

# Preliminary Studies on Personalized Preference Prediction from Gaze in Comparing Visualizations

Hamed R.-Tavakoli<sup>1</sup>(✉), Hanieh Poostchi<sup>1</sup>, Jaakko Peltonen<sup>1,2</sup>,  
Jorma Laaksonen<sup>1</sup>, and Samuel Kaski<sup>1</sup>

<sup>1</sup> Department of Computer Science, Aalto University, Espoo, Finland  
[hamed.r-tavakoli@aalto.fi](mailto:hamed.r-tavakoli@aalto.fi)

<sup>2</sup> School of Information Sciences, University of Tampere,  
Tampere, Finland

**Abstract.** This paper presents a pilot study on the recognition of user preference, manifested as the choice between items, using eye movements. Recently, there have been empirical studies demonstrating user task decoding from eye movements. Such studies promote eye movement signal as a courier of user cognitive state rather than a simple interaction utility, supporting the use of eye movements in demanding cognitive tasks as an implicit cue, obtained unobtrusively. Even though eye movements have been already employed in human-computer interaction (HCI) for a variety of tasks, to the best of our knowledge, they have not been evaluated for personalized preference recognition during visualization comparison. To summarize the contribution, we investigate: “How well do eye movements disclose the user’s preference?” To this end, we build a pilot experiment enforcing high-level cognitive load for the users and record their eye movements and preference choices, asserted explicitly. We then employ Gaussian processes along with other classifiers in order to predict the users’ choices from the eye movements. Our study supports further investigation of the observer preference prediction from eye movements.

## 1 Introduction

Humans face a preference dilemma in daily life — we are unremittingly choosing between alternatives and preferring one over others. Preference prediction has been an interesting topic of research in many areas such as image quality assessment, information retrieval, advertisement, recommender systems, human-computer interaction, etc. The preference prediction can be carried out by exploiting explicit user feedback as in recommender systems, which employ the explicitly recorded history of actions, as well as implicit feedback. The implicit feedback cues are the pieces of information that complement the explicit feedback and are often obtained unobtrusively.

The range of available implicit feedback information varies depending on the system and applications. For example, in an information retrieval system, the number of clicks, the time spent on a web-page, and revisits are exemplars of

traditional implicit feedback cues [1]. With the advent of new sensors for human-computer interaction, the source of implicit feedback can include human bio-signals such as heart-rate, brain signals, and eye movements. The eye movements are the topic of the current study. While eye movements have already been utilized in several systems as a source of implicit feedback, e.g., [2,3], to the best of the authors' knowledge, there exists no study signifying whether the eye movements carry any useful information conveying the preference of an observer for the task of comparing two visualizations. Thus, we are seeking the answer to "How well do eye movements disclose user's preference?"

The current paper presents our experiment setup and the preliminary results of preference prediction using pure eye movements. We explore the usefulness of several fixation-based features and demonstrate successful above-chance prediction using them. The results motivate further in-depth investigation of eye movements and their contribution in preference prediction.

## 2 Related Work

Inference from eye movements is well-recognized by the seminal work of Yarbus [4], in which it is hypothesized that the observers' eye movement patterns change with respect to the task. In recent studies, [5,6] demonstrated that, using features extracted from eye movements, it is possible to decode observers' task. [7] exploited a hidden Markov model architecture to encode the fixation locations for each of the seven tasks defined by Yarbus, including: wealth estimation, age estimation, remembering position of objects and people, etc., and achieved the state-of-the art in observer task decoding. In HCI, a highly related area is user activity recognition, where eye movements have been utilized. For example, [8] used features from electrooculography (EOG) in order to discriminate several activities of a computer user, including: reading, browsing, writing, watching video, and copying operation. Another relevant area is user interface design, e.g., web interfaces by learning the user's attention location [9].

The eye movements also convey the emotional state of the observer. For example, [10] demonstrated the correlation of positive affect and fixation duration. In a similar vein, [11] studied various eye movement properties, such as saccade angular behaviour and saccade length, with respect to the valence and arousal of the stimulus. In computer vision, [12] exploited eye movements for determining the pleasantness of an image. It was followed by [13] who analyzed the contribution of each feature for the image pleasantness recognition task.

