

Visual Analytics of Work Behavior Data - Insights on Individual Differences

S. Koldijk^{1,2}, J. Bernard³, T. Ruppert³, J. Kohlhammer³, M.A. Neerincx^{2,4} and W. Kraaij^{1,2}

¹ICIS, Radboud University, The Netherlands; ²TNO, The Netherlands;
³Fraunhofer IGD, Germany; ⁴II, Delft University of Technology, The Netherlands

Abstract

Stress in working environments is a recent concern. We see potential in collecting sensor data to detect patterns in work behavior with potential danger to well-being. In this paper, we describe how we applied visual analytics to a work behavior dataset, containing information on facial expressions, postures, computer interactions, physiology and subjective experience. The challenge is to interpret this multi-modal low level sensor data. In this work, we alternate between automatic analysis procedures and data visualization. Our aim is twofold: 1) to research the relations of various sensor features with (stress related) mental states, and 2) to develop suitable visualization methods for insight into a large amount of behavioral data. Our most important insight is that people differ a lot in their (stress related) work behavior, which has to be taken into account in the analyses and visualizations.

Categories and Subject Descriptors (according to ACM CCS): I.5.4 [Pattern recognition]: Applications—Signal processing; H.5.0 [Information Interfaces and Presentation]: General—

1. Introduction

Stress at work is a serious problem, in the worst case leading to burn-out. The goal of the SWELL project (<http://www.swell-project.net>) is to help employees to detect patterns in work behavior with potential danger to well-being in time. Trends like ‘quantified self’ show a potential of collecting personal sensor data (e.g. heart rate, activity patterns) for health improvement. Vast amounts of data from different sensors can be recorded, yielding a multi-modal, time oriented, and multivariate data set. Interpreting this data in a meaningful way is challenging. In this paper we describe how we applied visual analytics to a work behavior dataset. We alternated between data analysis to find structures in the data and visualizations to gain insights. This exploratory analysis was aimed at finding relations between work stress and behavior that can be measured with sensors.

Data: We used the multi-modal SWELL-KW dataset [KSV*14], which was collected in an experiment in which 25 people performed typical knowledge work (writing reports, making presentations, reading e-mail, and searching for information) for about 3 hours. We manipulated working conditions. In the neutral condition, participants were instructed to work as they would usually do. In one stressor condition, the participants got email interruptions and in the other stressor condition, participants worked under time

pressure. As ground truth, questionnaire ratings of the participants were accessed after each condition for task load, mental effort, emotion and perceived stress. Work behavior data was recorded with various sensors: computer logging, camera, Kinect 3D sensor, and physiological body sensors. For the analyses presented here we used the preprocessed feature dataset, which contained averages per minute for several extracted features regarding: computer interaction, facial expressions, body postures, heart rate (variability) and skin conductance. We also had topic labels for what the participants were working on [SVKK14]. The SWELL-KW dataset can thus be described as a multi-modal, time-oriented, and multivariate dataset. The dataset consists of 3000 data entries of 25 different participants (on average 120 data points per participant) and has 150 columns containing different features.

Research question: How can sensor data be used to gain insight into work behavior, specifically related to stress at work?

Contributions: 1) Research into the relations of various sensor features with (stress related) mental states; 2) Focus on differences between users and how to cope with them in data processing and visualization; 3) Towards visualizations for a large amount of behavioral data recorded with sensors.

2. Related work

2.1. Affective computing

Regarding the relation between stress and sensor data, Sharma and Gedeon [SG12] provide a compact survey. Often, **body sensors** are used to measure the physiological stress response directly (e.g. skin conductance [BHK*12], heart rate [HBV13]). There also is potential in using outward characteristics, such as facial expressions, postures or computer interactions as indicators for the user's mental state. **Facial expressions** are currently mainly used for inferring emotions, but people might also show facial expressions indicative of mental states. Craig et al. [CDWG08] looked at facial expressions while students worked with an online tutoring system. Association rule mining identified that frustration was associated with activity in facial action units 1, 2 (inner and outer brow raiser) and 14 (dimpler); confusion was associated with AU 4 (brow lowerer), 7 (lid tightener) and 12 (lip corner puller). Moreover, preliminary results by Dinges et al. [DRD*05] suggest that high and low stressor situations could be discriminated based on facial activity in mouth and eyebrow regions. Regarding **postures**, Kapoor and Picard [KP05] present research in which posture data was successfully used for estimating interest (vs. uninterest). Mental states are also being estimated from **computer interaction** data. Results by Vizer, Zhou and Sears [VZS09] indicate that stress can produce changes in typing patterns. We think also general human computer behavior, like task switching or browsing, might be indicative of mental states. In general, the affective computing community often uses (black-box) machine learning algorithms to classify sensor data into mental states (see [SG12]). Often one model is learned over all users. This work is different, as we are interested in finding underlying behavioral patterns related to stress, for individual users. Visualizing behavioral patterns may give users more insight and actionable information than just a stress labeling.

