Aalto University
School of Science
Bachelor's Programme in Science and Technology

visualization of human behaviour data for healthcare

Bachelor's Thesis

December 22, 2018

Simo Antikainen

Aalto University	ABSTRACT OF
School of Science	BACHELOR'S THESIS
Bachelor's Programme in Science and Technology	

Author:	Simo Antikainen
Title of thesis:	Visualization of human behaviour data for healthcare
Date:	Dec 22, 2018
Pages:	35
Major: **Code:**	Computer science SCI3027
Supervisor: **Instructor:**	Professor Eero Hyvönen Doctoral student Talayeh Aledavood, Staff Scientist Richard Darst (Department of Information and Computer Science)

The growing number of devices capable of collecting human behaviour data provide opportunities for its visualization in healthcare. This requires solutions to the three problems of breath of use, data complexity and the need for analytical validity. The thesis reviews solutions to these problems from the perspective of the patient and the clinician and offers instruction on their application.

Breath of use requires often the development of personalized visualization for the patient and support for visual analytics for the clinician. Complexity arising from the large quantities and qualitative nature of data can be managed through interactive visualization techniques and finding the right chart types. For temporal data these could include variants of line-charts, box-plots and heatmaps, whereas spatial data often requires the incorporation of time dimension, filtering and aggregation into map to support analysis. Analytical validity is can be attained through validating the data sources, pre-processing and the final visualization regarding accuracy and precision.

The practicality of the solutions was tested on a human behaviour data set using Plotly interactive visualization library for Python. It revealed the need for domain experience and multiple iteration cycles in producing appropriate visualizations for healthcare.

The thesis recommends evaluating the visualizations on how well they answer to the three problems. Interactive visualization techniques and appropriate chart types help in this provided they answer to the user needs.

Keywords:	Information visualization, healthcare, behaviour, digital phenotype, Python, Plotly
Language:	English

Aalto-yliopisto	KANDIDAATINTYÖN
Perustieteiden korkeakoulu	TIIVISTELMÄ
Teknistieteellinen kandidaattiohjelma	

Tekijä:	Simo Antikainen
Työn nimi:	Visualization of human behaviour data for healthcare
Päiväys:	22. joulukuuta 2018
Sivumäärä:	35
Pääaine:	Tietotekniikka
Koodi:	SCI3027
Vastuuopettaja:	Professori Eero Hyvönen
Työn ohjaaja(t):	Tohtorikoulutettava Talayeh Aledavood, Yliopistotutkija Richard Darst (Tietotekniikan laitos)

Ihmisen käyttäytymistietojen keräämiseen kykenevien laitteiden kasvava määrä tarjoaa mahdollisuuden kerätyn datan visualisoimiseksi osana terveydenhuoltoa. Tämä edellyttää ratkaisuja kolmeen ongelmaan: käyttökohteiden laajaan määrään, tietojen monimutkaisuuteen ja analyyttisen pätevyyden tarpeeseen. Opinnäytetyö tarkastelee näiden ongelmien ratkaisuja potilaan ja hoitohenkilökunnan näkökulmasta ja ohjeistaa niiden soveltamisesta.

Käytön laaja-alaisuuden hallitseminen edellyttää monesti personalisoitujen visualisointien kehittämistä potilaalle ja visuaalisen analytiikan tukemista hoitohenkilökunnalle. Tietojen suuresta määrästä ja laadullisista eroista johtuva monimutkaisuus edellyttää usein interaktiivisten visualisointitekniikoiden sekä sovelluskohteeseen sopivien diagrammien käyttöä. Ajallisen tiedon osalta näihin sisältyvät linja- ja laatikkokuviot, lämpökartat ja niiden muunnelmat, kun taas tilallinen tieto tarvitsee usein aikaulottuvuuden, suodatuksen ja aggregoinnin sisällyttäminen karttadiagrammiin analyysin tueksi. Analyyttinen validiteetti voidaan saavuttaa validoimalla tietolähteet, esikäsittely ja lopullinen tiedon esitys tarkkuuden suhteen.

Python-ohjelmointikielen plotly-visualisointi kirjastolla toteutettiin käytösdatan visualisointi projekti, jonka tarkoituksena oli testata ratkaisuja käytännössä. Tutkimus paljasti alakohtaisen asiantuntijuuden ja useiden iteraatiokertojen olevan hyödyllistä tiedon esityksen suunnittelussa terveydenhuollon tarpeisiin.

Työ suosittelee arvioimaan visualisoinnit sen suhteen, kuinka hyvin he vastaavat kolmeen ongelmaan. Interaktiiviset visualisointitekniikat ja asianmukaiset diagrammit auttavat tässä, mikäli ne on valittu käyttätarpeiden mukaan.

Avainsanat:	Informaation visualisointi, terveydenhuolto, käytös, digitaalinen fenotyyppi, Python, Plotly
Kieli:	Englanti

Contents

1 Introduction

The ubiquity of internet and mobile computing combined with the increase in internet of things devices have made it easy to collect large amounts of human behaviour data from a variety of sources (Marcengo and Rapp, 2013). On the cutting edge of behaviour data use is the emerging discipline of digital phenotyping which seeks to build continuous human phenotype profiles, by collecting different metrics from their personal digital devices (Onnela and Rauch, 2016). behaviour phenotype metrics such as mobility patterns, voice samples and social networks can then be used to diagnose and monitor mental health disorders, and to reveal possible links between expressed behaviour and underlying genetic causes (Onnela and Rauch, 2016). Phone based behaviour metrics have also been used to reveal correlations between user behaviour and physical symptoms of diseases such as common cold and influenza (Madan et al., 2012). The adoption of big data applications such as digital phenotyping in healthcare is inevitable result of the opportunities it provides (Murdoch and Detsky, 2013). One of these opportunities is the visualization of the data. Visualizations play a key role in behaviour data use in healthcare, because they act as a medium through which the clinician analyses the data. Patients gain value out of visualizations which increase their engagement to treatments (Gotz and Borland, 2016).

