Understanding Relevance: An fMRI Study

Yashar Moshfeghi[1], Luisa R. Pinto[1], Frank E. Pollick[2], and Joemon M. Jose[1]

[1] School of Computing Science, University of Glasgow, Glasgow, UK
[2] School of Psychology, University of Glasgow, Glasgow, UK
{Yashar.Moshfeghi,Frank.Pollick,Joemon.Jose}@glasgow.ac.uk,
l.pinto.1@research.gla.ac.uk

Abstract. Relevance is one of the key concepts in Information Retrieval (IR). A huge body of research exists that attempts to understand this concept so as to operationalize it for IR systems. Despite advances in the past few decades, answering the question "How does relevance happen?" is still a big challenge. In this paper, we investigate the connection between relevance and brain activity. Using functional Magnetic Resonance Imaging (fMRI), we measured the brain activity of eighteen participants while they performed four topical relevance assessment tasks on relevant and non-relevant images. The results of this experiment revealed three brain regions in the frontal, parietal and temporal cortex where brain activity differed between processing relevant and non-relevant documents. This is an important step in unravelling the nature of relevance and therefore better utilising it for effective retrieval.

1 Introduction

Relevance is a fundamental concept in information retrieval (IR), and a huge body of research exists that attempts to understand this concept so as to operationalize it for IR systems. Despite all the research done on relevance over the past 40 years, answering the question "How does relevance happen?" is still a big challenge for information science and retrieval (IS&R) communities [1]. This is due to two main problems associated with relevance. Firstly, notions relating to human judgement, such as relevance, are hard to grasp and difficult to define for the purpose of automatic interpretation by a system. Secondly, although there is some agreement about the concept of *relevance* per se, there is disagreement on *what should be considered as relevant* since it depends on individual preferences [1]. Saracevic considers relevance as a phenomenon that should not be studied as a formally defined entity, that is, the answer to a *"what is it?"* question. Instead, it should be regarded as a *"what is the nature of it?"* question, similar to those in natural sciences [1]. Understanding the brain regions involved in the relevance assessment process is the first step in unraveling the nature of relevance. In this paper, we try to locate the brain regions involved in the relevance assessment process.

The orientation of research into the concept of relevance has developed from considering only the relevance of documents to a query, to understanding and

P. Serdyukov et al. (Eds.): ECIR 2013, LNCS 7814, pp. 14–25, 2013.

modelling a more user oriented concept of relevance [2]. IR systems use an operationalized version of relevance. This concentrates on the retrieval algorithms that process information objects and match them with users' queries, attempting to maximise retrieval of relevant information objects (system side of relevance). Users of such systems then browse through retrieved results to find what they consider relevant (user side of relevance) depending on their context, cognition, and/or affect. The system and user sides of relevance are complementary, and in order to improve the performance of IR systems, the user and system sides should work together [1].

Therefore, IR systems employ feedback techniques to integrate the user and system sides of relevance. An example of such an approach is the relevance feedback technique where feedback is gathered through explicit [3], implicit [4], and/or affective feedback [5]. Despite the robustness of explicit feedback in improving retrieval effectiveness [3], it is not always applicable or reliable due to the cognitive burden that it places on users [6]. To overcome this cognitive burden, implicit feedback is proposed where relevance is inferred from the interactional data in an indirect and unobtrusive manner [7]. For example, researchers try to understand how task [4], dwell time [8] and click-through [7] relate to relevance. However, a problem occurs when actions are taken as an indication of relevance without sufficient evidence to support their effectiveness [9]. For example, Kelly and Belkin [10] show that the implicit feedback measures based on user interaction with the full content of documents can often be unreliable, and difficult to measure or interpret. However, if we can identify which brain regions are activated during the explicit relevance judgement, and how these activations changes for positive (relevant) and negative (non-relevant) feedback, we can use these finding as a direct way of measuring relevance, and hence can use as an effective feedback technique without intruding the user cognitive process. This paper focuses on the first step, which is to identify the brain regions activated while performing explicit relevance judgment tasks.