2.2. Visual analytics

The structures of our included data set fit well to the survey of Kehrer and Hauser for the visualization and visual analysis of multi-faceted scientific data [KH13]. Most relevant are the characterizations for multi-modal, multivariate, and (spatio-) temporal data. From a task-based perspective we borrow concepts from exploratory data analysis. In particular, most relevant classes of techniques are feature selection, visual comparison, and feature relation. In our approach we have to identify and **select features** relevant to mental states. Regarding human motion analysis, a visual-interactive exploratory search system was presented in [BWK*13]. The results of the data characterization phase give an indication on the complexity of the feature selection and (pre-) processing tasks for respective data sets. A profound overview of feature selection techniques in general is presented in the survey of Guyon and Elissee [GE03]. Well-known examples for the visual-interactive specification and selection of features are the Polaris system [SH00] and approaches presented by Doleisch et al. [DGH03] and May et al.

[MBD*11]. A **visual comparison** helps to show differences in behavior related to stress. A taxonomy for information visualization tools that support comparison tasks is provided by Gleicher et al. [GAW*11]. They distinguish between juxtaposition, superposition, and explicit encoding. Visual comparison tasks can also be supported through compact representations of multidimensional data with glyph designs. For a recent state-of-the-art report on glyph-based visualization we refer to Borgo et al. [BKC*13]. An important step our approach is the identification of **feature relations** within our multi-modal data set. Multiple linked views are an important class for supporting the visual-interactive exploration of relations. In addition, superposition-based overview visualizations can support users in revealing relations between multi-modal features. In the approaches of Bernard et al., the users are guided towards interesting relations between clusterings of time series data and attached meta-data attributes [BRS*12b, BRS*12a]. In addition, visual-interactive approaches addressing relation tasks for mixed data sets are presented in [KBH06, AAMG12, BSW*14].

3. Data analysis and visual analytics

We carried out an iterative process by moving back and forth between data analysis and data visualization. In this process we gained a variety of insights in the nature of the dataset and its challenges, which we then aimed to address in the next iteration. Here we describe this process in detail.

3.1. General overview

We started with some general analyses on the dataset [KSNK13]. A comparison of our working conditions (neutral, stressor email interruptions, stressor time pressure) with t-tests revealed that people showed differences in computer interactions under stressors. However, t-tests did not reveal an increase in experienced stress under stressors. As this result surprised us, we plotted the difference in stress for each participant in a bar chart. This visualization revealed that for half of the participants perceived stress was higher under a stressor, whereas for a quarter of the participants stress was lower. When comparing conditions with t-tests, simple averages over all participants were taken, which kept such effects hidden. This brings us to our first main insight: Data visualization enabled us to view all 25 individual users and eased the identification of effects within subgroups. To be able to look into the data of individual participants, we implemented a visualization for the SWELL-KW dataset (http://cs.ru.nl/~skoldijk/Visualization/ExperimentBrowser/Generic/Gantt_and_Numeric.html). Subjective experience data is displayed in relation to different sorts of data. For sensor data we use line charts to plot specific features over time. We applied machine learning techniques to the computer logging data to infer the current task [KvSNK12]. The categorical variables are plotted as Gantt charts (see Fig. 1). This is a starting point for giving insight into working behavior. The question that arises is: *Are there interesting relations between subjective variables and features measured by the sensors?*




Figure 1: SWELL-KW dataset browser (extract). Topic and task worked on depicted (green: main task, blue: help task), together with the skin conductance level (measure of stress).




Figure 2: Interactive chord diagram of the correlations found in the SWELL-KW dataset. On hovering over a variable, all its connections become highlighted. The following groups of variables are depicted: questionnaire (red), physiology (gray), facial expressions (green) and computer logging (yellow) (posture features omitted for readability).