Search target prediction from gaze is another related area. [14,15] studied the role of fixations in categorical search tasks. In a series of experiments, they investigated the number of fixations prior to finding a target and the percentage of fixations landing on the target. Later, they tried to predict the search tasks from fixations. Similar efforts have been made by others in different setups, e.g. [16,17]. In [17], an open-world setting is proposed, i.e., there is no assumption about the fixations and the target of interest. They, however, rely on book covers as search targets and perform the prediction based on the attended locations.

That is equivalent to applying an attention model to sample visual features followed by feature matching in order to spot a search target.

Perceptual image quality assessment is also a related area. In quality assessments, observers are often asked to choose from a pair of an original image and its perturbed version in order to provide a ground-truth, e.g. [18]. Motivated by the role of the human vision system, a group of algorithms proposed for such a task rely on the visual saliency and attention models, e.g. [19,20]. Under such a setup, the image quality assessment can be seen as a preference prediction task where the attention models replicate the observers' eye movement statistics.

The preference prediction task is a mixture of observer task decoding and emotion recognition. Some of the early systems utilizing eye movements in preference prediction can be found in the information retrieval community. SUIOR [2] is a gaze-based attentive information retrieval system. It uses gaze as an input to the system and, depending on what the user looks at, it fetches more alike information or, if the user looks at a headline long enough, the system will fetch the news content. A more predictive system is [21], which uses gaze as a clue for ranking documents in order to build a personalized recommendation system. In their system, the distance to gaze point is used as a weight for each document. The user browses the retrieved documents and the system re-ranks the unseen results after several iterations by learning from the history of gazed items. A step towards eye movement incorporation in user preference prediction is [3]. It employs probabilistic modelling fused with a collaborative filtering mechanism in order to perform proactive information retrieval where gaze is an implicit feedback cue.

While there have been efforts for incorporating the eye movements into preference prediction, there has not been much investigation about the usefulness of the gaze signal for such a purpose. Furthermore, the influencing parameters are not lucid enough. To address this shortcoming, we build a pilot experiment where the setup akin to image quality assessment requires comparison of two panels, consisting of textual information. In other words, instead of natural images, the panels depict visualizations of keyword clouds related to a given query term. An observer is then instructed to choose the visualization that he finds to be better. The use of such visualizations imposes a high-level cognitive load that requires both reading and thinking. Having the eye movements of several observers recorded, we try to predict the observers' preferences from their eye movements.

### 3 Data

To assess the usefulness of eye movements for the task of preference prediction, we require data reflecting the high-level cognitive state of the observers in a preference prediction task. We introduce a pilot experiment following the image quality assessment setup, where one evaluates two image panels displayed side by side. Nonetheless, we are exploiting visualization of textual information that requires careful study and minimizes the aesthetic effect of stimuli on the observers. Figure 1 shows an example visualization used in the experiments.



**Fig. 1.** Experiment setup. On the left, two visualization for the query term “supervised learning”; on the right, the heatmap of an observer’s eye movements and his preferred panel in red. (Color figure online)

**Participants.** Six participants (3 male and 3 female) took part in the experiment. All were computer science graduate students, majoring in machine learning, from Aalto University. The participants had normal or corrected-normal vision.

**Apparatus.** The observers sat 70 cm away from a 22-inch LCD monitor screen, subtended approximately to  $36^\circ \times 24^\circ$  of visual angle. A chin rest was used to minimize head movements. Stimuli were presented at 60 Hz at the resolution of  $1680 \times 1050$ . The eye movements were recorded using a SMI RED500 eye tracker with the spatial resolution of  $0.03^\circ$ . The sampling rate was 500 Hz. SMI’s standard 9 point calibration procedure was applied and we made sure that the spatial error is less than  $1^\circ$  before proceeding with the recording.

**Design and Procedure.** The observers were asked to assess two visualizations, shown simultaneously side by side, and choose the one which looks more appealing to them. The visualizations consist of keywords which corresponds to predefined query terms. For all the observers, the query terms are the same and the keywords are identical in the visualizations, where only their relative locations vary between the panels. For a given query term, we highlight several keywords in green. Then, the observers choose the view in which they find the set of relevant keywords are visualized better. When the observer has determined this, he/she signals the system to stop the eye movement recording procedure and then explicitly chooses the better visualization panel. To control the observers’ vigilance and selection, we recorded their explicit feedback for at least one more relevant keyword immediately after choosing their preferred visualization.