Recently, affective feedback has been proposed [5] where the idea is to capture facial expression [11], eye tracking [12], and physiological signals [13] (such as skin temperature) and use them as implicit relevance judgement. However, these methods can only help researchers to understand the concept of relevance to a certain level and are not considered to be very effective. It is intriguing to study what the underlying brain activity is which leads to the conclusion that a document is relevant or not. This is only possible by investigating what exactly happens inside the human brain while users are analysing the relevance of the retrieved documents. In this work, we try to examine human brain activity while they are assessing the relevance of information objects with respect to a predefined task. We devised a user study to investigate the underlying brain activity during the relevance assessment of information objects. In our experiment, we focus on topical relevance[1], since it is the basis for system relevance [1]. The results from this paper are an important step towards better understanding

[1] Topical relevance is a relation between representations of information object and the topic or subject under consideration [14].

the notion of relevance. The remainder of the paper is organised as follows: Section 2 and 3 outline our experimental methodology and the obtained results respectively, followed by the discussion of the results and conclusion in Section 4.

2 Experimental Methodology

Research Questions and Hypothesis. This paper studies the concept of relevance from a neuropsychology aspect by investigating the answer of the following research question: "are there any connections between relevance and brain activity?" In particular our hypothesis is that there exists brain regions for which the activation level are different depending on whether a user has encountered a relevant document or a non-relevant document for the same information need.

MRI Scanner. To test our hypotheses, we scanned participants using a research MRI device (shown in Figure 1 (A)) while they assessed the relevance of different images. The participants performed the experiment while they were lying supine in the bore of the MRI scanner. They could see the images presented to them by looking at the mirror above their eyes (shown in Figure 1 (C)) which reflected images from a data projector that was aimed down the bore of the scanner. Participants could make judgments about the images by pressing buttons on a button box (Figure 1 (B)) as instructed.

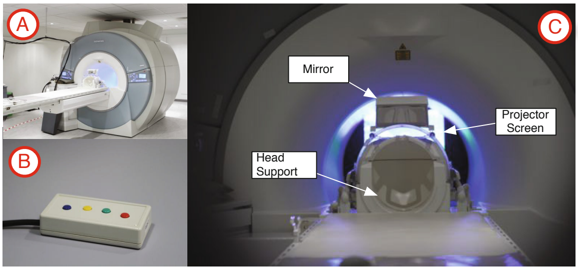

Fig. 1. MRI scanner (A), button box (B), and mirror and projector screen (C)

The MRI scanner can provide high resolution (approximately 1 mm) anatomical images of brain structure, and lower resolution (approximately 3mm) images of functional brain activity. These functional images of brain activity provided by the MRI scanner (fMRI images) can be produced every few seconds in series of 3D volumes (i.e., 4D data). The high spatial (but low temporal) resolution and whole brain coverage that characterises the fMRI technique allows spatial localisation of brain activation during the relevance assessment process.

The underlying principle behind fMRI measurements is the deoxygenation of haemoglobin as a consequence of an increase in neural activation [15]. The image contrast obtained through this form of MRI is called "BOLD" (Blood Oxygenation Level Dependent) because it varies with the amount of deoxyhemoglobin present in blood. Since the fMRI technique consists of the association of the hemodynamic activity related to brain tissue activity with mental processes it is considered an indirect approach [15].[2] The remainder of this section discusses the physical and technical limitations associated with an fMRI study.

Physical Limitations. As described earlier, during the scanning time, participants must lay down inside the MRI scanner. In addition, participants are required to maintain their heads in the same position without any movement, otherwise brain images may suffer distortion. Lighting conditions inside the tube, the participants position, and the overall laboratory environment, may cause discomfort and fatigue in participants which could be reflected in cognitive performance. To minimise the effects associated with physical limitations, the experimenter was in contact with participants from time to time to make sure they were meeting the requirements to proceed with the experiment. This contact was made using an intercom system since the experimenter was in an adjacent control room. The experimenters also had visual contact with the participant throughout the experiment using a window to the MRI scanner room.