3.2. Relating subjective experience to sensor data

To investigate which subjective experience variable (stress, task load, mental effort, emotion; from questionnaires) is most closely related to measurable sensor data we performed a correlation analysis. For easier interpretability we decided to visualize the 150x150 correlation matrix as a chord diagram (Figure 2, see also <http://cs.ru.nl/~skoldijk/Visualization/ExperimentAnalysis/CircleCorrelation-noKinect.html>). On the circle, several variables are depicted. When a correlation stronger than 0.3 was found in the data, a connection is drawn between two variables. In the image, many connections between variables of the same source (same color) can be identified. More interesting for us, however, are connections between subjective variables (red) and the sensors. The most promising relation seems to be the one between mental effort and several features resulting from facial expression analysis (green). The question now is: *How can we read someone's mental effort from facial activity?*

We used our dataset browser to plot the facial activity features most strongly correlated with mental effort, expecting to see clear trends, e.g. increase in lid tightening as indicator of increased mental effort. The results were disappointing. In general there is much fluctuation over time in the sensor data. Moreover, we see much variation between users: for some users specific action units are very active, whereas others show no activity in the specific facial region at all. The features with high correlations overall are not really insightful for individual participants. Some of the more 'extreme' users had a big influence in the correlation analysis. By using our visualization we were able to reveal individual differences, which need to be addressed. The question that arises is: *Can we group participants in their characteristics for further analysis?*

3.3. Typical user behavior groups

For the identification and comparison of different user groups we applied clustering techniques. Hierarchical clustering was used to reveal the amount of clusters (k) in the data and then k-means clustering was applied. We addressed each sensor separately and found that for each sensor the users were grouped differently. Clustering of computer activity data revealed that users differ in how they work: one group could be described as 'copy-pasters' (many special keys, a lot of mouse), the other as writers (many keystrokes). Clustering of facial activity data revealed that one group shows little facial activity, one group shows tight eyes and a loose mouth, whereas the other group shows wide eyes and a tight mouth. Clustering of body movement data revealed that one group sits rather still with their body and moves the arm a lot, another group moves the entire body, the last group just moves average. As users seem to differ in their general behavior, we normalized the data to make different users comparable: from each value, we subtracted the participant's average. Thus we are now focusing on difference scores, e.g. how much more than average someone is frowning. The question is: *Can we find strong features among all participants with this normalized dataset?*

3.4. Filtering the set of features

To find the most relevant features (from the available >150 features) to predict mental effort, we decided to apply (information gain based) feature selection to our normalized dataset. Again, we plotted the best features with our dataset browser. But even with the data normalized per user, the best features in general do not seem to give any insights for individual users. The change in behavior still is very individual. *We decided to focus on individual users and further explore their facial expressions.*

3.5. Exploration of facial activity patterns

To investigate whether there are any meaningful patterns in the facial activity of individual users, we used the Acume behavioral data visualization toolkit [MEKKP11]. The tool

generates heat maps from facial action unit data and facilitates the comparison of different users and different working conditions. A comparison of users showed large individual differences regarding which facial regions show activity in general. When comparing data within a user, however, differences in facial activity between working conditions become visible. Inspired by the Acume toolkit, we implemented a heat map visualization, which enables us to see patterns in different features in one glance (Figure 3). This is an improvement over separate line charts, but *to make the data better interpretable we think an avatar displaying the particular facial expressions would be an useful addition.*




Figure 3: A heat map of facial activity data of 2 participants. Different facial regions show changes in activity.



Figure 4: HapFACS avatar displaying different facial expressions of one participant while working

3.6. Details on demand: Visualizing facial activity

To render facial expressions on an avatar (see Figure 4) we used the HapFACS tool [AL13]. We added this visualization as a detail-on-demand concept to the heat map visualization (see <http://cs.ru.nl/~skoldijk/Visualization/ExperimentBrowser/heatmap/FacialActionUnits2/heatmap3.html>). In addition to seeing general patterns in the heat map, the user can hover over a point in time to get a representation of the actual facial expression. *Ideally we would want to add an indication of which facial expressions typically refer to high mental effort.*

3.7. Grouping typical facial expressions

To find typical facial expressions related to high mental effort, we applied a supervised Self-organizing Map (SOM). The SOM takes facial action unit data together with annotations on mental effort and produces a map in which the regions are ordered according to mental effort. Each cell of the SOM contains a typical facial expression, which was visually encoded with a Chernoff face metaphor (Fig. 5; alternatively the avatars could be used). As a result, we found several facial expressions typically associated with a high or low mental effort.




Figure 5: Left: Self organizing map, sorted on mental effort; Right: Facial expressions related to mental efforts.

4. Conclusion and future work

In an iterative approach, we used automatic data analysis procedures and visualization techniques in order to answer our research question: How can sensor data be used to gain insight into work behavior, specifically related to stress at work? We found that mental effort seems to be most closely related to facial expression features. Which specific facial action units are most informative differs per user. By clustering we were able to identify several user groups. Even after normalizing our data per user, individual differences remained. By means of a heat map we were able to visualize meaningful patterns in facial activity for an individual user. The visualization was made more insightful by rendering facial expressions on an avatar. Finally, we identified several facial expressions that are typically related to a low or high mental effort. We conclude that facial expressions may be a promising measurable outward characteristic that can be visualized to indicate mental state patterns during work.