All of the observers assess the same query terms and perform at least seven successive evaluations. This results in a total of 58 evaluations over all the observers and their iterations. There is no constraint on the duration of each evaluation, i.e., the observers can spend as long as they like to explore the visualization panels and discover the relevance of keywords with the given query term, meanwhile their eye movements are recorded. An example heatmap from an observer’s eye movements is depicted in Fig. 1.

## 4 Method

We are interested to determine the user's preference or choice from his eye movements by giving him two options to choose from. To this end, we exploit features extracted from the user's eye movements on the keyword clouds, described above.

### 4.1 Features

Features extracted from eye movements often fall into two categories: fixation-based features and saccade-based features. The first type is often demonstrated of more influential role compared to the saccade-based features in task decoding experiments, e.g., [7]. Therefore, in this work, we rely on fixation-based features extracted from fixation location, fixation duration, and pupil diameter during the fixation period.

**Fixation Location.** A key feature in determining the observers' task is fixation location [5–7], where the viewing pattern is a decisive factor in answering a question in regard to the Yarbus experiment [4]. In general, the fixation location not only conveys the attended object/area of interest, but it also carries the emotional message induced by the stimulus [22]. In our experiment, the fixated locations indicate the keywords, perceived by the user. Contrary to traditional task decoding approaches, which encode the exact fixation locations to maximize the role of the viewing pattern, we rather prefer to minimize the role of the viewing patterns in order to neutralize the effect of the aesthetic aspect of the visualizations. Therefore, we encode the fixation location as the *entropy* of the fixation density map.

**Fixation Duration.** Thus far the most cited feature, which is believed to convey the cognitive load of an observer is fixation duration. In particular, reading tasks are well demonstrated to influence the fixation duration [23, 24]. To encode the fixation duration, we empirically studied the duration minimum, maximum, mean, mode, and histogram representation of [12]. The histogram representation was working best for our data, similar to [13]. We performed a rapid optimization scheme for the number of bins, and a histogram of 200 bins was selected. We only report the results of the experiments with such a histogram.

**Fixation Dispersion.** An indicator of how gaze is dispersed during a fixation event. The fixation dispersion is caused by involuntary eye movements such as tremor, drift, and microsaccades, which are nuance saccadic eye movements. It is affected by various parameters, including: target's shape, size, color, and luminance [25]. We consider fixation dispersion as a potential indicator meanwhile deciding about a keyword. Akin to fixation duration, we tested several representations and eventually adopted a 10-bin histogram representation.

**Pupil Diameter.** The pupil diameter is associated with the working memory. There exists various states of the mind detectable by changes in pupil diameter [26], e.g., the recall process causes a dilation followed by an erosion in the pupil diameter [27]. We encoded the pupil diameter information as a 20-bin

histogram, which performed better than other representations. We must, however, signify that the pupil diameter is sensitive to environmental noise such as illumination changes and needs more careful setup in an HCI scenario.

## 4.2 Preference Prediction

Given all the information from previous preferences of the same user, we are interested to predict his preference for the  $i$ -th pair of visualization instances. For two visualizations and the preference relation  $v_i \succ \hat{v}_i$ , we transfer the previous instances to pairs of samples with labels  $+1$  and  $-1$ , for the selected and non-selected panels, respectively. Then, we train a classifier to predict the user preference for the  $i$ -th instance using the information from the previous  $(i - 1)$  preference records. Hence, the problem is a binary classification. In other words, suppose a feature vector  $\mathbf{x} \in \mathbb{R}^{D \times 1}$  corresponds to a binary class variable  $c \in \{-1, +1\}$ , where we have  $N = 2 \times (i - 1)$  observations, denoted as  $\mathcal{D} = \{(\mathbf{x}_j, c_j)\}_{j=1}^N$ , and  $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ ,  $\mathbf{c} = \{c_1, \dots, c_N\}$ . We are then interested in inferring a classification, denoted as  $\mathbf{c}_*$ , from the observations in order to assign a new feature vector  $\mathbf{x}_*$ , obtained from the  $i$ -th visualization instances, to one of the two classes with a certain degree of confidence. To this end, we employ a Gaussian Process (GP) [28], briefly explained in this section.