Technical Limitations. Obtaining a task that approximates a natural search task is constrained by the technical limitations of the fMRI system. For example, the temporal resolution of measurements is in the order of two seconds, and this is compounded by the linkage to the hemodynamic response [16]. While the underlying neural events may take place on the order of tens to hundreds of milliseconds [17], the haemodynamic response has a slow response time (i.e., seconds). When the cognitive events to be measured occur slowly, they can be measured in real-time through fMRI. In contrast, short time scale events may cause evoked fMRI signals to overlap, increasing the difficulty of determining individual events [18].

The design of the fMRI study itself may also alter the task context. For instance, the study employs a basic fMRI block-design that has proven useful in the detection of small differences in the BOLD signals between contrasting events [16]. However, one drawback of the block-design is that it is hard to ensure the average response to a trial-type is not biased by a specific context or history of preceding types [16]. To a certain extent, the block-design limits the order in which stimulus is presented, therefore, constraining complete randomisation of the presented items.

Finally, the interpretation of the results from measured hemodynamic signal is limited by the fact that hemodynamic response indicates an approximately linear relation to underlying neuronal activity [19]. There is a vast controversy surrounding the use of fMRI and interpretation of results [20]. Nonetheless, the high spatial resolution provided by fMRI can provide unique evidence of the organisation of neural interactions and the functional architecture of the typical human brain [21]. Based on these observations, we have devised an experiment to explore the brain activity during relevance assessment.

Design. A within-subjects design was used in this study. The independent variable was the relevance (with two levels: relevant, and non-relevant), which was

[2] For more information on fMRI please refer to [15].

controlled by viewing relevant and non-relevant images corresponding to a task. The dependent variable was brain activities revealed by the BOLD signal.

Task. Four search tasks were selected from the ImageCLEF 2009 - Photo Retrieval Task[3]. The topics were selected by considering the general knowledge of the population in mind, e.g., the concept "Olympic Torch" (topic 11) was chosen rather than "Brussels Airport" (topic 3), assuming that olympic torch is a more familiar concept than Brussels airport. For each topic shown in Table 1, five relevant and five non-relevant images were selected from the relevance assessment set (Qrel) provided for this ImageCLEF task[4]. All images were checked manually, prior to the experiment, to ensure that they clearly belonged to one of the relevant levels (i.e., "relevant" or "non-relevant") depending on their task.

Table 1. The information of the topics selected form the ImageCLEF 2009 - Photo Retrieval Task

Topic No.	Title	Description
2	Fortis Logo	Relevant images will show the Fortis bank building with logo. Images of Fortis logo or building without bank are not relevant. Images of other logos or buildings are not relevant.
11	Olympic Torch	Images are relevant if they show the Olympic torch with fire. Images of people are relevant only if the Olympic torch is clearly visible. Images of Olympic torch without fire are irrelevant.
12	Obama Clinton	Relevant images show photographs of Obama and Hillary Clinton. Images of those two with other people are relevant. Images of only either of them or none of them are irrelevant.
23	British Royals	Relevant images will show photographs of the British royal family. Relevant images must contain at least one member of the British royal family. Images showing royal families from other countries are irrelevant. Images showing other people are irrelevant.

The topics were selected so that they covered a variety of image recognition tasks, such as face recognition, text recognition, and object recognition. This was done in order to reduce the effect of a particular image recognition task from the brain activity results. The images were also selected such that there was a mixture of expectedly easy and difficult images in each block, reflected by the relative assessment processing time. By doing so, we attempted to reduce the effect of task difficulty an each topic. Finally, an instruction message was given with each task description as follow: *"Participant: Please watch the displays. Press button 1 (if the image is relevant) or 2 (if the image is irrelevant)."*

Procedure. This section outlines the flow of the study, from beginning to end. Ethical permission for the study was obtained from the Ethics Committee of the College of Science and Engineering, University of Glasgow. Participants were instructed of the duration of the experiment, which included approximately 6 minutes to perform all four relevance tasks, and approximately 10 minutes to obtain a scan of their anatomical structure. They were informed that they could