These opportunities give raise to a set of challenges related to visualization. Gotz notes (Gotz and Borland, 2016) that the healthcare sector there is a broad range of individuals interested in clinical data with different needs, from clinicians and patients to administrators and researchers, therefore visualizations should reflect these groups specific needs. Patients and clinicians needs alone give a raise to a large set of applications (Faisal et al., 2013), which include: treatment planning, examination of patients records and experiences, communication and decision making and life management and health monitoring. Also the complexity of the collected data and the strict safety requirements in clinical health application place another set of challenges for the visualizations (Gotz and Borland, 2016). Moreover the success of behaviour data applications hinges on the common users gathering of the data, for which engaging and non-inferential applications have to be developed (Rapp and Cena, 2014). Practical solution to these problems include the building of visual dashboards, which can display key metrics of patients to support clinician analysis needs (Stadler et al., 2016) and incorporating visual storytelling (Marcengo and Rapp, 2013) and interaction patterns (Wang et al., 2010) to the data graphics.

After gaining understanding of the domain specific challenges and solutions we can combine them with the long wealth of literature on the best practises of information visualization to produce the appropriate visualization. Edward Tufte in his famous book on quantitative information visualization (Tufte, 1986) outlined a set of principles of graphical excellence each quantitative data graphic should follow. In (Ware, 2004) the multiple ways visualization can make use of human perception to maximize the decision making power of visualization are realized. For most graph types the best practises are widely known (Heer et al., 2010). The created visualizations can then be analysed based on how effectively people gain insights using the visualizations. This can be achieved through benchmarking the effectiveness of achieving specific tasks and by analysing the number and quality of insights gained (North, 2006).

This thesis aims to answer: Which principles of information visualization are the most applicable to human behaviour visualization in healthcare, and what methods of visualization are appropriate for different behaviour data types, and what their design considerations are. The thesis aims to accomplish this though examining the ways patient and patient collected data has been visualized as part of healthcare and wellbeing, and how behaviour data fits into this application field. Through these applications the thesis motivates a set of visualization challenges and solutions to them. Solutions are specifically discussed to the problems of breath of use, data complexity and the need for analytic validity. Moreover the thesis examines specific visualization types as a solution to the rising challenges. Experimental part of the thesis aims to exemplify the challenges and their solutions through applying the learned to a real world human behaviour data-set collected as a part of digital phenotyping study for monitoring the state of a major depressive disorder.

Large scope of the thesis forces a set of restrictions on its scope. The focus in applications is placed on clinicians and patients who benefit the most of behaviour data visualization. The thesis does not explain mathematically the principles used for processing and displaying the behaviour data, nor does it discuss in detail the applications of behaviour visualization. The reviewed applications have extra focus on monitoring the mental state of the user, as the experimental part of the project involves visualization of behaviour features from a set of patients suffering from major depressive disorder. The specific visualization types discussed are limited to temporal and spatial nature of the human behaviour.

The thesis is structured in three parts: Part 2 discusses the application context of human behaviour data use in healthcare and wellbeing from the standpoint of the patient and the clinician and provides a summary of challenges and solutions to arising visualization problems. Part 3 details highly applicable visualization techniques for temporal and spatial human behaviour data, and part 4 details the application of the visualization methods to a real-world human behaviour data set.

2 Background

2.1 Self-tracking and health

There has been considerable interest in the development of mobile medical and health applications (mHealth) in the past decade (Boulos et al., 2014). One of the promising application fields of mHealth is using mobile devices to deliver behaviour interventions to help in the treatment of mental health disorders (Mohr et al., 2013) , chronic diseases and promoting healthier habits (Riley et al., 2011). The goal of behaviour change plays a key role in a larger field of personal informatics, which refers to a branch of science concerned with "collecting personally relevant information for the purposes of self-reflection and gaining self-knowledge" (Li et al., 2010). It is known mostly thought the quantified self movement (Marcengo and Rapp, 2013) which consist of people and companies interested with self-tracking and its applications. The goal of these applications is to offer a change for self-reflection and monitoring by visualizing the collected data. This increased self-knowledge is then supposed to offer changes for self-improvement and change. There exists many applications for the tracking of behaviour data, but it has had most success in the fields of fitness, mood and healthcare (Marcengo and Rapp, 2013). Healthcare sector has understood the importance of gathering the wealth of personal health information quantified-self applications can be used to gather coining the term quantified-patient (Appelboom et al., 2014).

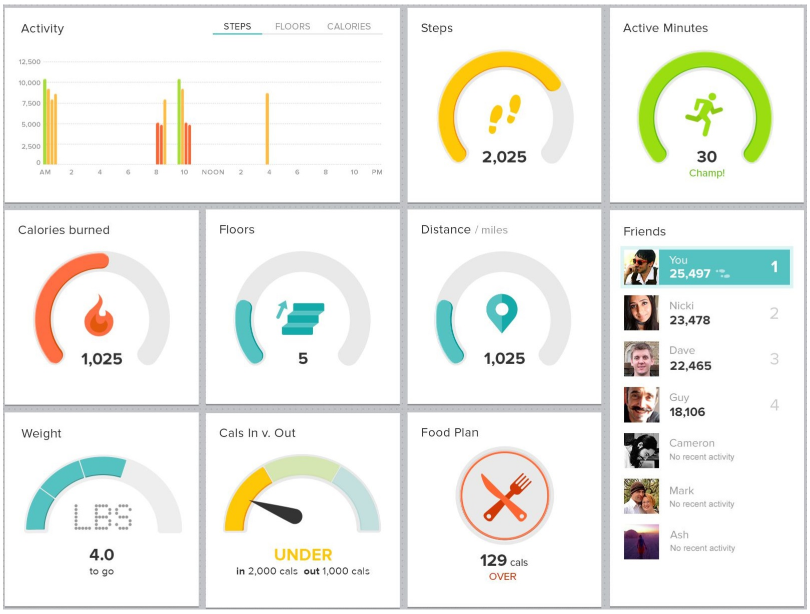

Figure 1: Fitbit dashboard of fitness and eating habit insights

Large part of the growing popularity of self-tracking can be placed on wearable self-tracking devices such as Fitbit and Jawbone which allow easy and non-inferring collection of activity measures (Crawford et al., 2015). Modern self-tracking devices are not limited to just tracking heart-rate or movement. Fibit charge3 offers tracking of sleep stages and offers insights on how to improve sleep hygiene. For women health it provides features such as period tracking and ovulation prediction. All of this data is then visualized through dashboards (Figure.1), which display goal progression and statistics tracked (Fitbit Inc, 2018). Other more healthcare specific include asthma inhaler sensor Propeller (Propeller Health, 2018), which tracks the information between inhaler use and location for the purpose of gaining insight on what triggers asthma attacks and for the prevention of asthma attacks. And Empatica smart watch, which detects relations between grand mal seizures and physiological measures (Empatica, 2018). Crawford (Crawford et al., 2015) notes that the value proposition of these devices hinges on the claim that the collected data is more accurate and reliable than subjective self-collected data, which of there is often no clinical proof. Moreover the users often have a limited degree of control over the data which might lead to the data being used against the users best interest for example in the form of denied insurance payments and disability claims (Crawford et al., 2015).