A Gaussian Process, denoted as  $\mathcal{GP}(\mu(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'))$ , is a stochastic process determined by a mean function  $\mu(\mathbf{x})$ , and a kernel function  $k(\mathbf{x}, \mathbf{x}')$ . While for a pair of an observation and a real-valued output, there exists an easy analytical predictive distribution, there is no straight analytically tractable solution for predictive distribution of categorical data. Thus, we need to employ the GP prior on a mapping from the input observations to a set of latent decision margin variables and apply an approximation technique, such as Laplace approximation [29] or expectation propagation [30] for inference. We choose the latter scheme under a probit model, which results in the predictive probability distribution.

$$p(c_* | \mathbf{x}_*, \mathbf{X}, \mathbf{c}) = \Phi\left(\frac{\mathbf{k}_*^T (\mathbf{K} + \tilde{\Sigma})^{-1} \tilde{\boldsymbol{\mu}}}{\sqrt{1 + k_* - \mathbf{k}_*^T (\mathbf{K} + \tilde{\Sigma})^{-1} \mathbf{k}_*}}\right), \quad (1)$$

where  $\mathbf{k}_*$  is the vector of kernel responses for  $\mathbf{x}_*$  and each training point  $j$ , and  $k_*$  is the kernel self-response over  $\mathbf{x}_*$ ,  $\tilde{\boldsymbol{\mu}}$  is the vector of  $\tilde{\mu}_j$  and  $\tilde{\Sigma}$  is the diagonal with  $\tilde{\Sigma}_{jj} = \tilde{\sigma}_i^2$ . The tilde indicates that the parameters are corresponding to the local likelihood approximations. (Please consult [28] for the derivation.)

To determine the appropriate kernel, we empirically evaluated three kernels, including: linear, exponential and squared exponential kernels of which the squared exponential was performing the best and the results are reported using it. We use the implementation of [31] to estimate (1) and determine the preference. As alternatives to GP, we also study the performance of  $k$ -nearest neighbour classifier, for  $k = 3$  (3NN), the logistic regression (lreg), the robust boosting classifier, and the SVM with linear and RBF kernels.

## 5 Results

To predict the observer preference for a given visualization instance, we train the classifier on all the previous instances of the data. That is, for the  $i$ -th evaluation, there exists  $2 \times (i - 1)$  training samples. We preserve the original order of evaluations for each observer and guaranteed that there exists at least 5 iterations in the training. To be more accurate, for the first evaluation, we train the 6th iteration on the data from the all the 5th prior iteration. For the second evaluation, we retrain over the data from iteration 1 to 6, and predict iteration 7. We continue this process until all the available iterations of a user are used for the prediction and evaluation of his preference. The feature parameters are decided on the first 5 iterations by taking the 4-th and the 5-th iterations as validation, meanwhile training on iteration 1 to 3. We, however, fix the classifier parameters empirically due to the limited data.

To evaluate the performance of a classifier and a feature, we report the accuracy of the predictions for all the instances of all observers in the preference prediction task, that is, the ratio of the number of correct evaluations to the number of total evaluations. It is worth noting that for each observer only his own preference record is used. In order to obtain an insight about the difficulty of the preference prediction, we also extract *baseline features* from the visualizations, where the average distance to a query term in each visualization is used as a feature. Then the same classification scheme is employed. We identify such features as ‘baseline’ features in the rest of the paper.

Table 1 summarizes the performance of each feature in observer preference prediction in a two panel visualization comparison setup. Using the 3NN classifier as a classification baseline, we learn that the fixation-based features are doing

**Table 1.** Comparing fixation-based features and baseline features using various classifiers. For each feature, the best and runner-up accuracy values are highlighted with green and red colors, respectively.

Features	GP	3NN	lreg	Boosting	SVM	
					Linear	RBF
Entropy of fixation density map	62.50 %	33.33 %	50.00 %	45.83 %	62.50 %	37.50 %
Histogram of fixation duration	62.50 %	41.67 %	66.67 %	33.33 %	33.33 %	45.83 %
Histogram of fixation dispersion	50.00 %	29.17 %	37.50 %	33.33 %	41.67 %	54.17 %
Histogram of pupil diameter	62.50 %	33.33 %	66.67 %	50.00 %	33.33 %	41.67 %
Fixation-based feature fusion	58.33 %	50.00 %	45.83 %	41.67 %	66.67 %	37.50 %
Baseline feature	58.33 %	16.67 %	54.17 %	33.33 %	37.50 %	41.67 %

significantly better than baseline features. On the average, a similar behaviour is also observed for most of the classifiers, albeit not with all the fixation features. The performance of the fixation-based features indicates that eye movements carry somewhat meaningful information for predicting the user preference.