[3] For more information please visit `http://www.imageclef.org/2009/Photo`

[4] The image collection used in ImageCLEF 2009 - Photo Retrieval Task were taken from Belga Press Agency (Copyright) and its usage is subject to conditions and agreement with CLEF.

leave at any point in time during the experiment and would still receive payment (the payment rate was £6/hr). They were then asked to sign a consent form. Before participating, participants underwent a safety checks to guarantee that they did not possess any metal items inside or outside of their body, or any other contraindications for scanning, such as certain tattoo inks. They were then provided with gear (similar to a training suit) to wear for the duration of the experiment to avoid interference from any metal objects in their clothes with the fMRI signal. Next, as a training process they were given an example task and a corresponding set of example images in order to familiarise themselves with the procedure. Once they had successfully completed their training task, participants entered the fMRI machine and the experimenter adjusted the settings of the machine to maximise their comfort and vision. Each participant had to complete four search tasks. For each task, the relevant and non-relevant images were shown in separate blocks. The block design is highly efficient in detecting differences between two conditions (i.e., relevant or non-relevant) although it does have the potential to lead to additional cognitive strategies on how relevance is assigned to elements of a block [18]. A schematic representation of the experiment is illustrated in Figure 2. The order in which each participant was introduced to the four tasks was randomised. For each task, the order in which each participant was introduced to its corresponding image blocks (i.e., relevant or non-relevant) was also randomised. Both randomisations were done to reduce any bias these would introduce into the brain data. The images inside each block were also randomly presented to the participants for the same reason.

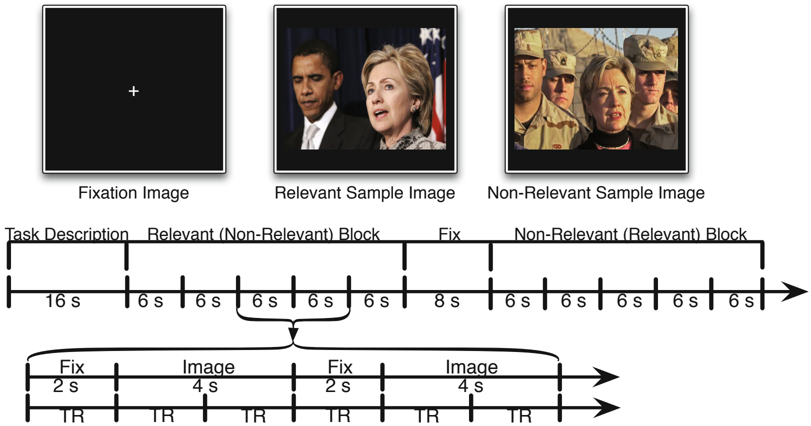

Fig. 2. A schematic representation of the experiment

As each image was presented to the participants, they were asked to assess the relevance of the image with respect to the described task, and to provide feedback using one of the two buttons on the button box. The association between the two buttons being relevant or non-relevant was changed for each participant. Prior to each image, a fixation image was presented for two seconds and then the actual image was presented for four seconds. The fixation image contained a fixation cross, indicating to participants that they needed to fixate on the cross and prepare for the next stimuli. Between each image block an eight second interval was used to reduce the effect of the brain activities of one block being extended to the subsequent block. During this eight second interval, the fixation image was presented to the user. Once the second image block of a task was finished, the next task was selected randomly from the remaining tasks. This process continued until the four tasks were all presented to the participant. Finally, the participants were asked to fill out the exit questionnaire.

Apparatus. The images were presented using Presentation® software[5], and projected using a LCD projector onto a translucent screen, while participants watched them in an angled mirror in the MRI scanner.

fMRI Data Acquisition. All fMRI data was collected using a 3T Tim Trio Siemens scanner at the Centre for Cognitive Neuroimaging. The functional run took 320 seconds. A functional T2*-weighted MRI run was acquired for all combinations of task and condition (TR 2000ms; TE 30ms; 32 Slices; $3mm^3$ voxel; FOV of 210, imaging matrix of 70×70). This resulted in 160 brain volumes obtained during the run. An anatomical scan was performed at the end of the scanning session that comprised a high-resolution T1-weighted anatomical scan using a 3D magnetisation prepared rapid acquisition gradient echo (ADNI- MPRAGE) T1-weighted sequence (192 slices; $1mm^3$ voxel; Sagittal Slice; TR = 1900ms; TE = 2.52; 256×256 image resolution).