There is great potential in making use of self-tracking data other than in the form of direct behaviour change in healthcare as West (West et al., 2016) demonstrates. The data could be utilized to deliver evidence-based medicine by increasing the data available from the patient and to enhance the process of clinical decision making for the diagnosis and monitoring of the patient condition. Furthermore, the data could be used to communicate the symptoms of the disease to the patient. Choosing suitable application for this purpose can be an arduous process as review of fitness and health applications by Higgins (Higgins, 2016) shows. First, the applications often lack evidence-based care and clinician involvement and should be evaluated clinically on their effectiveness and safety before their use. Second, these applications should consider the wide range of different sensor that act as possible data sources, and the possible ecosystem of applications that can make use of the data to provide a more complete picture of the user health. Self-tracking applications, such as Exist (Hello Code, 2018) offer leads on how self-tracking benefits could be maximized. It integrates a wide range of information ranging from fitness trackers to social media use statistics, and provides visual insights based on the data to improve its users mood (Figure.2), physical activity and sleep quality. Still before an application such as Exist should be used as part of healthcare it should be evaluated according to factors outlined by Higgins, as it has not been validated to be used in healthcare.

The central issue in personal informatics applications is how to keep the user engagement up. This is challenged by the large non-expert audiences which are not well invested in self-tracking and do not understand what it can offer for them, and on the other side by

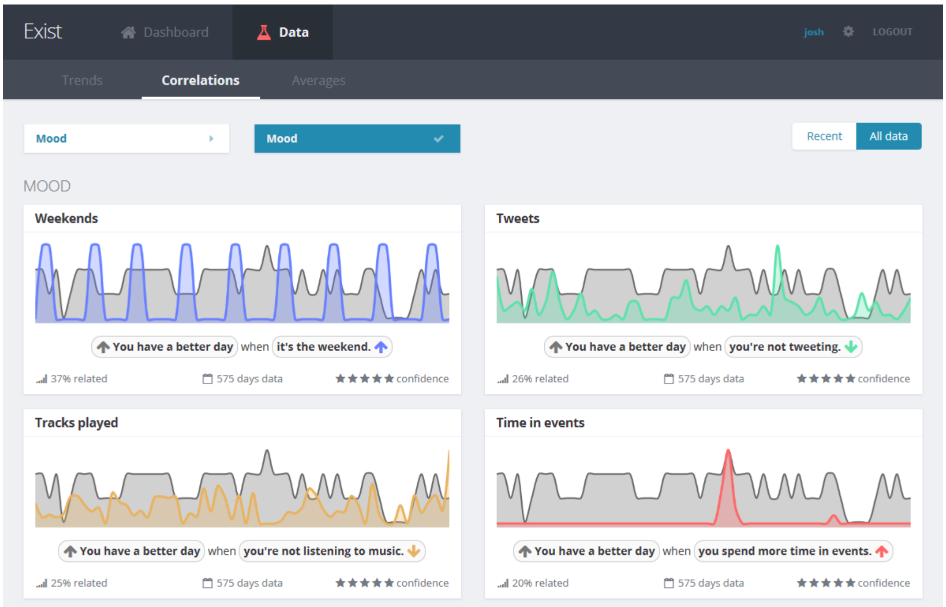

Figure 2: Exist self-tracking application mood insight dashboard (Hello Code, 2018).

the immature ecosystem which does not match user needs (Rapp and Cena, 2014). One of the ways to increase user engagement is to systemically research which type of data and what kind of visualization provide the valuable insights to act upon (Epstein et al., 2014) and (Choe et al., 2015). In a study of how non-expert audiences use personal informatics applications (Rapp and Cena, 2016) several ways where found to improve visualizations for non-expert users. The data should reflect the user in a personal manner. This can be achieved with avatars which make it easy for the user to form an emotional bond to the data leading to increased motivation for data collection and self-reflection. The applications should also provide custom information in the form of reports, goals and suggestion for the user to inform him of the benefits of the self-tracking. This information should leverage the power of storytelling to engage the non-expert user. The applications should allow users to relive the past, because memories add meaning to data. This can be achieved by letting the user re-experience the past through spatial visualizations (Rapp and Cena, 2016). This problem could be integral in the design of the patient facing side of any eHealth applications, which relies on regular user engagement and data gathering to deliver the service.

Research by (Cuttone et al., 2014) proposed a set of more general design heuristics to facilitate the goal of self-reflection and change in personal informatics applications. The data should be interpretable at a glance since users often want answers with little effort and time. This can be achieved with the help of dashboards and summaries which transform complex data sets into simple sets of readings from which it is easy to interpret the most essential information. In contrast to simplicity, the visualizations should also have the option to increase in complexity depending on the user needs and domain experience. The visualizations should allow the exploration of global trends and periodic patterns in time series data. Global trends can be explored with line plots, calendar and spiral heatmaps provide a natural way to present periodic data. The visualization should enable the discovery of trends in multiple data streams by supporting multivariate analysis of the data. This can be achieved by using appropriate visualization techniques such as scatterplots, correlograms and small multiples and by overlaying line plots. Finally, the visualizations should support the iterative process in exploratory data analysis where questions and answers found in the data lead to further exploration of the data. The iterative process can be supported by the use of interactive techniques such as zoom and filter, details on demand and history and annotations (Cuttone et al., 2014). In practice these heuristics would have to take into consideration the mobile use context of self-tracking. Mobile devices pose a set of technical problems for the visualization work (Marcengo and Rapp, 2013). The small size of mobile device displays means that we must find methods to maximize the amount of visualized information, and the small processing power of the devices places limits on the amount of data processed. The mobile use context of the applications means that visualizations will be accessed in a variety of different lighting conditions challenging perception and that the human attention span will not often be fully utilized by the visualizations.

