted in Section 4. A summary of the evaluation results is provided in Section 5, before we conclude the paper.

2 Related Work

Different approaches have been pursued to advance the understanding of images by annotating image regions. They deal with the *manual labeling* of image regions, *automatic labeling* based on segmentation of images into regions or the use of visual similarity for tag recommendation, and finally analyzing *gaze information* obtained from eye-trackers.

Manual labeling of image regions is the simplest solution to improve the annotation of images. For example, the photo sharing platform Flickr allows its users to manually mark image regions by drawing rectangle boxes on it and writing a comment to it. Other web platforms like LabelMe [14] allow the more precise creation of regions by drawing polygons on images. These regions can be annotated with a tag. The same principle is used for a “Games with a purpose” called Squigl³. It triggers the human play instinct in order to obtain image regions [16]. Two randomly selected users team up to mark a region on the same image. The users score when the region traces match.

With respect to the *automatic labeling*, Rowe [13] presents an approach to find the visual focus of an image by applying image processing in terms of segmentation and low-level features. Goal is to link the visual focus with the image caption. This approach is well-suited for images with a single object[13]. However, it has many limitations concerning the position and characteristics of the shown object. Regarding the use of visual similarity for tag recommendation, we find an approach by Li et al. [9] that recommends tags for an unlabeled image by using low-level similarity with already tagged images and obtaining relevant tags from these images. Identifying objects in images based on computer vision is a challenging task. A large amount of training data — consisting of images and labeled image regions — is needed to deliver good results (e.g. [1]).

Gaze information has been mainly used for image retrieval. Klami et al. [7] asked users to decide if at least one image of an array of four images fits the search task “sports”. They were able to identify individual images from the array that belongs to the task by means of analyzing the users' eye movements. They reach a higher accuracy than a random selection of images does. Similar, GaZIR [8] is an image retrieval system leveraging implicit user feedback from eye movements. A comparison with explicit user feedback by clicking on relevant images and a random baseline are promising. However,

³<http://www.gwap.com/squigl-a/>

they also show that the results are very noisy and for some images even poor[8]. Pasupa et al. [10] apply a support vector machine (SVM) algorithm using eye-tracking information together with content-based features to rank images. Hardoon and Pasupa [3] have recently extended this approach by using images with gaze data as training set for ranking images when no eye-tracking data is available. The ranking is conducted using tensor kernels in a SVM. Hajimirza and Izquierdo [2] developed an image retrieval system such as the previous ones and used it for image annotation. Relevant images are obtained from a list of images displayed to the user in a certain search context. However, no information with respect to image regions is obtained or created.

Privitera and Stark [11] have investigated the detection of image regions by looking at the visual foci obtained from users' gaze data. They compare the image regions identified by the gaze data with different image processing algorithms in order to automatically identify such regions. Goal is to predict the users gaze path by using image processing. Also this work in principle associates image regions with fixation points from the gaze path, Privitera and Stark just relate the two without being interested in labeling the identified image regions like we aim at. Santella et al. [15] present a method for semi-automatic image cropping using gaze information in combination with image segmentation. However, the cropped area is not further analyzed or annotated. Klami et al. [5, 6] present an approach to identify image regions relevant in a specific task using gaze information, its combination with low-level features, and low-level features only. Based on several users' gaze paths, heat maps are created that identify the regions in the image that are of importance in a given task. The work also revealed that the region identified depends on the task given to the subject before viewing the image. Jaimes et al. [4] carried out a preliminary analysis of identifying common gaze trajectories in order to classify images into five, predefined semantic categories. These semantic categories are handshake, crowd, landscape, main object in uncluttered background, and miscellaneous. The general assumption is that similar viewing patterns occur when different subjects view different images in the same category. To this end, a generic object-definition model is provided that allows the users to specify the relation of objects in the images like persons, hands, and so in an image showing a handshake situation. The results are encouraging and they determine that it may be possible to construct an automatic image category classifier from the approach. However, constructing the object-definition model is tedious and the number of classes is limited. In addition, an object classifier needs to be provided for each object category in the definition model in order to actually be able to classify new images. Finally, the work of Ramanathan et al. [12] aims at localizing affective objects and actions in images by using gaze information. Thus, the image regions that are affecting the users are identified and correlated with given concepts from an affection model. The affective image regions are

identified using segmentation and recursive clustering of the gaze fixations. However, also a general labeling of image regions showing specific objects is not conducted.

The related work shows that it is in principle possible to relate image regions with gaze path information. However, they are limited with respect to their general applicability as they are focused on a specific, small set of concepts in a specific task and they do not aim at providing a general solution for assigning image regions with tags for the regions. In order to provide such a solution, gaze information and in particular the fixations of the users' eye movements need to be analyzed and put in relation with image regions.

3 Experiment Design

Goal of our work is to investigate the possibilities to assign tags to image regions by analyzing the users' gaze information. From a preliminary experiment, we have learned that establishing such tag-to-region-assignments from users' gaze paths is a complex task. In this experiment, we have asked 12 subjects to tag 12 images on the photo sharing platform Flickr. The gaze paths of the subjects differ, i.e., some subjects constantly switch between looking at the image and entering a tag, some occasionally change the gaze, whereas others look at the image first and then type in all tags.

Thus, the setup of our experiment had to be designed such that the users' gaze paths are obtained in a more controlled manner. In our experiment application, we show existing tags to the subjects instead of asking them to enter own tags. In addition, the experiment application is designed such that first a tag and subsequently an image is shown to the subjects. The subjects were asked to decide whether or not an object described by the tag is shown on the image.

In Sections 3.1 and 3.2, we describe the subjects and data set used in our experiment before we explain the detailed experiment setup and design of our experiment application in Section 3.3. A brief experiment evaluation regarding efficiency, effectiveness, and satisfaction is conducted in Section 3.4. It provides information regarding the validity of our experiment. However, central part of our experiment evaluation is the detailed gaze path analysis described in Section 4.

3.1 Subjects

20 subjects (4 female) have participated in our experiment. The age of the subjects is between 23 to 40 years (average: 29.6 years). Their professions are undergraduate students (6), PhD students (12), and office clerks (2).

All subjects were familiar with tagging like on the photo sharing platform Flickr.

3.2 Data Set

As data set, we use LabelMe⁴ with 182.657 user contributed images (downloaded August 2010). The LabelMe community has manually created image regions by drawing polygons into the images and tagging them (see related work in Section 2). These manually created and annotated regions are used as ground truth in our experiment. The LabelMe simplification tool was applied to deal with synonyms and to remove labeling noise. By using this tool, terms describing the same objects (“person”, “person walking”, “pedestrian”) are summarized to one label and many spelling mistakes and abbreviations are removed.

For our experiment, we have randomly selected 50 images from the LabelMe data set. The images selected for our experiment have a minimum resolution of 1000x700 pixels and contain at least two labeled regions. For this image set, two different sets of tags were created, with one tag per image. So we got 100 different image-tag-pairs. The two tag sets are needed for the part of our analysis where we differentiate between several regions in the same image. 56 of the selected tags were “true”, that means an object described by the tag can actually be seen on the image. The other 44 tags were “false”. The true tags are obtained from the image themselves whereas incorrect tags are taken from other images in the LabelMe data set. We have manually replaced images from the dataset when a) the randomly selected false tags by coincidence correlate to some actually visible parts of the image and thus were true tags. We also replaced images where b) the tags were incomprehensible or expert knowledge is required and nonsense tags. In some cases there is c) a tag associated to a region like bicycle but multiple bicycles are depicted on the image and not all regions are explicitly marked as such. Thus, not all instances of the object the tag is referring to are actually labeled in the data set. Please note, that the selection does not mean that the removed images are not suitable for our approach per se. They just could not be used as ground truth in our experiment.

3.3 Experiment Setup

The setup of the experiment application consists of three steps as illustrated in Figure 1.

1. First, the tag together with the question “Can you see the following thing on the image?” is presented to the subjects (see Figure 1, left). After pressing the “space” button, the application continues with the next screen.
2. In this screen, a small blinking dot in the upper middle is displayed for one second (see Figure 1, middle). The subjects were asked to look

⁴<http://labelme.csail.mit.edu/>

at that point in order to let all subjects start viewing the images from the same position. The red dot is placed above the actual image that is shown in the third screen.