Questionnaires. At the end of the experiment, the participants were introduced to an *exit questionnaire*, which gathered background and demographic information. It also enquired about previous experience with fMRI type user studies as well as participants general comments for the user study.

Pilot Studies. Prior to running the actual user study, a pilot study was performed using 4 participants to confirm that the process worked correctly and smoothly. A number of changes were made to the presentation slides based on feedback from the pilot study. After the pilot, it was determined that the participants were able to complete the user study without problems and that the system was correctly logging participants' interaction data.

3 Results Analysis

A study with the procedure explained in Section 2 was conduced over 43 days from 20 April, 2012 to 12 June, 2012. Eighteen participants (8 males and 10 females) with an average age of 27.77 ($\pm$ 4.46) participated in this study. Participants were recruited from the participant database at the School of Psychology, University of Glasgow. All of the participants had experience using search engines and providing explicit relevance judgement.

Data Preprocessing. The fMRI data were preprocessed using FMR preprocessing tools in Brain Voyager QX [22] and involved 3D Motion Correction with

[5] Presentation® software (Neurobehavioral systems, Inc.), `http://www.neurobs.com`

Trilinear detection/sinc interpolation. This was followed by normalisation of functional scans into the common Talairach space [23], and corregistration of functional and anatomical data. Finally, spatial smoothing using a Gaussian kernel of 6mm, linear trend removal, and temporal high pass filter with a cutoff of 2 cycles, were applied to create volume timecourses (VTCs) for each participant.

Log Analysis. In this section we analyse the accuracy of the relevance judgement provided by the participant during the course of the experiment. This is an indication of how well they performed each task and therefore how trustworthy the captured brain activity is across each relevance block. In other words, to what extent participants' brain activity was inline with the stimuli presented to them: the best scenario is that participants judge all images in a relevant block as relevant and all images in a non-relevant block as non-relevant. The results show that the overall accuracy of the participants across all tasks was 0.92 ($\pm$ 0.09) where the accuracy over relevant and non-relevant blocks was 0.94 ($\pm$ 0.08) and 0.9 ($\pm$ 0.12) respectively. Further analysis showed consistent performance over each task among participants where tasks associated with topic 2 and 11 were performed the worst (0.83, $\pm$ 0.15) and the best (0.98, $\pm$ 0.03) respectively.

General Linear Model (GLM) Analysis. Analysis began with a first-level analysis on the data of individual participants using multiple linear regression of the BOLD-response time course in each voxel, using two predictors (Display Type: Relevant, Non-relevant). To achieve this, for each participant's data a BrainVoyager protocol file (PRT) was derived that represented the onset and duration of the events for the different conditions. Predictors' time courses were adjusted for the hemodynamic response delay by convolution with a hemodynamic response function. Following this a second-level analysis was performed with a random effects analysis of variance using Display Type as a within-participants factor. Activations are reported at a threshold of $P < 0.001$ (uncorrected), and were corrected for multiple comparisons using the cluster-size threshold of $P < 0.05$ (for details see [24]) based on a 3D extension of the randomization procedure described in [25]. In this method a whole-brain correction criterion was based on an estimate of the statistical map's spatial smoothness and on an iterative procedure (Monte Carlo simulation) to estimate cluster-level false-positive rates (i.e. the theoretical number of false positive voxels that are activated in each random map). After 1000 iterations, the minimum cluster-size that yielded a cluster-level false- positive rate of 5% was used to threshold the statistical map.

The results based on all 18 participants for the effect of the factor Type of Display are shown in Figure 3 plotted on an average brain and in Table 2. These results showed three clusters in the right hemisphere of the brain, which upon closer examination revealed greater activation for relevant than non-relevant stimuli. These included the medial aspect of the superior frontal gyrus, the inferior parietal lobe and the posterior region of the inferior temporal gyrus near to occipital cortex.