2.2 Visual analytics in healthcare

The use of behaviour data visualization in clinical healthcare falls largely under the larger umbrella term of visual analytics, which seeks to enhance analytical reasoning through the use of advanced interactive visual interfaces (Thomas and Cook, 2006). The field combines techniques from data mining, machine-learning, statistics, human computer interaction and human cognition to enhance the reasoning power of the visualizations. The produced visualizations could then help clinicians monitor patient progression, by integrating multiple different data types together, and help patients to compare their health measurements against other patients. Researchers could benefit from the visual tools, which let them derive insights from aggregate patient data (Caban and Gotz, 2015). The main application context of visual analytics for behaviour visualization is through different clinical decision support systems, which support clinicians in clinical decision making by turning collected health data and knowledge into helpful insights (Musen et al., 2014). The visualizations itself can be as simple as showing the results of a patient filled mood questionnaire in an telemedicine application to a doctor, or complicated as helping to infer relationship between medicine received and symptoms of diseases and disorders the patient suffers from (Plaisant et al., 2003).

One of the major application fields of visual analytics in healthcare is data-mining and visualizing the data collected in electronic healthcare record systems (EHRs) for the purposes of clinical decision making and research (Rind et al., 2013). Rind (Rind et al., 2013) notes that EHRs are a valuable resource for clinical decision making, because they centralize the necessary patients medical history and current symptoms and treatment into a easy access location. These systems possibly will provide amble opportunities for behaviour data visualization as the incorporation of social and behaviour characteristics of the patient in the systems is seen increasingly important (Estabrooks et al., 2012) and (Adler and Stead, 2015). Of these EHR visualization applications one of the most researched is Lifelines and its successors Lifelines 2 (West et al., 2015). Lifelines purpose was to produce temporal visualization of a selected persons whole medical history for the purpose of inferring relationships (Figure. 3) between different parts of medical records (Plaisant et al., 2003). Lifelines 2 refined this concept to work population level by allowing the clinician to find patients based on shared temporal events (Wang et al., 2008). Other EHR visualization systems such as VISITORS (Klimov and Shahar, 2005) allow the cohort level analysis of patient populations based on shared temporal events. Although the approach of inferring relationships between temporal events such as the symptoms and the medication received is well understood, it is not often the case with behaviour data. For example there is no direct link between how depressed a patient is and how much he moves or uses social media (Pratap et al., 2018), which means we have to be careful at not making misleading visualizations.

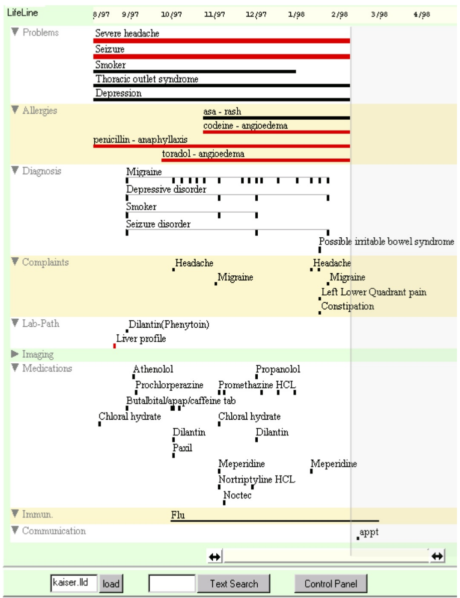

Figure 3: Lifelines temporal visualization of patient history (Plaisant et al., 1996).

One of the most promising sources for visual analytics of behaviour data could rise as a part of applying digital phenotyping to deliver telemedicine. It leverages a wide range of behaviour data collected to build human behaviour profiles. The difference between digital phenotyping and previous approaches such as Ecological momentary assessment (EMA) in symptom and behaviour tracking is that it can make use of passively collected data from mobile devices in addition to active data, whereas previously the data was collected mainly actively in the form of surveys (Insel, 2017). This wide range of captured behaviour would lead to a large range of visualization opportunities, which could be used to deliver better personalized and evidence-based healthcare. Research on digital phenotyping offers opportunities to look at what the possible metrics of the visualization are. In mental health alone digital phenotyping has revealed correlations between mobility features derived from global positioning system (GPS) data and depressive symptom severity (Saeb et al., 2015), and prediction of bipolar states and state change through the combined use of mobility and social interactivity features such as the total number of phone calls and the average speaking length (Grünerbl et al., 2015).

Interactive visualization in healthcare has two other major domain specific challenges in addition to the problems caused by the breath of use, the problem of data complexity and the need for analytical validity (Gotz and Borland, 2016). The large variety of collected data means that the visualizations must allow the easy navigation of multidimensional data, and that they must allow the easy handling and display of multiple data types and their relationships. The temporal nature of the data means that visualizations must display interesting temporal patterns over time. To help in this, the visualizations could use summarizing and prioritization techniques such as temporal aggregation. Data complexity is also a problem in the form of data quality and heterogeneity. Problems arise when we are required to discern between events that have not between recorded versus which have not happened. Different data sources also have differences in bias and precision especially in the case of data collected from mobile devices. Most of these concerns were addresses in a study discussing self-collected data application to healthcare by West (West et al., 2016), which also raised the question of how data literate the clinicians itself are to analyse the data, as they often lack the training to address the missing standards for data preparation in self-collected data, and lack the tools for its analysis. Clinicians are also not often familiar with the wide variety data collected.

With medical data analytical validity is of utmost importance (Gotz and Borland, 2016). They have to satisfy a set of both statistical and visualization requirements to satisfy the high safety standards required. Visualizations must be careful not to show incorrect causality between variables and representative samples visualized must have minimal selection bias. Moreover interactive visualizations should offer a possibility to detect selection bias and filter it out. West (West et al., 2016) offers possible solutions for these problems. First the devices, which collect the data are checked for data quality as they often are not approved for medical use. With data requiring active user collection it should be checked for relevance and completeness, as the user often logs unimportant data for the clinician or simply does not record it. Second, the visualizations itself should follow the standards of data presentation in healthcare, such as the standard way to represent heartbeat that are often lacking in self-tracking applications. Moreover, the visualization could be audited on how intuitive they are to use as they are interpreted through many different audiences.