3. Finally, the image is shown to the subjects (see Figure 1, right). While viewing the image, the subjects were asked to judge whether the previously shown tag is true or false and making their decision by pressing the “y” (yes) or “n” (no) key.




Figure 1: Experiment Setup

The steps are repeated for each image-tag-pair and an additional first image-tag-pair. The first image-tag-pair is used to introduce the application to the subjects. The gaze path belonging to this first image-tag-pair is not used in the gaze analysis. Each subject evaluates one of the two sets consisting of 50 image-tag-pairs from the data set described above. True image-tag-pairs are mixed with false image-tag-pairs in order to keep the subjects concentrated. The image-tag-pairs have been selected such that each image is evaluated 10 times at the end of the experiment.

The subjects were told that the goal of the experiment is not to measure their efficiency in conducting the experiment task. They could take as much time as they liked to make a decision, but they were asked to press the “y” or “n” key once they have made their decision. Subsequently to the experiment, the subjects were asked to provide subjective feedback in a questionnaire. Besides recording the raw gaze data, we have also measured some quantitative data during the experiment. This data is the time the subjects took to make a decision per image and the chosen answer.

The experiment was performed on a screen with a resolution of 1680x1050 pixels. The subjects’ gaze was recorded with a Tobii X60 eye-tracker at a data rate of 60Hz. It has an accuracy of 0.5 degree. The experiment application was running in Microsoft’s Internet Explorer as a simple web application.

3.4 Experiment Application Evaluation

In this section, the effectiveness, efficiency, and satisfaction of our subjects while participating in our experiment is described.

3.4.1 Effectiveness

We have measured how many tag-to-region-pairs have been correctly identified by the subjects. Correctly identified means that a true tag is confirmed with “yes” and that a false tag is decided with “no” in our experiment application. In total, we have received 1000 answers, 10 answers per image-tag-pair. 5.7% of the given answers of all subjects were incorrect. The proportion of wrong answers is the same for true and false tags. The highest number of wrong answers for one image-tag-pair is 5, i.e., half of the users did not correctly identify whether the tag given is true or false. However, the high number of incorrect answers for single images is not an issue with our experiment as we only analyze the gaze paths of subjects having successfully identified a tag as true or false. Thus, we only consider image-tag-pairs with a true tag and where the subjects have given the correct answer.

3.4.2 Efficiency

The average answer time over all images is about 3,003 ms (shortest answer time was 204 ms and the longest was 25,163 ms). 50% of the answers were given in a time between 1,413 ms and 3,920 ms. For true tags, the average answer time over all subjects and all images is 2,818 ms, for false tags it was almost twice as long with 3,854 ms. Also the number of fixations is higher for false tags (13 fixations in average) than for true tags (9.6 fixations). This means that the subjects look longer and more precisely on images where there is no object related to the tag provided.

3.4.3 Satisfaction

We have asked the subjects to express the feelings they had during the experiment. The answers are provided on a 5-point-Lickert scale where a value of 1 means strongly disagree and a value of 5 stands for strongly agree. Concerning the statement “It was easy to decide on an answer.”, the subjects answered on average with a score of 3.85 (SD: 0.59). 15 subjects agree or strongly agree with the statement. Most of them also felt positive regarding the correctness of the answers they gave. 16 persons agree or strongly agree with the statement “I am confident that I have provided the right answer.” (average: 3.95, SD: 0.76). Most of the subjects felt comfortable during the evaluation (average: 4.4, SD: 0.75). 11 strongly agreed and 6 agreed to the statement. Thus, we assume that the results obtained from the experiment application are not influenced by side effects like users feeling discomforted in front of an eye-tracker.

4 Gaze Analysis

In this section, the detailed analysis of the eye-tracking data obtained from our experiment is presented. As said above, we have analyzed only the gaze paths for images with a true tag and for which the subject gave a correct answer. 547 gaze paths have been collected during the experiment that fulfill these requirements. 476 (87 %) of these gaze paths have at least one fixation inside or near the correct region. The preprocessing of the raw eye-tracking data was performed with the fixation filter offered by Tobii Studio with the default velocity threshold of 35 pixels and a distance threshold of 35 pixels. The extracted fixations are the base for our measure analysis.

Section 4.1 explains how we have calculated the precision of the tag-to-region assignments. The comparison of nine existing and four new fixation measures with respect to their performance for a tag-to-region assignment is presented in Section 4.2. In Section 4.3 and Section 4.4, the influence of extending the region boundaries and weighting of small image regions on the precision is considered and summarized in Section 4.5. The results are compared with two baselines in Section 4.6. The influence of the second screen in our experiment application with the blinking dot to let all subjects start viewing from the same position is investigated in Section 4.7. Section 4.8 shows that the accidental identification of favorite regions from gaze paths, not corresponding to this region, is rare. In Section 4.9, we examine if it is possible to identify several regions in one image. Typical characteristics of regions with correct vs. incorrect assigned tags are considered in Section 4.10. The influence aggregating gaze paths over several subjects is investigated in Section 4.11.

4.1 Calculating the Precision of Tag-to-Region-Assignments

The procedure for calculating the precision of the tag-to-region assignments is illustrated in Figure 2.

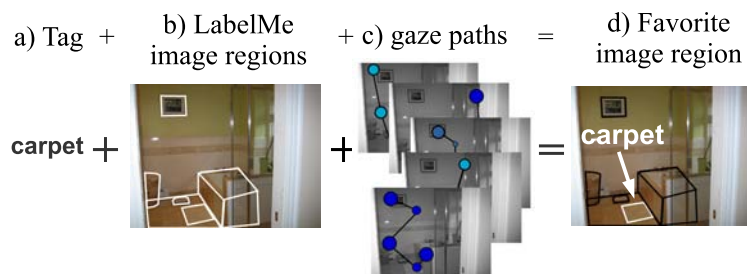


Figure 2: Overview of Calculating the Tag-to-Region-Assignments

The single steps performed for this calculation are:

1. For every LabelMe region in an image (b) a value for a fixation measure is calculated for every gaze path (c).
2. For every region, the measure results for every gaze path are summed up. From this, we obtain an ordered list of image regions for a fixation measure that determines the favorite region (d).
3. The label of the favorite region is compared with the tag (a) that was given to the subject in the experiment. If the label and tag match, the assignment is true positive (*tp*) otherwise it is false positive (*fp*). We have summed up the total number of correct and incorrect assignments over all images and calculate the precision P for the whole image set using the following formula:

$$P = \frac{tp}{tp + fp} \quad (1)$$

4.2 Comparison of Fixation Measures

Research question 1 *Which eye-tracking measure provides the best tag-to-region assignments?*

We compare nine existing and four new fixation measures. An overview of all fixation measures used in our comparison can be found in Table 1. To compare the measures, we calculate the precision P over all images for all gaze paths.

The measure `firstFixation` counts the number of fixations on the image before fixating on a region r . The favorite is the region that was fixated first. That means the region with no previous fixations on the image. We have used a modification of this measure to examine also the last fixations `fixationsAfter` [5]. 96% of the gaze paths have fixations after making a decision, due to the inherent reaction time of the experiment setup. The average duration of the recording after making the decision is 834 milliseconds. We have investigated the fixations around the moment of decision with the new measures `fixationsBeforeDecision` and `fixationsAfterDecision`. `fixationsAfterDecision` includes also fixations at the moment of decision. The `fixationDuration` describes the sum of the duration of all fixations on a region r . The Tobii measure `firstFixationDuration` considers the order of the fixations and describes the duration of only the first fixation on a region r . In addition, we have used the new measure `lastFixationDuration`. It provides the duration of the last fixation on the region. The standard measure `fixationCount` counts the fixations on a region r . A `visit` describes the time between the first fixation on a region and the next fixation outside. The three standard measures `maxVisitDuration`, `meanVisitDuration` and `visitCount` are based on visits. The last measure `saccLength` [8] provided good results for the relevance feedback in image search. Thus, we have also considered it in our experiments. The

assumption is that moving the gaze focus over a long distance (i.e., long saccade) to reach an image region r shows high interest in a region and makes the region a favorite.