As summarized in Table 1, the classification performance of GP and logistic regression are above chance for at least two features. The linear SVM performs above chance for the fixation location, while it is not doing well for the other features. While logistic regression achieves maximum accuracy for the histogram of fixation duration and the pupil diameter, the GP performs on average better over all the features indicating that it is more robust than logistic regression in handling various features.

## 6 Discussion and Conclusion

We performed a pilot study in order to investigate the feasibility of observer preference prediction from his eye movements. To this end, we designed an experiment imposing cognitive load on the observers by asking them to evaluate two visualizations meanwhile recording their eye movement signal. The observer preferences were recorded explicitly after the evaluation process was over, preventing interaction bias in the eye movement recordings.

The pilot study consisted of six observers of which we empirically noticed that the prediction of the preference of two individuals was more difficult than others. This indicates the effect of individual differences and necessitates a larger number of observers in the later studies.

The preliminary results of the current pilot study support overall preference decoding from eye movements of the observers. The experiments were, however, carried out under simplified conditions where the statistics of the eye movements over two panels were exploited. While such a simplification facilitates gaze point to item association, it is not always possible to have such a user interface. Therefore, future studies will need to investigate more sophisticated user interface scenarios, where a well-designed user interfaces and robust gaze estimation algorithms are necessary.

We did not study saccade-based features, such as saccade length, saccade velocity, etc. The saccade-based features are, however, capable of conveying observers' cognitive load, albeit not as well as the fixation-based features. Showing that preference prediction is doable by fixations, the saccade-based features are also worth being investigated and need to be addressed later.

To summarize, we performed a pilot study to predict observer preference from his implicit gaze feedback. The preference prediction seems to be a difficult task, where the baseline features, extracted from the visualization data are outperformed by the eye movements of the observers. Overall, the results motivate further investigation of eye movements for preference prediction.

**Acknowledgement.** The authors would like to acknowledge the support of the Finnish Center of Excellence in Computational Inference Research (COIN), the Revolution of Knowledge Work 2 project, and Academy of Finland decision 295694.



## References

1. Kelly, D., Teevan, J.: Implicit feedback for inferring user preference: a bibliography. *SIGIR Forum* **37**, 18–28 (2003)
2. Maglio, P.P., Barrett, R., Campbell, C.S., Selker, T.: SUITOR: an attentive information system. In: *Proceedings of the 5th International Conference on Intelligent User Interfaces* (2000)
3. Puolamäki, K., Salojärvi, J., Savia, E., Simola, J., Kaski, S.: Combining eye movements and collaborative filtering for proactive information retrieval. In: *Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (2005)
4. Yarbus, A.L.: *Eye Movements and Vision*. Plenum Press, New York (1967)
5. Borji, A., Itti, L.: Defending Yarbus: eye movements reveal observers' task. *J. Vis.* **14**, 29 (2014)
6. Kanan, C., Ray, N.A., Bseiso, D.N.F., Hsiao, J.H., Cottrell, G.W.: Predicting an observer's task using multi-fixation pattern analysis. In: *Proceedings of the Symposium on Eye Tracking Research and Applications* (2014)
7. Haji-Abolhassani, A., Clark, J.J.: An inverse Yarbus process: predicting observers' task from eye movement patterns. *Vis. Res.* **103**, 127–142 (2014)
8. Bulling, A., Ward, J., Gellersen, H., Troster, G.: Eye movement analysis for activity recognition using electrooculography. *PAMI* **33**, 741–753 (2011)
9. Buscher, G., Cutrell, E., Morris, M.R.: What do you see when you're surfing? Using eye tracking to predict salient regions of web pages. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2009)
10. Tichon, J.G., Mavin, T., Wallis, G., Visser, T.A.W., Riek, S.: Using pupillometry and electromyography to track positive and negative affect during flight simulation. *Aviat. Psychol. Appl. Hum. Factors* **4**, 23–32 (2014)
11. Simola, J., Fevre, K.L., Torniainen, J., Baccino, T.: Affective processing in natural scene viewing: valence and arousal interactions in eye-fixation-related potentials. *NeuroImage* **106**, 21–33 (2015)
12. Tavakoli, H.R., Yanulevskaya, V., Rahtu, E., Heikkilä, J., Sebe, N.: Emotional valence recognition, analysis of salience and eye movements. In: *22nd International Conference on Pattern Recognition* (2014)
13. R.-Tavakoli, H., Atyabi, A., Rantanen, A., Laukka, S.J., Nefti-Meziani, S., Heikkilä, J.: Predicting the valence of a scene from observers' eye movements. *PLoS ONE* **10**, e0138198 (2015)
14. Zelinsky, G.J., Peng, Y., Samaras, D.: Eye can read your mind: decoding gaze fixations to reveal categorical search targets. *J. Vis.* **13**, 10 (2012)
15. Zelinsky, G., Adeli, H., Peng, Y., Samaras, D.: Modelling eye movements in a categorical search task. *Philos. Trans. R Soc. Lond. B Biol. Sci.* **368**, 20130058 (2013)
16. Borji, A., Lennartz, A., Pomplun, M.: What do eyes reveal about the mind? Algorithmic inference of search targets from fixations. *Neurocomputing* **149**, Part B, 788–799 (2015)
17. Sattar, H., Mller, S., Fritz, M., Bulling, A.: Prediction of search targets from fixations in open-world settings. In: *2015 IEEE Conference on Computer Vision and Pattern Recognition* (2015)
18. Ponomarenko, N., Jin, L., Ieremeiev, O., Lukin, V., Egiazarian, K., Astola, J., Vozel, B., Chehdi, K., Carli, M., Battisti, F., Kuo, C.C.J.: Image database TID2013: peculiarities, results and perspectives. *Sig. Process. Image Commun.* **30**, 57–77 (2015)