A more relevant set of seven challenges for the use of interactive visualization use in clinical healthcare was outlined by (Shneiderman et al., 2013): Visualizations should offer clinicians timely information in the right format. The diverse set of data needs a set of specialized visual tools. The visual tools should facilitate team decision-making. Visualizations should help to characterize and understand similarities. Visualizations should help to understand cause and effect relationships and comparative-effectiveness. Visualizations should predict the future, risk and uncertainty. Visualizations should be evaluated so they match the safety criteria in place(Shneiderman et al., 2013).

To increase the decision-making power of these healthcare visualizations the information seeking in the data should be as efficient as possible. One of the ways to increase information acquisition and the amount of insights from healthcare data is to add different types of interaction patterns into the visualizations (Wang et al., 2010). These interaction patterns are often concerned with helping to visualize the high volume and variety temporal data that can arise as part of healthcare applications (Du et al., 2017), and is often typical of human behaviour data. One of the interaction patterns which has been used as solution to the problem of visualizing this type of data in healthcare (Bade et al., 2004) and personal informatics (Cuttone et al., 2014) is the Shneiderman information-seeking mantra. Shneiderman mantra (Shneiderman, 1996) of overview first, zoom and filter, then details on detail is a set of generalisable interactive visualization rules, which are meant to act as the first step in maximizing the information seeking efficiency.

For the organization of the patient data into relevant and easily actionable visualizations the use of clinical dashboards has been researched in the past decade (Dowding et al., 2015). These dashboards visualize the necessary key performance metrics without the need for further visual analysis of the data. Stadler (Stadler et al., 2016) demonstrated that the use of clinical dashboards lead to improved efficiency of clinical analysis leading to significant time savings and enabled wide range of end-user to access data to derive insights from. The use of behaviour data in these clinical dashboards could adapt the style seen in self-tracking dashboards such as (Figure.1) and (Figure.2) to a medical context. Still dashboards are not a catch all solution for the visualization in healthcare as Dowding explains (Dowding et al., 2015) the characteristics of dashboards which improve decision making are not understood well enough. Their use might also lead to tunnel vision where the measured data is given over-importance at the expense of unmeasured data.

3 Methods of behaviour data visualization

In the last part we saw how human behaviour data appears as high-density temporal and spatial data. The following review of visualizations for these types of data is ordered according to the task by data type taxonomy presented in (Shneiderman, 1996), where the ordering of the visualizations is based on the type of data they are representing. Focus is given on techniques that help the viewer manage the large density of the data. This density is often not a problem as Tufte (Tufte, 1986, p.161-162) explains the human eye is able to make a large number of distinction in a small area, and therefore the visualization should be packed densely with data. Furthermore, multiple variants of the visualization are introduced to cope with the large amount of use cases, and common pitfalls, which endanger analytic validity are discussed.

3.1 Temporal data

3.1.1 Line-chart

Line charts main purpose is to display trends over time (MSKTC, 2018). Its x-axis consists of time units distributed in even intervals, and the y-axis represent the range of variables. Time series is then plotted in the corresponding points of the chart and the closest individual variables are connected across the time-axis using lines. In design of line charts the y-axis should be always extend to zero or some other expected baseline and the intervals of time in x-axis should always be constant. If multiple line chart is used the different lines should always be meaningfully related to each to avoid the creating correlations and causality which does not exist. There exist many alternatives to the traditional line chart including index charts and stacked graphs (Heer et al., 2010). Index chart allows us to compare the relative changes of multiple variables against a baseline, for example the percentage changes in stock prices against some date seen in (Figure: 4). Stacked graphs aggregate multiple line charts values together but can only be used when the variables have positive values and the units of the variables are the same. Readability of individual values can be improved by using bar chart and stacked bar chart in place of their line equivalences.

Plotting multiple lines in the same line chart can make it difficult to infer differences between the lines. To manage the density of data graphics in line charts and many other visualizations small multiples (Tufte, 1986, p.170-75) can be used to represent different variables in multiple small charts instead of plotting them into the same chart. Possible large range of y-axis values can be managed using horizon graphs (Heer et al., 2009) see (Figure:5), which mirror negative values into the positive range, bands the y-axis and superimposes the bands on top of each other. This drastically cuts the y-axis size at the

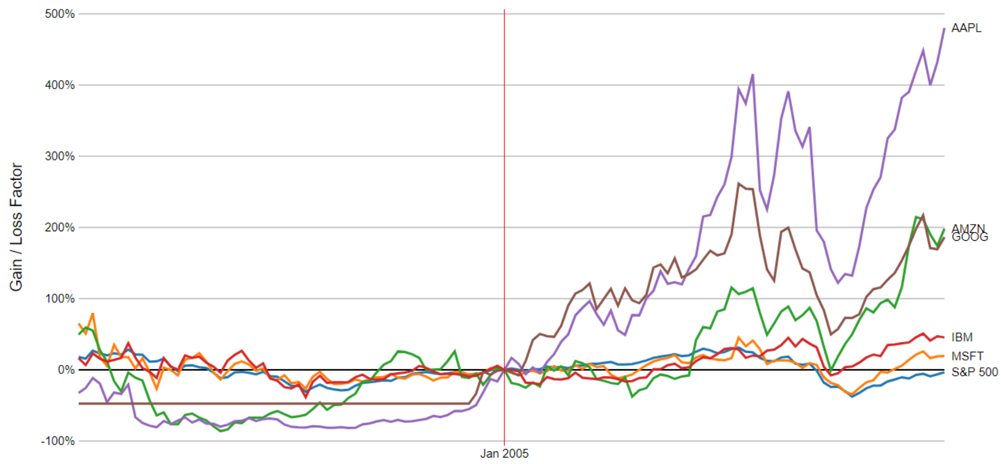

Figure 4: Index chart of stock prices (Group, 2018b).

cost of making the value reading from chart slightly slower and accurate.