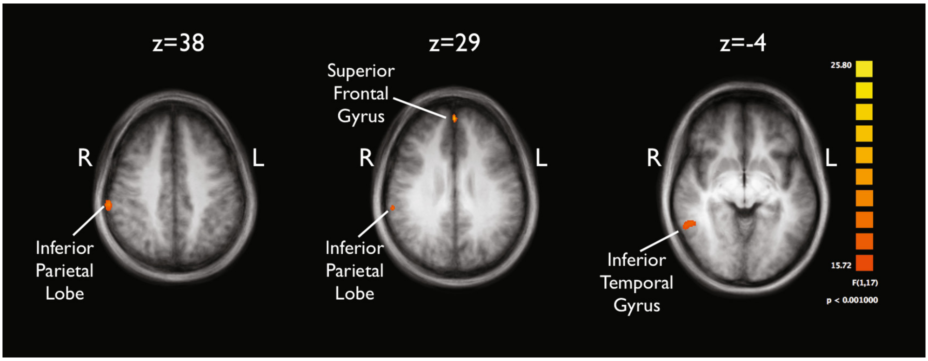

Fig. 3. The three activations are shown projected onto the average anatomical structure for three transverse sections

Table 2. Further details of the activations are provided including their anatomic label, location, Brodmann Area (BA), effect size and volume

		Talairach Coordinates				Effect size		Number of voxels
Brain Area	Hemisphere	X	Y	Z	BA	F(1,17)	p-value	mm^3
Inferior Parietal Lobe	Right	61	-30	38	40	24.80	.0001	462
Inferior Temporal Gyrus	Right	49	-45	-4	37	20.84	.0003	336
Superior Frontal Gyrus	Right	1	51	29	9	24.31	.0001	89

4 Discussion and Conclusion

In this work, we studied the brain regions activated during the process of judging relevance. The results of our experiment contrasting relevant and non-relevant displays show activation in three clusters in the right hemisphere of the brain. These include the medial aspect of the superior frontal gyrus, the inferior parietal lobe and the posterior region of the inferior temporal gyrus near to occipital cortex. Our analysis show that these three regions show greater activations for relevant stimuli than for non-relevant stimuli.

The fMRI data showed activity in frontal, parietal and temporal cortices of the brain related to judging relevance. We can consider these activations at either a modular level of description, where we assume functional specialization of specific brain areas, or as part of networks that distribute a particular function over the whole brain. We will first consider these brain areas in terms of their functional specialization [26]. The region of the superior frontal gyrus where activity was found is known to be involved in processes of memory, cognitive control and inferential reasoning. These regions are also thought to involve amodal representations[6] that form more abstract representations. Specialization

[6] Amodal representations is the way the brain codes multiple inputs such as words and pictures to integrate and create a larger conceptual idea, independent of a particular modality.

involving high level processes is common in many regions of the frontal lobe. The right inferior parietal lobe where activity was found has been implicated in many processes, and this is consistent with the general role of parietal cortex in multisensory representations and visuomotor control. It has also been implicated in same-different discrimination tasks as well as working memory. The region of temporal cortex where activity was found is near the border with occipital (visual) cortex and its activity has been shown to reflect visual processing. For example, a recent experiment has shown activity near this region to involve the processing of natural scenes [27].

Beyond the activity of individual brain areas we can consider the organisation of activity within functional networks. The finding of frontal, parietal and temporal regions in the right hemisphere are consistent with several models of how the task demands of judging relevance might map onto functional brain networks. To be successful at judging relevance, a participant must use executive function to encode the instructions as a cognitive set, and attention to maintain the relevance task as a cognitive set for the duration of an instruction period. During this time they also must encode the visual features of each incoming stimulus and use visuo-spatial working memory to determine whether the stimulus is a target or a distracter. Visuo-spatial working memory has been discussed as a visuospatial sketchpad that is lateralized to the right hemisphere and involves frontal, parietal and occiptitotemporal regions [28–30]. Accordingly, if instead of images the relevance task involved text or speech we can speculate that a left-hemisphere mechanism would be revealed involving aspects of the phonological loop component of working memory [28–30]. Moreover, recent research into visual working memory has attempted to model the interplay of activity in frontal, parietal and occipitotemporal regions in how we recognize and categorize complex visual input [31, 32]. Finally, attention has also been shown to generally activate fronto-parietal circuits [33] and thus the fronto-parietal activity we found might simply reflect greater attention to the relevant than non-relevant displays rather than specific aspects of representing relevance.