Determining the favorite region depends on the kind of measure. For some measures, the area with the highest value in terms of time the subjects are looking at the region (e.g., `fixationDuration`) is the favorite region. For other measures, the lowest value defines the favorite region (e.g., `firstFixation`). For our analysis, only fixations on the image are considered. Fixations on the experiment page but outside the image are ignored.

Results We have received the best results for the measure `meanVisitDuration` with precision $P = 0.54$. That means, 54% of the image regions selected by the gaze analysis belonged to the tag that was shown to the subjects in advance. Two measures reach the second best value ($P = 0.53$): `fixationsBeforeDecision` and `lastFixationDuration`. With $P = 0.5$ the measure `fixationDuration` provides the third best result. The lowest precision values are $P = 0.21$ and $P = 0.26$ for `firstFixation` and `secondFixation`. In Figure 3 the detailed results for all measures are displayed.

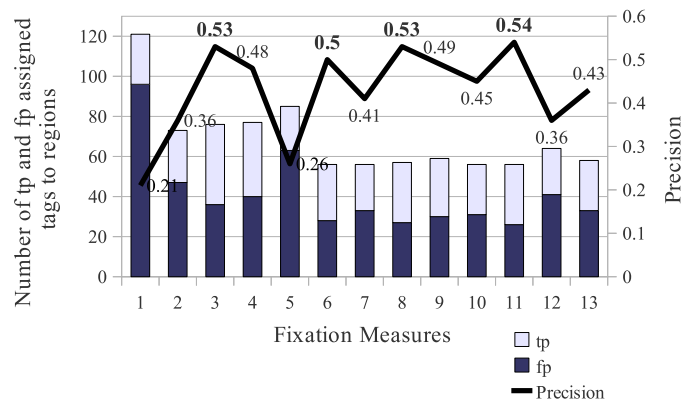


Figure 3: Precision Values for the Fixation Measures from Table 1

4.3 Extension of Region Boundaries

Research question 2 *Can the extension of region boundaries improve the assignment results?*

One reason why the identification of image regions is difficult is the inaccuracy of the eye-tracker. For the Tobii device the standard error is 1-0.5°. With a distance of 60 centimeters from the eye to the screen, this inaccuracy equates one centimeter on the screen or about 35 pixels. To compensate this factor, the region boundary can be extended. By this,

fixations near to a region are also considered belonging to the region. We have investigated the influence of this extension on the precision for the best performing measure `meanVisitDuration` (see Section 4.2). Values for the region extension $d = 1 \dots 35$ pixels are analyzed.

Results The precision increases when applying the extension parameter. The best result is precision $P = 0.61$ for $d = 13$ as shown in Figure 4. This equates to an improvement of about 13%.

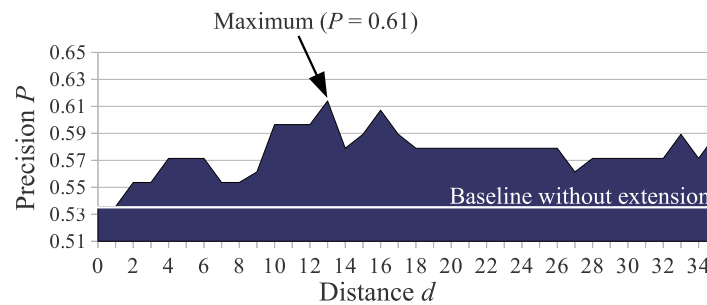


Figure 4: Influence of extension parameter

Interpretation It is reasonable to include the extension of region boundaries in the calculation of tag-to-region assignments. The precision is fluctuating depending on the extension value d chosen. Reasonable values for d are between 10 and 17. This parameter, however, should be looked into more detail and with different data sets to be able to better explain the fluctuation of the graph.

4.4 Weighting of Small Regions

Research question 3 *Can the weighting of small image regions improve the assignment results?*

Larger image regions have the benefit that it is more likely that fixations are located on them. Because of that, we analyze if the tag-to-region assignment can be improved when adding a weight to prefer smaller regions. We have compared the results of an exponential weighting function and a linear weighting function. The weighting of the exponential function was too strong, i.e., the preference for small regions was too high. Good results are obtained using a linear function. In the following, we consider the linear weighting function *weighted-fm* on image region r :

$$weighted-fm(r) = \begin{cases} fm(r) * weight(s_r) & \text{if } s_r \leq S \\ fm(r) & \text{else} \end{cases} \quad (2)$$

with

$$weight(s_r) = -\frac{t-1}{S} - s_r + t$$

In Equation 2, the function $fm(r)$ describes the calculation of a measure for a region r . s_r is the size of a region r in percentage of the whole image size. For example, $s_r = 30$ means that the region area is 30% of the size of the whole image. The measure is weighted with a factor only when $s_r \leq S$, where S is a predefined threshold. Thus, only image regions up to a specific size gain from the weighting function. The weighting factor itself is calculated depending on the threshold S and the maximum weighting value t . In our analysis, we have investigated the parameters S and t of the weighting function by calculating the precision values for all images for $S = 1 \dots 25\%$ and $t = 1 \dots 50$. An example of applying the *weighted-fm* for $S = 5$ and $t = 4$ is shown in Figure 5. Regions of size between 0 and 5% of the actual image are weighted with a factor between 1 and 4. We have computed the weights for the best performing measure `meanVisitDuration`.




Figure 5: Example weighting for $S = 5$ and $t = 4$

Results The best precision value applying the weighting function on the fixation measure `meanVisitDuration` is $P = 0.59$. The worst result is $P = 0.48$. These values are provided by different parameter combinations. In Figure 6, the results for the two weighting parameters are displayed. The precision value of $P = 0.54$ (see Section 4.2) is presented in white as baseline. Values higher than and lower than 0.54 are highlighted in the figure. The areas with the greatest distances to the baseline are labeled as “Maxima” and “Negative results”.

From the results in Figure 6, one can see that the influence of parameter S is higher than the influence of t . For values $t < 10$, the precision values

are varying very strong. For $4 < S < 10$ the precision is clearly higher than the baseline. For $S > 13$ the precision is lower than the baseline without the weighting function. Only for $S > 21$, again precision results near the baseline are reached.

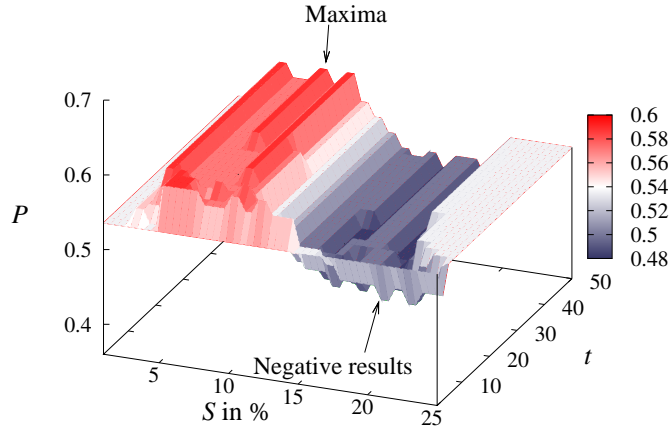


Figure 6: Influence of weighting function

Interpretation The usage of the weighting function can improve the results. However the precision can also decrease. For S we suggest values between 4 and 10. Like with the region extension, further investigations are necessary to better explain the fluctuation of the graph.

4.5 Combination of Region Extension and Weighting Function

Finally, we use the best performing fixation measure `meanVisitDuration` and combine the region extension and weighting function. When taking the image region extension and the weighting function into account, we receive the best results for precision $P = 0.67$ for d between 10 and 13, S between 3 and 13, and t between 2 and 25. Thus, the best precision can be obtained in some interval of the region extension parameter and weighting function. The worst result $P = 0.52$ was received for $d = 14$ and $d = 18$ (with S between 8 and 14, t between 2 and 25).