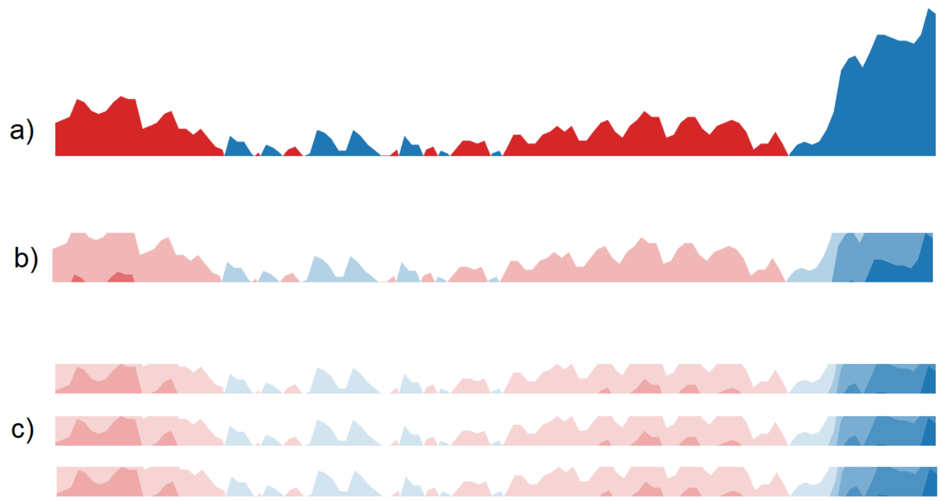

Figure 5: a) Mirrored negative values in red b) The value ranges are superimposed on top each other to form horizon graphs c) Small multiples of horizon graphs. Adapted from (Group, 2018a).

3.1.2 Box-plot

Box-plot (Figure:6) is used to visualize summary statistics on categories of variables. The traditional box plot uses the median, upper and lower quartile and the highest and the lowest value of the data to visualize the level, spread and symmetry within the category (Williamson et al., 1989). The categories of variables visualized can be belong to time-intervals or individual days making the box-plot a temporal visualization technique. The traditional box-plot has been developed extensively since its conception with modification such as the variable width box-plot, which uses the width of the box to represent the size of category and the notched box-plot, which uses a notch in the box to visualize the confidence interval around the median (Potter et al., 2006). Further developments include the addition of probability density function of the distribution using information using kernel density estimation in the violin-plot (Hintze and Nelson, 1998). Box and violin-plots miss vital information about the number of the variables in the plot to see if the visualization is statistically meaningful. Individual variables can be displaying on the side of the box-plot using a strip chart or the chart can be place inside the violin-plot creating a bean-plot (Kampstra and others, 2008).

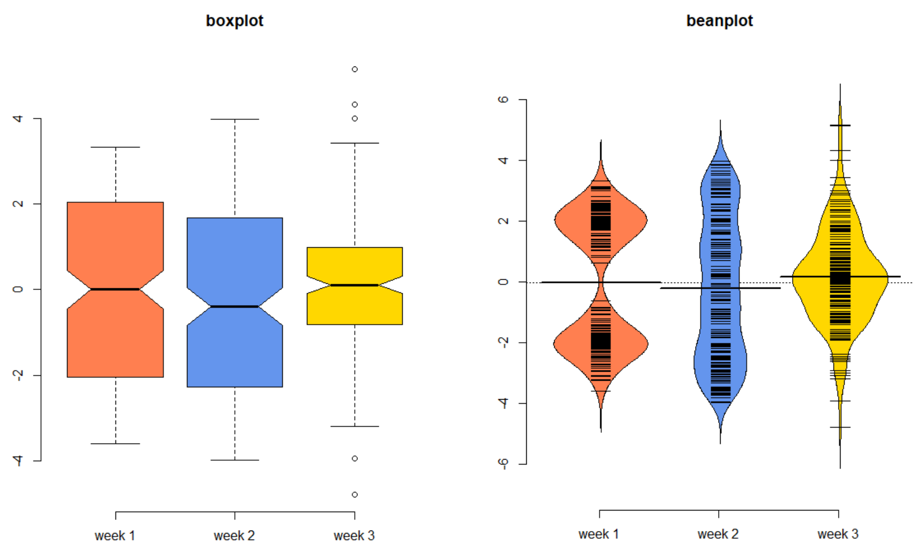

Figure 6: Box-plot comparison to bean-plot. Box-plot misses information about the weekly distribution spread, which is contained in bean-plot strip chart black stripes and the probability density function in color around the strip chart

3.1.3 heatmap

Heatmaps (Figure:7) are used to represent two dimensional columns of data by shades of colours. The density of the displayed information and the intuitiveness of the colour mapping makes it proper for displaying high-throughput data (Gehlenborg and Wong, 2012), such as temporal data arising from behaviour tracking. Two principles (Gehlenborg and Wong, 2012) are vital in heatmap design: choosing how to order the rows and columns relative to each other, and the choice of colours used. The problem of ordering can be solved using clustering. The colour mapping used should use a divergent colour gradient if we want to highlight values at both extremes and grayscale should be used when

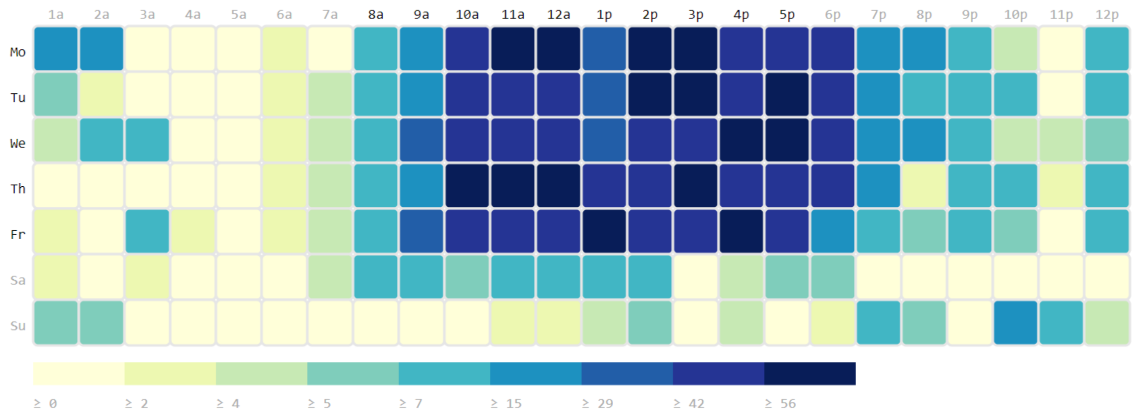

Figure 7: Heatmap of weekly activity ordered by day and hour (May, 2018).

highlighting one extreme. When one of the axes of the data has an inherent ordering, which is the case with time-series data a more proper visualization could be the parallel coordinates plot. It visualizes multivariate data points in parallel axis columns and the points are connected with each other (Heer et al., 2010).