19. Zhang, L., Shen, Y., Li, H.: VSI: a visual saliency-induced index for perceptual image quality assessment. *IEEE Trans. Image Process.* **23**, 4270–4281 (2014)
20. Zhang, W., Borji, A., Wang, Z., Callet, P.L., Liu, H.: The application of visual saliency models in objective image quality assessment: a statistical evaluation. *IEEE Trans. Neural Netw. Learn. Syst.* **27**, 1266–1278 (2016)
21. Xu, S., Jiang, H., Lau, F.C.: Personalized online document, image and video recommendation via commodity eye-tracking. In: *Proceedings of the 2008 ACM Conference on Recommender Systems, RecSys 2008*, pp. 83–90. ACM, New York (2008)
22. Wadlinger, H., Isaacowitz, D.: Positive mood broadens visual attention to positive stimuli. *Motiv. Emot.* **30**, 87–99 (2006)
23. Just, M., Carpenter, P.: A theory of reading: from eye fixations to comprehension. *Psychol. Rev.* **87**, 329–354 (1980)
24. Vitu, F., McConkie, G., Kerr, P., O'Regan, J.: Fixation location effects on fixation durations during reading: an inverted optimal viewing position effect. *Vis. Res.* **41**, 3513–3533 (2001)
25. Thaler, L., Schütz, A., Goodale, M., Gegenfurtner, K.: What is the best fixation target? The effect of target shape on stability of fixational eye movements. *Vis. Res.* **76**, 31–42 (2013)
26. Kahneman, D., Beatty, J.: Pupil diameter and load on memory. *Science* **154**, 1583–1585 (1966)
27. Johnson, E.L., MillerSingley, A.T., Peckham, A.D., Johnson, S.L., Bunge, S.A.: Task-evoked pupillometry provides a window into the development of short-term memory capacity. *Front. Psychol.* **5**, 218 (2014)
28. Rasmussen, C.E., Williams, C.K.I.: *Gaussian Processes for Machine Learning*. MIT Press, Cambridge (2006)
29. Williams, C.K.I., Barber, D.: Bayesian classification with Gaussian processes. *IEEE Trans. Pattern Anal. Mach. Intell.* **20**, 1342–1351 (1998)
30. Minka, T.: *A family of algorithm for approximate Bayesian inference*. Ph.D. thesis, MIT (2001)
31. Vanhatalo, J., Riihimäki, J., Hartikainen, J., Jylänki, P., Tolvanen, V., Vehtari, A.: GPstuff: Bayesian modeling with Gaussian processes. *J. Mach. Learn. Res.* **14**, 1175–1179 (2013)