4.6 Compare with Baselines

Research question 4 *Is the gaze-based assignment of tags-to-regions better than a naive and random baseline?*

To analyze the quality of our tag-to-region assignment for the best measure `meanVisitDuration`, we compare it with the results from two baselines. In these and the following computations, we use the parameters $S = 5$, $t = 4$, and $d = 13$. We use the same baselines as they have been applied to evaluate relevance feedback from gaze information in [8] and [7]. We compare the number of *tp* and *fp* image-tag-pair assignments calculated from the baseline “naive” (a) and the baseline “random” (b) with the mere measure `meanVisitDuration` (c) and the `meanVisitDuration` measure including region extension and weighting (d). The naive baseline makes the assumption that the largest area in an image should be the asked one. It could be that the LabelMe users try to label larger image areas first. The random baseline randomly chooses one of the labels.

Results As the results in Figure 7 show, the naive approach has a precision of $P = 0.16$ and the random baseline of $P = 0.21$ compared to the gaze-based approach with a precision of $P = 0.54$ and the extended and weighted of $P = 0.67$. The precision values are calculated from the *tp* and *fp* tag-to-region pairs from equation 1.

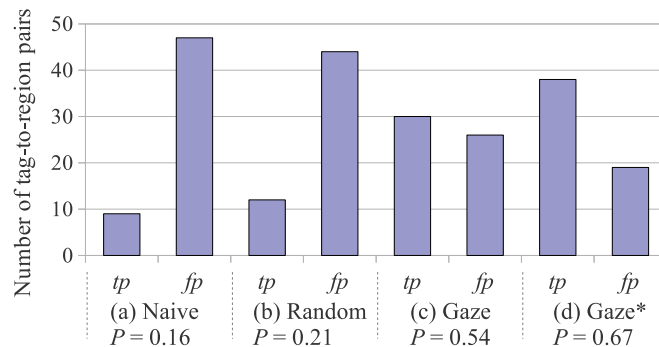


Figure 7: Compare *tp* and *fp* values for `meanVisitDuration` and `meanVisitDuration` with regions extension and weighting and two baselines

Interpretation The identification of assignments based on gaze information or on gaze information including extension and weighting performs better than both baseline approaches. Applying Chi-square upper-tailed tests show that the gaze assignments are significantly better than the baselines (all with $\alpha < 0.001$).

4.7 First Five Fixations

In our experiment setting, the subjects were asked to look at a red dot—placed above the image position—before the image appeared on the screen

(see Section 3.3). This was done to let all subjects start viewing the image from the same position. The measure `secondFixation` was added to investigate negative influence of only the first fixation. Removing all measures from our calculation, which are explicitly taking the first fixations into account (`firstFixation`, `secondFixation`, `firstFixationDuration`) improves the average precision over all measures. To illustrate this, Figure 8 shows the first five fixations over all subjects and all images. One can see, that the first fixations are centered in the middle of the images. Later, the fixations are better distributed over the images. This shows, that the first fixations are not applicable in this context (see also lowest precision in Section 4.2).




Figure 8: First five fixations accumulated over all subjects and all images

4.8 Tag-to-region assignments from gaze paths not corresponding to the region

Research question 5 *How often is a region identified as favorite region, without having shown the corresponding tag to the subjects?*

To answer this research question, the data of 20 subjects, each viewing one of two sets of 50 image-tag-pairs, is analyzed. Every image has two tags t_i and t_j assigned, one for each group, with $t_i \neq t_j$. The tags correspond to regions r_i and r_j in the image ($r_i \neq r_j$). As tags could be true or false, all combinations of true and false tags between the two groups of subjects appear. In our data set, there are 16 true-true image-tag-pairs (tags for both groups are true), 24 true-false image-tag-pairs (one tag is true, one false), and 10 false-false image-tag-pairs. In this section, we use the true-true and true-false pairs to answer the research question. For every true tag, shown to one of the two groups of 10 subjects, we analyze the assignment from the gaze paths of the other 10 subjects. It is calculated how often a region r_i (described by the tag t_i) is identified as favorite from the gaze paths corresponding to tag t_j . All results are calculated with measure `meanVisitDuration`, including extension and weighting.

Results The tp_i and fp_i values show the results for the assignment of tag t_i to region r_i by the gaze paths g_{t_i} . The tp_i assignments are the correct ones: the favorite region is described by the tag presented to the user. The gaze paths g_{t_j} were gained from the subjects, viewing tag t_j before viewing the

image. tp_j counts how often the region r_i is identified as favorite from gaze path g_{t_j} . 8 of 57 of the analyzed images-tag-pairs belong to tp_j (precision $P=0.12$). In the case fp_j , the region r_i is *not* identified as favorite. The fp_j assignments are the correct ones: the investigated region is not described by the tag presented to the user and it was not identified as favorite.




Figure 9: Assignment of tag t_i to region r_i (tp_i and fp_i) versus tag t_j to region r_i (tp_j and fp_j)

Interpretation The precision for assignments based on gaze paths corresponding to the given tag is $P = 0.73$; for tags corresponding to another region it is $P = 0.12$. The difference between these values is significant (Upper-tailed Chi-square test, $\alpha < 0.001$). These results show, that the quality of the assignments is based on the tag, shown to the subjects during the experiment.

4.9 Compare Tag-to-Region Assignments for Multiple Regions in One Image

Research question 6 *Can different regions be identified in one image?*

In this section, we investigate if it is possible to identify different regions from gaze paths corresponding to different tags in the same image. We use the measure `meanVisitDuration`, including extension and weighting, to calculate the results.

Results For 16 images with two correct tags, the favorite image regions were calculated. In 6 images, two correct image regions were identified. This is a proportion of 38%. In Figure 10, some examples with two correctly identified regions are shown. As the figure shows, the two tags *sky* and *sea* could be distinguished in the upper image. Also the tags *water pot* and *teas* in the lower image could be identified using gaze information.

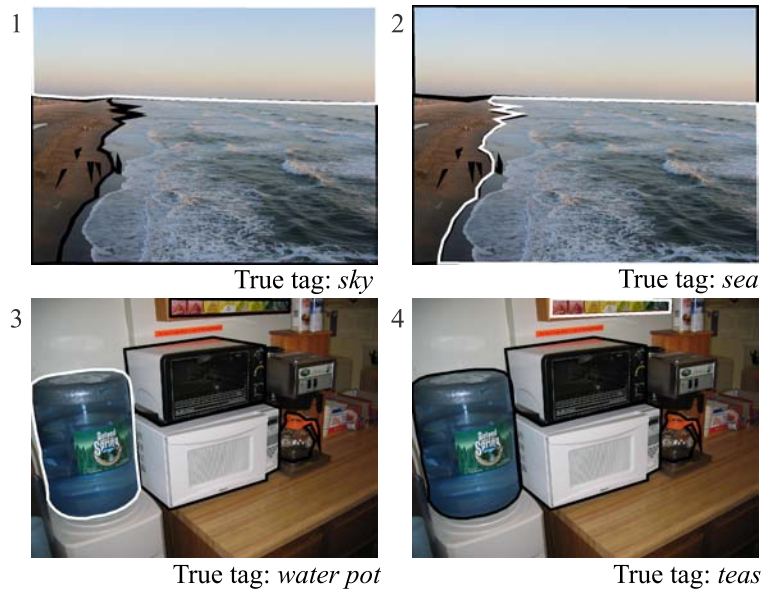


Figure 10: Example images with two correctly identified regions (white boundaries)

Interpretation The average probability to identify the correct region in one image is 67%. For two images, the probability of identifying correct assignments for both tags is 44%. With a value of 38% for two image regions in one image, the probability is close to the probability for two image regions in two different images. Thus, it is possible to identify different image regions in one image with an accuracy close to the accuracy of the single assignments.

4.10 Analyze Image Region Characteristics

Research question 7 *Are there typical region sizes and region positions for correctly versus incorrectly assigned tags-to-regions? What are typical characteristics for incorrect assignments?*

The best precision value $P = 0.67$ from measure `meanVisitDuration` (including extension and weighting) is calculated from 38 *tp* and 19 *fp* assignments. In this section, the assigned regions are examined concerning size and position. We also look into negative examples, i.e., regions that could not be successfully assigned.

Size The average size of the image regions, corresponding to the tags given to the subjects, is 23,415 pixels square. The average image size for correct assigned regions (*tp*) is 96,472, for incorrect assigned regions (*fp*) 106,541 pixels square.