3.2 Spatial data

3.2.1 Spatial heatmap

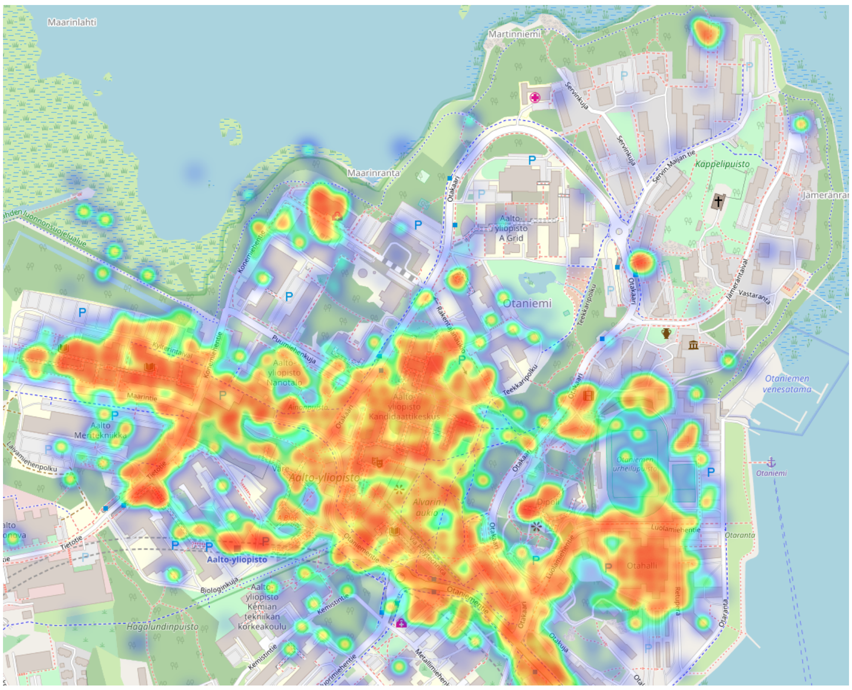

Figure 8: Heatmap of personal movement around Otaniemi area.

Spatial heatmaps (Bojko, 2009) represent the density of two-dimensional observations using colours (Figure:8). Bojko explains that they are a natural choice for spatial data for two reasons. First, the human viewers easily associate the cool to hot colorscale with density. Second, it aggregates large quantities of spatial data into a form from which it is simple to discern the shape of the distribution. This adaptability has lead spatial heatmaps to have a wide range of applications in visualizing human behaviour data from movement tracking (Giannotti et al., 2011) to eye gaze (Duchowski et al., 2012) and cursor movement (Huang et al., 2011). Although these maps help in communicating the data Bojko (Bojko, 2009) notes that density alone cannot help explain or analyse the data in most use cases. He notes that spatial heatmaps easily lead to drawing false conclusion if they are not compared to other heatmaps. Moreover, even if comparative heatmaps are generated the visualization type does not lend to itself for meaningful comparison between the different charts if data analysis is not performed first.

3.2.2 Spatio-temporal data and spatial trajectories

The need for incorporating temporal dimension into spatial data was acknowledged by Hägerstraand (Hägerstraand, 1970) as time context is integral to human actions. His thoughts lead to creation of a field of time geography (Sui, 2012), which researches human activities through space and time. Time geography's benefits to place-based health research are illustrated by Rainham (Rainham et al., 2010). Through tracking individuals' activities through a space-time path, it is possible to explore the space-time patterns and their possible relation to health or demographic characteristics. Possible opportunities for visualization could arise within finding correlations between location and weight-related behaviours (Zenk et al., 2011) and (Rainham et al., 2010), and location and mental health (Wheaton and Clarke, 2003) and (Vallée et al., 2011). Self-tracking applications, such as Sports Tracker (Sports Tracking Technologies, 2018) allow the user to produce charts of their activities ordered by space and time and share them to the wider community for the purpose of motivation and gaining location based insights.

Spatio-temporal data calls for many different approaches for visualization. For example in the case of spatial trajectories Dykes (Dykes and Mountain, 2003) made use of a tool which plotted movement behaviour data points on a map coloured by time. His tool allows the selection of point by time-period and provides the possibility of identifying space-time patterns, which are characterized by the movement type, such as driving or walking, and time period. The tool provides opportunities to explore how the user's relation to environment has changed through time and how environmental constraints affect movement behaviour. Often these trajectories become cluttered by the large volume of data points and there is a need for generalization of movement for which the movement data is often aggregated together along the path and direction (Adrienko and Adrienko, 2011).

4 Case study: Visualizing metrics of behaviour

4.1 Introduction

The case study was done using behaviour data collected from a pilot study between Aalto University's complex system group and Helsinki University Central Hospital of using digital phenotyping to monitor the state of major depressive disorder, which utilized the Non-Intrusive Individual Monitoring Architecture (Niima) platform (Aledavood et al., 2017) to collect and distribute the patient data. The pilot study involved 14 patients and 23 healthy controls. Data collected involved active data in the form of starting and exit surveys, and daily questionnaires for two weeks. Passive data up came from Aware framework phone measures, Phillips Actiwatch and Murata bed sensors. From these the Aware framework phone data was recorded up to one year and the other two were recorded for two weeks. The goal of the study was to visualize key behaviour metrics of mood, sleep, sociability and physical activity that the clinician knows how to interpret, and to demonstrate the possibilities for visualization that might arise in systems that integrate large quantities of health data, such as EHR-systems. The visualizations itself drew ideas from simplistic self-tracking and clinical dashboard design, with emphasis on the display of large quantities of data in a small space. Idea was to make the visualizations accessible to both the clinician and the patient, while following possible clinical guidelines related to them. The visualizations made use of niimpy behaviour data gathering package (CxAalto, 2018), which was used to query and pre-process the collected data.