Our results of activations from determining relevant versus non-relevant images are broadly consistent with the determination of relevance involving a right hemisphere system of visuospatial processing. Taken as a whole, this would implicate the type of visuospatial sketchpad notion implicated by theories of working memory. However, the current results do not have the power to distinguish between this and other alternative explanations. For example, the activations found in frontal and parietal regions might simply reflect greater attention being focused on the relevant stimuli rather than the actual determination of relevance. Additionally, although the definition of relevant versus non-relevant was suitably complex so that it would require complex higher order representations, it is still possible that the activity in the occipitotemporal cortex area could represent a filter that can flexibly categorize a scene as relevant or non-relevant. The results of the current research establish an important first step in being able to isolate how these different areas are involved in the neural processes involved in deciding the relevance of sensory input.

In conclusion, we examined human brain activity while participants were assessing the relevance of images with respect to a pre-defined task. We devised a user study to investigate the underlying brain activities during the process of topical relevance assessment of image documents. Results of an experiment on eighteen participants showed that three brain regions were activated during the process of relevance judgement and the intensity of their activation was statistically significantly different when participants assess relevant documents than the situation when they assess non-relevant ones. These regions were found in the right hemisphere of the brain and include the medial aspect of the superior frontal gyrus, the inferior parietal lobe and the posterior region of the inferior temporal gyrus near to occipital cortex. In particular we localised brain regions involved in the relevance assessment process, which is helpful to unravel the nature of relevance. Besides informing a theoretical understanding of relevance, these results might ultimately be useful in the development of new relevance feedback techniques. Therefore, the results obtained are an important step towards better understanding the notion of relevance and its potential applications. As future work, we will continue to study underlying brain activity during an information retrieval and seeking process using different document types (e.g. text, video).

Acknowledgement. This work was supported by the EU LiMoSINe project (288024).

References

1. Saracevic, T.: Relevance: A review of the literature and a framework for thinking on the notion in information science. Part II: nature and manifestations of relevance. JASIST 58(13), 1915–1933 (2007)
2. Mizzaro, S.: Relevance: The whole history. JASIS 48(9), 810–832 (1997)
3. Koenemann, J., Belkin, N.J.: A case for interaction: a study of interactive information retrieval behavior and effectiveness. In: SIGCHI, pp. 205–212. ACM (1996)
4. White, R., Kelly, D.: A study on the effects of personalization and task information on implicit feedback performance. In: CIKM, pp. 297–306 (2006)
5. Arapakis, I., Athanasakos, K., Jose, J.M.: A comparison of general vs personalised affective models for the prediction of topical relevance. In: SIGIR, pp. 371–378 (2010)
6. White, R.W.: Implicit Feedback for Interactive Information Retrieval. PhD thesis, University of Glasgow (2004)
7. Joachims, T., Granka, L., Pan, B., Hembrooke, H., Gay, G.: Accurately interpreting clickthrough data as implicit feedback. In: SIGIR, pp. 154–161 (2005)
8. Kelly, D., Belkin, N.J.: Display time as implicit feedback: understanding task effects. In: SIGIR, pp. 377–384. ACM (2004)
9. Nichols, D.M.: Implicit rating and filtering. In: Proceedings of the Fifth DELOS Workshop on Filtering and Collaborative Filtering, pp. 31–36 (1997)
10. Kelly, D., Belkin, N.: A user modeling system for personalized interaction and tailored retrieval in interactive ir. ASIST 39(1), 316–325 (2005)
11. Arapakis, I., Moshfeghi, Y., Joho, H., Ren, R., Hannah, D., Jose, J.M.: Enriching user profiling with affective features for the improvement of a multimodal recommender system. In: CIVR (2009)