The region size of favorite regions (tp or fp) is about five times larger than the average region size. Thus, tag-to-region assignments are preferably conducted for larger regions. It is also interesting to notice that the average region size of tp and fp assignments is very close. Thus, it seems that the region size has no influence on the correctness or incorrectness of the assignments.

Position In Figure 11, the positions of the regions in the images are depicted. The images are divided into nine uniform areas. The numbers shown in the areas describe the percentage of the regions or parts of the regions located in the image areas. In the first diagram of Figure 11(a), the positions of all regions corresponding to true tags are depicted. 49% of the regions touch the center field of the image. In the upper third of the image areas, there is only one fourth of the regions located. In the lower areas, there is about one third.

In Figure 11(b) and (c), the correctly and incorrectly assigned regions are summarized. One can see in Figure 11(b), that the percentage of correctly assigned regions is almost evenly spread over the nine areas. This indicates that correct assignments of regions is independent of the image area. Many incorrectly assigned regions, however, are in the center of the images as shown in Figure 11(c). This higher percentage of wrongly assigned regions might be caused by a concentration of fixations in the center of the images. This concentration has been observed during the first fixations on the images as shown in Figure 8.

25%	25%	25%
46%	49%	43%
33%	33%	32%

(a) Region distribution
of all tags

31%	29%	32%	26%	32%	11%
33%	39%	37%	58%	74%	42%
32%	39%	37%	53%	47%	37%

(b) Correctly
assigned regions

(c) Incorrectly
assigned regions

Figure 11: Image areas in which the tag-to-region-assignments are located

Qualitative Analysis of Incorrect Assignments Some examples of incorrect assignments can be seen in Figure 12. The white boundaries show the object that corresponds to the tag given to user. The black boundaries show the objects determined as favorite from the gaze information. From an

qualitative analysis of 19 wrongly assigned tags to regions, we have identified the following characteristics of the images with incorrect assignments:

- Some images show scenes with a small correct object and a wrongly selected favorite object also small and located next to the correct object (cf. images 1 and 2). This problem can be based on the accuracy of the eye-tracker. Five images belonged to this category.
- In some images, the correct object is displayed within another object (cf. image 3, *lamp* inside *wall*). In these cases, the outer region is identified as favorite. That means our weighting function does not work for all occurrences of smaller regions. Also five images belong to this category.
- Further images show scenes with an object that seems to be easy to identify. For example larger objects like *road* (cf. image 4), *sky* or *tree* might be perceived even in the corner of the human eye or based on context knowledge (e.g., sky is above sea is above sand in a beach scene). Seven images belong to this category.

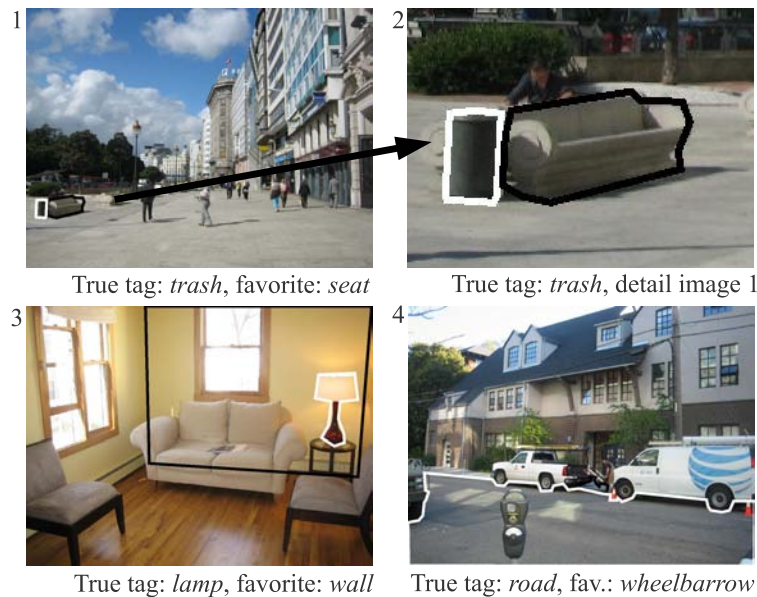


Figure 12: Examples of image-tag-pairs with given tags (white boundaries) and *fp* assignments (black boundaries)

4.11 Aggregation of Gaze Paths

Research question 8 *How is the precision influenced by the aggregation of gaze paths?*

For the calculations in this paper, the gaze paths of up to 10 users are aggregated (only gaze paths of users who gave a wrong answer were ignored). In this section, we investigate how strong the influence of the aggregation over multiple subjects on the precision is. We present precision values for aggregations of 1 to 10 subjects for the measure `meanVisitDuration`, including extension and weighting. Precision P is calculated for every possible subset of subjects and averaged for all subgroups of the same size.

Results As Figure 13 shows, the influence of the number of users is very high. With the gaze paths of only a single user, we have received an average precision (over all users and all images) of $P = 0.31$. For the aggregated data for all 10 users we got a precision $P = 0.67$, this corresponds to an improvement of 109%. The biggest improvements take place between the first group sizes. For example between one and two users per group we have an improvement of 46%. In comparison, between nine users and ten users per group, there is only an improvement of 4%.

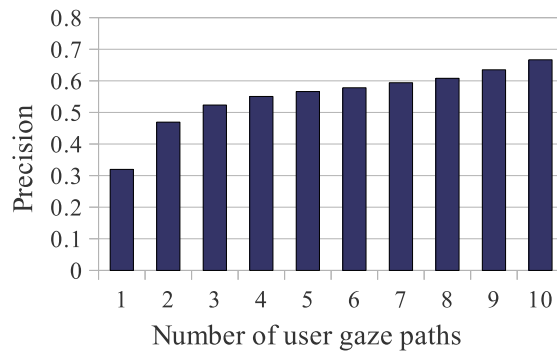


Figure 13: Effect of Aggregating Gaze Paths

Interpretation The results based on multiple gaze paths are considerably better than the ones calculated from only a few gaze paths. However the improvement of the precision gets lower when aggregating more gaze paths. A clear trade of can yet not be made and should be investigated. Compared with the two baselines from Section 4.6, the results for single users are still significant better than the naive or random baseline (Chi-square upper-tailed test with $\alpha < 0.001$ for the naive approach and $\alpha < 0.002$ for the random approach).

5 Summary of Evaluation Results

In Section 4.2, we have shown that the best measure is `meanVisitDuration` with precision $P = 0.54$. In general, good results are obtained for stan-

standard measures delivering aggregations of the whole gaze path. The first fixations of a gaze path are not applicable in our application. Also the new invented measures, which take the moment of decision into consideration, provide good results. The new measure `fixationsBeforeDecision` and `lastFixationDuration` provides the second best results and should be considered in future experiments. The best performance is gained by a combination of extension of region boundaries (Section 4.3) and weighting of small image regions (Section 4.4) with an improvement of the precision from $P = 0.54$ to $P = 0.67$ (Section 4.5). We have shown that the first fixations in a gaze path are not suitable in our approach in Section 4.7.

We have compared the results with two baselines (Section 4.6). The detailed analysis shows that the identified assignments are much better than the naive baseline or random baseline. The correctly identified correlations are not found by chance: in Section 4.8 we have shown that the correlations strongly depends on the tag, shown to the users. In addition, we are able to assign different regions to different tags in one image with a precision of 38% (Section 4.9). In Section 4.10 we have analyzed typical characteristics for the assigned regions. The result is that size and position of an area do not have an influence on the correctness of the assignment in principle.

The aggregation of user gaze paths in Section 4.11 shows the potential to improve the number of correct assignments.

6 Conclusions

In this paper, we have shown that it is possible to identify image regions by analyzing gaze paths of users viewing the image with a given tag at a precision of 67%. In addition, we have shown that two different regions can be differentiated in the same image with an accuracy of 38%. Possible application of our results are image tagging or image search.

The results are gained in a controlled experiment with manually segmented images from the LabelMe data set. We have used LabelMe instead of applying automatic segmentation based on low-level features because of the additional error that would have been introduced in the experiment by automatic segmentation.