The benefit of incorporating visual analytics to our problem, instead of a black box machine learning approach, was to gather a deeper understanding of the structures in our data and to gain insights from individual users' data. Important lessons learned are: 1) There are many individual differences. People experience stressors differently, people differ in their usual work behavior and people show different outward behavior under stress; 2) Machine learning techniques that build one general model over all users seem not to make sense under these conditions. Models should be trained on individual users or groups of similar users. 3) A direct mapping from low level sensor data to subjective experience is hard. We rather suggest to first interpret low level data to a higher level e.g. raw computer interactions to tasks, facial action unit activity to meaningful facial expressions.

The presented results can serve as a baseline for a variety of future approaches. In future work, detailed analysis of specific users or user groups should be done, ideally with a dataset containing more labeled data per user. We now mainly focused on facial expressions. The combination of data from different sensors might be interesting to analyze and visualize. The presented analyses and visualizations should be integrated in an interactive visual analytics system. Finally, a user-study should be performed to evaluate whether the resulting visualizations can indeed help employees to detect alarming patterns in work behavior.

Acknowledgment: This publication was supported by the Dutch national program COMMIT (project P7 SWELL).

References

- [AAMG12] ALSALLAKH B., AIGNER W., MIKSCH S., GROLLER M.: Reinventing the contingency wheel: Scalable visual analytics of large categorical data. *Visualization and Computer Graphics, IEEE Transactions on* 18, 12 (2012), 2849–2858. 2
- [AL13] AMINI R., LISETTI C.: Hapfacs: An open source api/software to generate facs-based expressions for ecas animation and for corpus generation. In *Affective Computing and Intelligent Interaction (ACII), 2013 Humaine Association Conference on* (2013), IEEE, pp. 270–275. 4
- [BHK*12] BAKKER J., HOLENDERSKI L., KOCIELNIK R., PECHENIZKIY M., SIDOROVA N.: Stess@ work: From measuring stress to its understanding, prediction and handling with personalized coaching. In *Proceedings of the 2nd ACM SIGHIT International health informatics symposium* (2012), ACM, pp. 673–678. 2
- [BKC*13] BORGIO R., KEHRER J., CHUNG D. H., MAGUIRE E., LARAMEE R. S., HAUSER H., WARD M., CHEN M.: Glyph-based visualization: Foundations, design guidelines, techniques and applications. In *Eurographics State of the Art Reports* (2013), EG STARs, Eurographics Association, pp. 39–63. 2
- [BRS*12a] BERNARD J., RUPPERT T., SCHERER M., KOHLHAMMER J., SCHRECK T.: Content-based layouts for exploratory metadata search in scientific research data. In *Proceedings of the 12th ACM/IEEE-CS Joint Conference on Digital Libraries* (New York, NY, USA, 2012), JCDL '12, ACM, pp. 139–148. 2
- [BRS*12b] BERNARD J., RUPPERT T., SCHERER M., SCHRECK T., KOHLHAMMER J.: Guided discovery of interesting relationships between time series clusters and metadata properties. In *Proceedings of the 12th International Conference on Knowledge Management and Knowledge Technologies* (New York, NY, USA, 2012), i-KNOW '12, ACM, pp. 22:1–22:8. 2
- [BSW*14] BERNARD J., STEIGER M., WIDMER S., LÜCKE-TIEKE H., MAY T., KOHLHAMMER J.: Visual-interactive Exploration of Interesting Multivariate Relations in Mixed Research Data Sets. *Computer Graphics Forum* 33, 3 (2014), 291–300. 2
- [BWK*13] BERNARD J., WILHELM N., KRUGER B., MAY T., SCHRECK T., KOHLHAMMER J.: Motionexplorer: Exploratory search in human motion capture data based on hierarchical aggregation. *Visualization and Computer Graphics, IEEE Transactions on* 19, 12 (2013), 2257–2266. 2
- [CDWG08] CRAIG S. D., D'MELLO S., WITHERSPOON A., GRAESSER A.: Emote aloud during learning with autotutor: Applying the facial action coding system to cognitive-affective states during learning. *Cognition and Emotion* 22, 5 (2008), 777–788. 2