12. Lorigo, L., Haridasan, M., Brynjarsdóttir, H., Xia, L., Joachims, T., Gay, G., Granka, L., Pellacini, F., Pan, B.: Eye tracking and online search: Lessons learned and challenges ahead. JASIST 59(7), 1041–1052 (2008)
13. Arapakis, I., Konstas, I., Jose, J.M.: Using facial expressions and peripheral physiological signals as implicit indicators of topical relevance. In: MM, pp. 461–470 (2009)
14. Ingwersen, P., Järvelin, K.: The turn: Integration of information seeking and retrieval in context, vol. 18. Springer (2005)
15. Ogawa, S., Lee, T., Kay, A., Tank, D.: Brain magnetic resonance imaging with contrast dependent on blood oxygenation. PNAS 87(24), 9868–9872 (1990)
16. Friston, K., Buechel, C., Fink, G., Morris, J., Rolls, E., Dolan, R.: Psychophysiological and modulatory interactions in neuroimaging. NeuroImage 6, 218–229 (2007)
17. Robinson, D., Rugg, M.: Latencies of visually responsive neurons in various regions of the rhesus monkey brain and their relation to human visual responses. Biological Psychology 26(1), 111–116 (1988)
18. Huettel, S., Song, A., McCarthy, G.: Functional Magnetic Resonance Imaging. Sinauer Associates (2009)
19. Heeger, D., Ress, D., et al.: What does fmri tell us about neuronal activity? Nature Reviews Neuroscience 3(2), 142–151 (2002)
20. Bandettini, P.: Twenty years of functional mri: The science and the stories. NeuroImage (2012)
21. Logothetis, N.: What we can do and what we cannot do with fmri. Nature 453(7197), 869–878 (2008)
22. Goebel, R.: BrainVoyager QX, Vers.2.1, Brain Innovation B.V., Maastricht, Netherlands
23. Talairach, J., Tournoux, P.: Co-planar stereotaxic atlas of the human brain, vol. 147. Thieme, New York (1988)
24. Goebel, R., Esposito, F., Formisano, E.: Analysis of functional image analysis contest (fiac) data with brainvoyager qx: From single-subject to cortically aligned group general linear model analysis and self-organizing group independent component analysis. Human Brain Mapping 27(5), 392–401 (2006)
25. Forman, S., Cohen, J., Fitzgerald, M., Eddy, W., Mintun, M., Noll, D.: Improved assessment of significant activation in functional magnetic resonance imaging (fmri): use of a cluster-size threshold. MRM 33(5), 636–647 (1995)
26. Kandel, E.R., Schwartz, J.H., Jessell, T.M., Siegelbaum, S.A., Hudspeth, A.: Principles of Neural Science, 5th edn. McGraw-Hill (2012)
27. Peelen, M., Fei-Fei, L., Kastner, S.: Neural mechanisms of rapid natural scene categorization in human visual cortex. Nature 460(7251), 94–97 (2009)
28. Baddeley, A., Hitch, G., et al.: Working memory. The Psychology of Learning and Motivation 8, 47–89 (1974)
29. Baddeley, A.: The episodic buffer: a new component of working memory? Trends in Cognitive Sciences 4(11), 417–423 (2000)
30. Mildner, V.: The Cognitive Neuroscience of Human Communication. Taylor and Francis (2008)
31. Christophel, T., Hebart, M., Haynes, J.: Decoding the contents of visual short-term memory from human visual and parietal cortex. The Journal of Neuroscience 32(38), 12983–12989 (2012)
32. Linden, D., Oosterhof, N., Klein, C., Downing, P.: Mapping brain activation and information during category-specific visual working memory. Journal of Neurophysiology 107(2), 628–639 (2012)
33. Corbetta, M., Shulman, G., et al.: Control of goal-directed and stimulus-driven attention in the brain. Nature Reviews Neuroscience 3(3), 215–229 (2002)