The next step will be to apply the experiment on automatically segmented images. Such automatic segmentation can be improved by using the gaze information as it has been done by Santella et al [15]. In the future work, we also plan to apply Support Vector Machines for training combinations of fixation measures, extension, and weighting.

7 Acknowledgments

The research leading to this paper was supported by the European Commission under contract FP7-216444, Petamedia Peer-to-peer Tagged Media. The authors thank the subjects for participating in our experiments and giving us feedback. We thank our student Chantal Neuhaus for supporting us in implementing the experiment.

References

- [1] R. J. Campbell and P. J. Flynn. A survey of free-form object representation and recognition techniques. *CVIU*, 81(2):166–210, 2001.
- [2] S. Hajimirza and E. Izquierdo. Gaze movement inference for implicit image annotation. In *Workshop on Image Analysis for Multimedia Interactive Services*. IEEE, 2010.
- [3] D. Hardoon, K. Pasupa, and S. Szedmak. Image ranking with implicit feedback from eye movements. In *Symposium on Eye Tracking Research & Applications*, 2010.
- [4] A. Jaimes. Using human observer eye movements in automatic image classifiers. *SPIE*, 2001.
- [5] A. Klami. Inferring task-relevant image regions from gaze data. In *Workshop on Machine Learning for Signal Processing*. IEEE, 2010.
- [6] A. Klami, S. Kaski, K. Pasupa, S. Szedmak, S. Gunn, and D. Predicting relevance of parts of an image. 2009.
- [7] A. Klami, C. Saunders, T. De Campos, and S. Kaski. Can relevance of images be inferred from eye movements? In *Multimedia information retrieval*. ACM, 2008.
- [8] L. Kozma, A. Klami, and S. Kaski. GaZIR: gaze-based zooming interface for image retrieval. In *Multimodal interfaces*. ACM, 2009.
- [9] X. Li, C. G. M. Snoek, and M. Worring. Annotating images by harnessing worldwide user-tagged photos. In *Acoustics, Speech, and Signal Processing*, pages 3717–3720. IEEE, 2009.
- [10] K. Pasupa, C. Saunders, S. Szedmak, A. Klami, S. Kaski, and S. Gunn. Learning to rank images from eye movements. In *Workshops on Human-Computer Interaction*, 2009.
- [11] C. M. Privitera and L. W. Stark. Algorithms for defining visual regions-of-interest: Comparison with eye fixations. *IEEE Transactions on pattern analysis and machine intelligence*, 22(9):970–982, 2000.

- [12] S. Ramanathan, H. Katti, R. Huang, T.-S. Chua, and M. Kankanhalli. Automated localization of affective objects and actions in images via caption text-cum-eye gaze analysis. *Multimedia*, 2009.
- [13] N. Rowe. Finding and labeling the subject of a captioned depictive natural photograph. *IEEE Transactions on Knowledge and Data Engineering*, pages 202–207, 2002.
- [14] B. C. Russell, A. Torralba, K. P. Murphy, and W. T. Freeman. LabelMe: a database and web-based tool for image annotation. *Journal of Computer Vision*, 77(1):157–173, 2008.
- [15] A. Santella, M. Agrawala, D. DeCarlo, D. Salesin, and M. Cohen. Gaze-based interaction for semi-automatic photo cropping. In *CHI*, page 780. ACM, 2006.
- [16] L. von Ahn, R. Liu, and M. Blum. Peekaboom: a game for locating objects in images. In *CHI*. ACM, 2006.

No	Name	Description	Favorite & Unit	Origin
1	firstFixation	Number of times the subject fixates on the image before fixating on region r for the first time	min count	Tobii
2	fixationsAfter	Number of times the subject fixates on the image after last fixation on region r	min count	[5]
3	fixationsBeforeDecision	Number of times the subject fixates on the image after the last fixation on r and before the decision	min count	New
4	fixationsAfterDecision	Number of times the subject fixates on the image the decision and before the fixation on region r	min count	New
5	secondFixation	Number of times the subject fixates on the image before fixating on region r for the first time without the first fixation on the image	min count	New
6	fixationDuration	Sum of the duration of all fixations on r	max seconds	Tobii
7	firstFixationDuration	Duration of the first fixation on r	max seconds	Tobii
8	lastFixationDuration	Duration of the last fixation on r	max seconds	New
9	fixationCount	Number of times the subject fixates on r	max count	Tobii
10	maxVisitDuration	Maximum visit length on r	max seconds	Tobii
11	meanVisitDuration	Mean visit length on r	max seconds	Tobii
12	visitCount	Number of visits within r	max count	Tobii
13	saccLength	Length of saccade before fixation on r	max centimeter	[8]

Table 1: Eye-tracking measures for a region r

Bisher erschienen

Arbeitsberichte aus dem Fachbereich Informatik

(<http://www.uni-koblenz-landau.de/koblenz/fb4/publications/Reports/arbeitsberichte>)

Tina Walber, Ansgar Scherp, Steffen Staab, Towards the Understanding of Image Semantics by Gaze-based Tag-to-Region Assignments, Arbeitsberichte aus dem Fachbereich Informatik 8/2011

Alexander Kleinen, Ansgar Scherp, Steffen Staab, Mobile Facets – Faceted Search and Exploration of Open Social Media Data on a Touchscreen Mobile Phone, Arbeitsberichte aus dem Fachbereich Informatik 7/2011

Anna Lantsberg, Klaus G. Troitzsch, Towards A Methodology of Developing Models of E-Service Quality Assessment in Healthcare, Arbeitsberichte aus dem Fachbereich Informatik 6/2011

Ansgar Scherp, Carsten Saathoff, Thomas Franz, Steffen Staab, Designing Core Ontologies, Arbeitsberichte aus dem Fachbereich Informatik 5/2011

Oleg V. Kryuchin, Alexander A. Arzamastsev, Klaus G. Troitzsch, The prediction of currency exchange rates using artificial neural networks, Arbeitsberichte aus dem Fachbereich Informatik 4/2011

Klaus G. Troitzsch, Anna Lantsberg, Requirements for Health Care Related Websites in Russia: Results from an Analysis of American, British and German Examples, Arbeitsberichte aus dem Fachbereich Informatik 3/2011

Klaus G. Troitzsch, Oleg Kryuchin, Alexander Arzamastsev, A universal simulator based on artificial neural networks for computer clusters, Arbeitsberichte aus dem Fachbereich Informatik 2/2011

Klaus G. Troitzsch, Natalia Zenkova, Alexander Arzamastsev, Development of a technology of designing intelligent information systems for the estimation of social objects, Arbeitsberichte aus dem Fachbereich Informatik 1/2011

Kurt Lautenbach, A Petri Net Approach for Propagating Probabilities and Mass Functions, Arbeitsberichte aus dem Fachbereich Informatik 13/2010

Claudia Schon, Linkless Normal Form for ALC Concepts, Arbeitsberichte aus dem Fachbereich Informatik 12/2010

Alexander Hug, Informatik hautnah erleben, Arbeitsberichte aus dem Fachbereich Informatik 11/2010

Marc Santos, Harald F.O. von Kortzfleisch, Shared Annotation Model – Ein Datenmodell für kollaborative Annotationen, Arbeitsberichte aus dem Fachbereich Informatik 10/2010

Gerd Gröner, Steffen Staab, Categorization and Recognition of Ontology Refactoring Pattern, Arbeitsberichte aus dem Fachbereich Informatik 9/2010

Daniel Eißing, Ansgar Scherp, Carsten Saathoff, Integration of Existing Multimedia Metadata Formats and Metadata Standards in the M3O, Arbeitsberichte aus dem Fachbereich Informatik 8/2010

Stefan Scheglmann, Ansgar Scherp, Steffen Staab, Model-driven Generation of APIs for OWL-based Ontologies, Arbeitsberichte aus dem Fachbereich Informatik 7/2010

Daniel Schmeiß, Ansgar Scherp, Steffen Staab, Integrated Mobile Visualization and Interaction of Events and POIs, Arbeitsberichte aus dem Fachbereich Informatik 6/2010

Rüdiger Grimm, Daniel Pähler, E-Mail-Forensik – IP-Adressen und ihre Zuordnung zu Internet-Teilnehmern und ihren Standorten, Arbeitsberichte aus dem Fachbereich Informatik 5/2010