- [DGH03] DOLEISCH H., GASSER M., HAUSER H.: Interactive feature specification for focus+context visualization of complex simulation data. In *Proceedings of the Symposium on Data Visualisation 2003* (Aire-la-Ville, Switzerland, Switzerland, 2003), VISSYM '03, Eurographics Association, pp. 239–248. 2
- [DRD*05] DINGES D. F., RIDER R. L., DORRIAN J., MCGLINCHEY E. L., ROGERS N. L., CIZMAN Z., GOLDENSTEIN S. K., VOGLER C., VENKATARAMAN S., METAXAS D. N.: Optical computer recognition of facial expressions associated with stress induced by performance demands. *Aviation, space, and environmental medicine* 76, Supplement 1 (2005), B172–B182. 2
- [GAW*11] GLEICHER M., ALBERS D., WALKER R., JUSUFI I., HANSEN C. D., ROBERTS J. C.: Visual comparison for information visualization. *Information Visualization* 10, 4 (2011), 289–309. 2
- [GE03] GUYON I., ELISSEEFF A.: An introduction to variable and feature selection. *J. Mach. Learn. Res.* 3 (Mar. 2003), 1157–1182. 2
- [HBV13] HOGERVORST M. A., BROUWER A.-M., VOS W. K.: Physiological correlates of stress in individuals about to undergo eye laser surgery. In *Affective Computing and Intelligent Interaction (ACII), 2013 Humaine Association Conference on* (2013), IEEE, pp. 473–478. 2
- [KBH06] KOSARA R., BENDIX F., HAUSER H.: Parallel sets: Interactive exploration and visual analysis of categorical data. *IEEE Transactions on Visualization and Computer Graphics* 12, 4 (2006), 558–568. 2
- [KH13] KEHRER J., HAUSER H.: Visualization and visual analysis of multifaceted scientific data: A survey. *Visualization and Computer Graphics, IEEE Transactions on* 19, 3 (March 2013), 495–513. 2
- [KP05] KAPOOR A., PICARD R. W.: Multimodal affect recognition in learning environments. In *Proceedings of the 13th annual ACM international conference on Multimedia* (2005), ACM, pp. 677–682. 2
- [KSNK13] KOLDIJK S., SAPPELLI M., NEERINCX M., KRAAIJ W.: Unobtrusive monitoring of knowledge workers for stress self-regulation. In *User Modeling, Adaptation, and Personalization*. Springer, 2013, pp. 335–337. 2
- [KSV*14] KOLDIJK S., SAPPELLI M., VERBERNE S., NEERINCX M., KRAAIJ W.: The swell knowledge work dataset for stress and user modeling research. In *Proceedings of the 16th ACM International Conference on Multimodal Interaction* (2014), ACM. 1
- [KvSNK12] KOLDIJK S., VAN STAALDUINEN M., NEERINCX M., KRAAIJ W.: Real-time task recognition based on knowledge workers' computer activities. In *Proceedings of the 30th European Conference on Cognitive Ergonomics* (2012), ACM, pp. 152–159. 2
- [MBD*11] MAY T., BANNACH A., DAVEY J., RUPPERT T., KOHLHAMMER J.: Guiding feature subset selection with an interactive visualization. In *Visual Analytics Science and Technology (VAST), 2011 IEEE Conference on* (Oct 2011), pp. 111–120. 2
- [MEKKP11] MCDUFF D., EL KALIOUBY R., KASSAM K., PICARD R.: Acume: A new visualization tool for understanding facial expression and gesture data. In *Automatic Face & Gesture Recognition and Workshops (FG 2011), 2011 IEEE International Conference on* (2011), IEEE, pp. 591–596. 3
- [SG12] SHARMA N., GEDEON T.: Objective measures, sensors and computational techniques for stress recognition and classification: A survey. *Computer methods and programs in biomedicine* 108, 3 (2012), 1287–1301. 2
- [SH00] STOLTE C., HANRAHAN P.: Polaris: A system for query, analysis and visualization of multi-dimensional relational databases. In *Proceedings of the IEEE Symposium on Information Visualization 2000* (Washington, DC, USA, 2000), INFOVIS '00, IEEE Computer Society, pp. 5–. 2
- [SVKK14] SAPPELLI M., VERBERNE S., KOLDIJK S., KRAAIJ W.: Collecting a dataset of information behaviour in context. In *Proceedings of the 4th Workshop on Context-Awareness in Retrieval and Recommendation* (2014), ACM, pp. 26–29. 1
- [VZS09] VIZER L. M., ZHOU L., SEARS A.: Automated stress detection using keystroke and linguistic features: An exploratory study. *International Journal of Human-Computer Studies* 67, 10 (2009), 870–886. 2