Christoph Ringelstein, Steffen Staab, PAPEL: Syntax and Semantics for Provenance-Aware Policy Definition, Arbeitsberichte aus dem Fachbereich Informatik 4/2010

Nadine Lindermann, Sylvia Valcárcel, Harald F.O. von Kortzfleisch, Ein Stufenmodell für kollaborative offene Innovationsprozesse in Netzwerken kleiner und mittlerer Unternehmen mit Web 2.0, Arbeitsberichte aus dem Fachbereich Informatik 3/2010

Maria Wimmer, Dagmar Lück-Schneider, Uwe Brinkhoff, Erich Schweighofer, Siegfried Kaiser, Andreas Wieber, Fachtagung Verwaltungsinformatik FTVI Fachtagung Rechtsinformatik FTRI 2010, Arbeitsberichte aus dem Fachbereich Informatik 2/2010

Max Braun, Ansgar Scherp, Steffen Staab, Collaborative Creation of Semantic Points of Interest as Linked Data on the Mobile Phone, Arbeitsberichte aus dem Fachbereich Informatik 1/2010

Marc Santos, Einsatz von „Shared In-situ Problem Solving“ Annotationen in kollaborativen Lern- und Arbeitsszenarien, Arbeitsberichte aus dem Fachbereich Informatik 20/2009

Carsten Saathoff, Ansgar Scherp, Unlocking the Semantics of Multimedia Presentations in the Web with the Multimedia Metadata Ontology, Arbeitsberichte aus dem Fachbereich Informatik 19/2009

Christoph Kahle, Mario Schaarschmidt, Harald F.O. von Kortzfleisch, Open Innovation: Kundenintegration am Beispiel von IPTV, Arbeitsberichte aus dem Fachbereich Informatik 18/2009

Dietrich Paulus, Lutz Priese, Peter Decker, Frank Schmitt, Pose-Tracking Forschungsbericht, Arbeitsberichte aus dem Fachbereich Informatik 17/2009

Andreas Fuhr, Tassilo Horn, Andreas Winter, Model-Driven Software Migration Extending SOMA, Arbeitsberichte aus dem Fachbereich Informatik 16/2009

Eckhard Großmann, Sascha Strauß, Tassilo Horn, Volker Riediger, Abbildung von grUML nach XSD soamig, Arbeitsberichte aus dem Fachbereich Informatik 15/2009

Kerstin Falkowski, Jürgen Ebert, The STOR Component System Interim Report, Arbeitsberichte aus dem Fachbereich Informatik 14/2009

Sebastian Magnus, Markus Maron, An Empirical Study to Evaluate the Location of Advertisement Panels by Using a Mobile Marketing Tool, Arbeitsberichte aus dem Fachbereich Informatik 13/2009

Sebastian Magnus, Markus Maron, Konzept einer Public Key Infrastruktur in iCity, Arbeitsberichte aus dem Fachbereich Informatik 12/2009

Sebastian Magnus, Markus Maron, A Public Key Infrastructure in Ambient Information and Transaction Systems, Arbeitsberichte aus dem Fachbereich Informatik 11/2009

Ammar Mohammed, Ulrich Furbach, Multi-agent systems: Modeling and Virification using Hybrid Automata, Arbeitsberichte aus dem Fachbereich Informatik 10/2009

Andreas Sprotte, Performance Measurement auf der Basis von Kennzahlen aus betrieblichen Anwendungssystemen: Entwurf eines kennzahlengestützten Informationssystems für einen Logistikdienstleister, Arbeitsberichte aus dem Fachbereich Informatik 9/2009

Gwendolin Garbe, Tobias Hausen, Process Commodities: Entwicklung eines Reifegradmodells als Basis für Outsourcingentscheidungen, Arbeitsberichte aus dem Fachbereich Informatik 8/2009

Petra Schubert et. al., Open-Source-Software für das Enterprise Resource Planning, Arbeitsberichte aus dem Fachbereich Informatik 7/2009

Ammar Mohammed, Frieder Stolzenburg, Using Constraint Logic Programming for Modeling and Verifying Hierarchical Hybrid Automata, Arbeitsberichte aus dem Fachbereich Informatik 6/2009

Tobias Kippert, Anastasia Meletiadou, Rüdiger Grimm, Entwurf eines Common Criteria-Schutzprofils für Router zur Abwehr von Online-Überwachung, Arbeitsberichte aus dem Fachbereich Informatik 5/2009

Hannes Schwarz, Jürgen Ebert, Andreas Winter, Graph-based Traceability – A Comprehensive Approach. Arbeitsberichte aus dem Fachbereich Informatik 4/2009

Anastasia Meletiadou, Simone Müller, Rüdiger Grimm, Anforderungsanalyse für Risk-Management-Informationssysteme (RMIS), Arbeitsberichte aus dem Fachbereich Informatik 3/2009

Ansgar Scherp, Thomas Franz, Carsten Saathoff, Steffen Staab, A Model of Events based on a Foundational Ontology, Arbeitsberichte aus dem Fachbereich Informatik 2/2009

Frank Bohdanovicz, Harald Dickel, Christoph Steigner, Avoidance of Routing Loops, Arbeitsberichte aus dem Fachbereich Informatik 1/2009

Stefan Ameling, Stephan Wirth, Dietrich Paulus, Methods for Polyp Detection in Colonoscopy Videos: A Review, Arbeitsberichte aus dem Fachbereich Informatik 14/2008

Tassilo Horn, Jürgen Ebert, Ein Referenzschema für die Sprachen der IEC 61131-3, Arbeitsberichte aus dem Fachbereich Informatik 13/2008

Thomas Franz, Ansgar Scherp, Steffen Staab, Does a Semantic Web Facilitate Your Daily Tasks?, Arbeitsberichte aus dem Fachbereich Informatik 12/2008

Norbert Frick, Künftige Anforderungen an ERP-Systeme: Deutsche Anbieter im Fokus, Arbeitsberichte aus dem Fachbereich Informatik 11/2008

Jürgen Ebert, Rüdiger Grimm, Alexander Hug, Lehramtsbezogene Bachelor- und Masterstudiengänge im Fach Informatik an der Universität Koblenz-Landau, Campus Koblenz, Arbeitsberichte aus dem Fachbereich Informatik 10/2008

Mario Schaarschmidt, Harald von Kortzfleisch, Social Networking Platforms as Creativity Fostering Systems: Research Model and Exploratory Study, Arbeitsberichte aus dem Fachbereich Informatik 9/2008

Bernhard Schueler, Sergej Sizov, Steffen Staab, Querying for Meta Knowledge, Arbeitsberichte aus dem Fachbereich Informatik 8/2008

Stefan Stein, Entwicklung einer Architektur für komplexe kontextbezogene Dienste im mobilen Umfeld, Arbeitsberichte aus dem Fachbereich Informatik 7/2008

Matthias Bohnen, Lina Brühl, Sebastian Bzdak, RoboCup 2008 Mixed Reality League Team Description, Arbeitsberichte aus dem Fachbereich Informatik 6/2008

Bernhard Beckert, Reiner Hähnle, Tests and Proofs: Papers Presented at the Second International Conference, TAP 2008, Prato, Italy, April 2008, Arbeitsberichte aus dem Fachbereich Informatik 5/2008

Klaas Dellschaft, Steffen Staab, Unterstützung und Dokumentation kollaborativer Entwurfs- und Entscheidungsprozesse, Arbeitsberichte aus dem Fachbereich Informatik 4/2008

Rüdiger Grimm: IT-Sicherheitsmodelle, Arbeitsberichte aus dem Fachbereich Informatik 3/2008

Rüdiger Grimm, Helge Hundacker, Anastasia Meletiadou: Anwendungsbeispiele für Kryptographie, Arbeitsberichte aus dem Fachbereich Informatik 2/2008

Markus Maron, Kevin Read, Michael Schulze: CAMPUS NEWS – Artificial Intelligence Methods Combined for an Intelligent Information Network, Arbeitsberichte aus dem Fachbereich Informatik 1/2008

Lutz Priebe, Frank Schmitt, Patrick Sturm, Haojun Wang: BMBF-Verbundprojekt 3D-RETISEG Abschlussbericht des Labors Bilderkennen der Universität Koblenz-Landau, Arbeitsberichte aus dem Fachbereich Informatik 26/2007

Stephan Philippi, Alexander Pinl: Proceedings 14. Workshop 20.-21. September 2007 Algorithmen und Werkzeuge für Petrinetze, Arbeitsberichte aus dem Fachbereich Informatik 25/2007

Ulrich Furbach, Markus Maron, Kevin Read: CAMPUS NEWS – an Intelligent Bluetooth-based Mobile Information Network, Arbeitsberichte aus dem Fachbereich Informatik 24/2007

Ulrich Furbach, Markus Maron, Kevin Read: CAMPUS NEWS - an Information Network for Pervasive Universities, Arbeitsberichte aus dem Fachbereich Informatik 23/2007

Lutz Priebe: Finite Automata on Unranked and Unordered DAGs Extended Version, Arbeitsberichte aus dem Fachbereich Informatik 22/2007

Mario Schaarschmidt, Harald F.O. von Kortzfleisch: Modularität als alternative Technologie- und Innovationsstrategie, Arbeitsberichte aus dem Fachbereich Informatik 21/2007

Kurt Lautenbach, Alexander Pinl: Probability Propagation Nets, Arbeitsberichte aus dem Fachbereich Informatik 20/2007

Rüdiger Grimm, Farid Mehr, Anastasia Meletiadou, Daniel Pähler, Ilka Uerz: SOA-Security, Arbeitsberichte aus dem Fachbereich Informatik 19/2007

Christoph Wernhard: Tableaux Between Proving, Projection and Compilation, Arbeitsberichte aus dem Fachbereich Informatik 18/2007

Ulrich Furbach, Claudia Obermaier: Knowledge Compilation for Description Logics, Arbeitsberichte aus dem Fachbereich Informatik 17/2007

Fernando Silva Parreiras, Steffen Staab, Andreas Winter: TwoUse: Integrating UML Models and OWL Ontologies, Arbeitsberichte aus dem Fachbereich Informatik 16/2007

Rüdiger Grimm, Anastasia Meletiadou: Rollenbasierte Zugriffskontrolle (RBAC) im Gesundheitswesen, Arbeitsberichte aus dem Fachbereich Informatik 15/2007

Ulrich Furbach, Jan Murray, Falk Schmidsberger, Frieder Stolzenburg: Hybrid Multiagent Systems with Timed Synchronization-Specification and Model Checking, Arbeitsberichte aus dem Fachbereich Informatik 14/2007

Björn Pelzer, Christoph Wernhard: System Description: "E-KRHyper", Arbeitsberichte aus dem Fachbereich Informatik, 13/2007

Ulrich Furbach, Peter Baumgartner, Björn Pelzer: Hyper Tableaux with Equality, Arbeitsberichte aus dem Fachbereich Informatik, 12/2007

Ulrich Furbach, Markus Maron, Kevin Read: Location based Informationssysteme, Arbeitsberichte aus dem Fachbereich Informatik, 11/2007

Philipp Schaer, Marco Thum: State-of-the-Art: Interaktion in erweiterten Realitäten, Arbeitsberichte aus dem Fachbereich Informatik, 10/2007

Ulrich Furbach, Claudia Obermaier: Applications of Automated Reasoning, Arbeitsberichte aus dem Fachbereich Informatik, 9/2007

Jürgen Ebert, Kerstin Falkowski: A First Proposal for an Overall Structure of an Enhanced Reality Framework, Arbeitsberichte aus dem Fachbereich Informatik, 8/2007

Lutz Prieße, Frank Schmitt, Paul Lemke: Automatische See-Through Kalibrierung, Arbeitsberichte aus dem Fachbereich Informatik, 7/2007

Rüdiger Grimm, Robert Krimmer, Nils Meißner, Kai Reinhard, Melanie Volkamer, Marcel Weinand, Jörg Helbach: Security Requirements for Non-political Internet Voting, Arbeitsberichte aus dem Fachbereich Informatik, 6/2007

Daniel Bildhauer, Volker Riediger, Hannes Schwarz, Sascha Strauß, „grUML – Eine UML-basierte Modellierungssprache für T-Graphen“, Arbeitsberichte aus dem Fachbereich Informatik, 5/2007

Richard Arndt, Steffen Staab, Raphaël Troncy, Lynda Hardman: Adding Formal Semantics to MPEG-7: Designing a Well Founded Multimedia Ontology for the Web, Arbeitsberichte aus dem Fachbereich Informatik, 4/2007

Simon Schenk, Steffen Staab: Networked RDF Graphs, Arbeitsberichte aus dem Fachbereich Informatik, 3/2007

Rüdiger Grimm, Helge Hundacker, Anastasia Meletiadiou: Anwendungsbeispiele für Kryptographie, Arbeitsberichte aus dem Fachbereich Informatik, 2/2007

Anastasia Meletiadiou, J. Felix Hampe: Begriffsbestimmung und erwartete Trends im IT-Risk-Management, Arbeitsberichte aus dem Fachbereich Informatik, 1/2007

„Gelbe Reihe“

(<http://www.uni-koblenz.de/fb4/publikationen/gelbereihe>)

Lutz Prieße: Some Examples of Semi-rational and Non-semi-rational DAG Languages. Extended Version, Fachberichte Informatik 3-2006

Kurt Lautenbach, Stephan Philippi, and Alexander Pinl: Bayesian Networks and Petri Nets, Fachberichte Informatik 2-2006

Rainer Gimnich and Andreas Winter: Workshop Software-Reengineering und Services, Fachberichte Informatik 1-2006

Kurt Lautenbach and Alexander Pinl: Probability Propagation in Petri Nets, Fachberichte Informatik 16-2005

Rainer Gimnich, Uwe Kaiser, and Andreas Winter: 2. Workshop "Reengineering Prozesse" – Software Migration, Fachberichte Informatik 15-2005

Jan Murray, Frieder Stolzenburg, and Toshiaki Arai: Hybrid State Machines with Timed Synchronization for Multi-Robot System Specification, Fachberichte Informatik 14-2005

Reinhold Letz: FTP 2005 – Fifth International Workshop on First-Order Theorem Proving, Fachberichte Informatik 13-2005

Bernhard Beckert: TABLEAUX 2005 – Position Papers and Tutorial Descriptions, Fachberichte Informatik 12-2005

Dietrich Paulus and Detlev Droege: Mixed-reality as a challenge to image understanding and artificial intelligence, Fachberichte Informatik 11-2005

Jürgen Sauer: 19. Workshop Planen, Scheduling und Konfigurieren / Entwerfen, Fachberichte Informatik 10-2005

Pascal Hitzler, Carsten Lutz, and Gerd Stumme: Foundational Aspects of Ontologies, Fachberichte Informatik 9-2005

Joachim Baumeister and Dietmar Seipel: Knowledge Engineering and Software Engineering, Fachberichte Informatik 8-2005

Benno Stein and Sven Meier zu Eißel: Proceedings of the Second International Workshop on Text-Based Information Retrieval, Fachberichte Informatik 7-2005

Andreas Winter and Jürgen Ebert: Metamodel-driven Service Interoperability, Fachberichte Informatik 6-2005

Joschka Boedecker, Norbert Michael Mayer, Masaki Ogino, Rodrigo da Silva Guerra, Masaaki Kikuchi, and Minoru Asada: Getting closer: How Simulation and Humanoid League can benefit from each other, Fachberichte Informatik 5-2005

Torsten Gipp and Jürgen Ebert: Web Engineering does profit from a Functional Approach, Fachberichte Informatik 4-2005

Oliver Obst, Anita Maas, and Joschka Boedecker: HTN Planning for Flexible Coordination Of Multiagent Team Behavior, Fachberichte Informatik 3-2005

Andreas von Hessling, Thomas Kleemann, and Alex Sinner: Semantic User Profiles and their Applications in a Mobile Environment, Fachberichte Informatik 2-2005

Heni Ben Amor and Achim Rettinger: Intelligent Exploration for Genetic Algorithms – Using Self-Organizing Maps in Evolutionary Computation, Fachberichte Informatik 1-2005