4.2 Methods

For plotting the charts Plotly (Plotly, 2018) was chosen, because it includes interactive visualization patterns such as the aforementioned Shneiderman information-seeking mantra and the ability to build dashboards. Its style defaults seemed sensible and easily modifiable. Plotly also included a slew of other interactive visualization techniques such as data transformations and custom controls such as the possibility of changing style and the type of the visualization. Techniques other than the Shneiderman mantra were not used as incorporating them can require considerable amount of custom code for each plot.

The visualization itself were split into four categories. For mood the heatmap (Figure. 9) visualizes daily mood values on a Likert scale from 0 to 6 (Komulainen et al., 2014). The motivation was to display the questionnaire values and show their progression to the clinician. For sleep the bed sensor status (Figure.10) and the Actiwatch actigraph

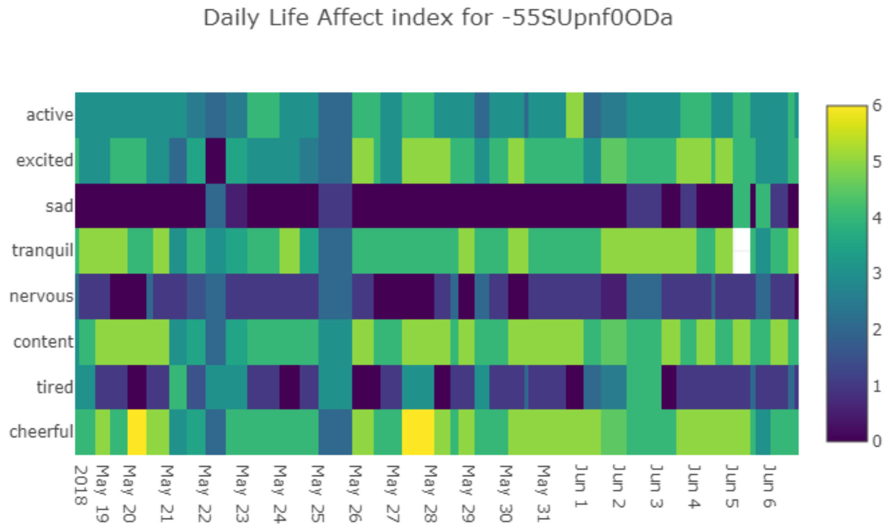

Figure 9: Daily life Affect index

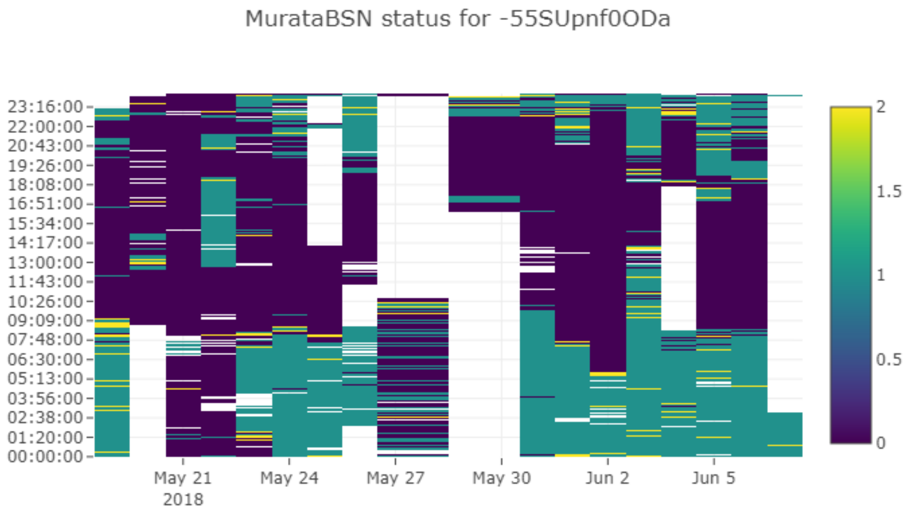

Figure 10: MurataBSN status

(Figure.11) were visualized using a temporal heatmap. Visualization of sleep status is valuable as sleep disorders are a major symptom of depression and insomnia risk factor in the development and recurrence of depression (Nutt et al., 2008). Moreover, depressed people often suffer from abnormal sleep patterns, such as delayed onset of sleep and difficulties with maintaining sleep. These abnormalities can be visualized using actigraphy, which measures how much the person is moving (Morgenthaler et al., 2007). Sleep status visualization drew idea on how the sleep status is visualized in self-tracking application, such as the Fitbit Charge line of trackers and traditional actigraphy. Determining stages of sleep was not tried as this would have required considerable amount of pre-processing and it is know that actigraphy accuracy suffers when the user suffers from sleeping disorders (Sadeh and Acebo, 2002).

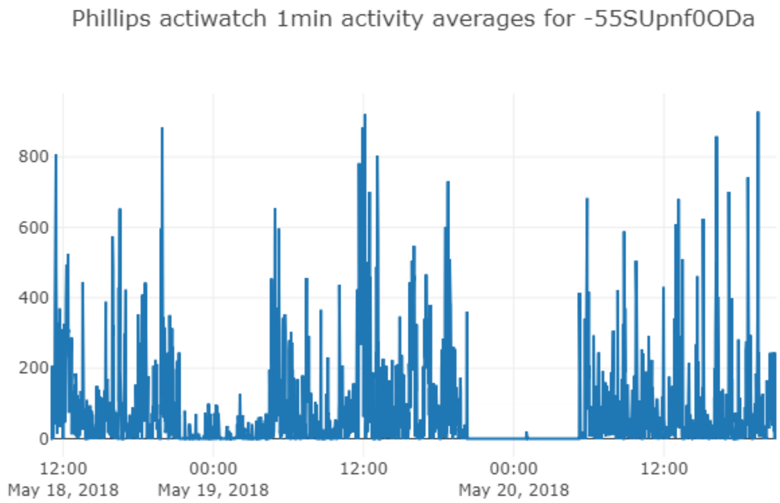

Figure 11: Actiwatch 1 minute activity averages

Actigraphy can also be used to measure the daily activity (Figure. 14), which is useful as there is evidence of physical exercise being valuable for the prevention and treatment of depression (Ströhle, 2008). Motivation for the visualization was drawn from Fitbit dashboard visualization of total activity seen in (Figure. 1). For sociability the amount of daily phone calls and mobile phone usage was visualized. Plots for mobile phone application usage can be seen in (Figure. 12) and (Figure. 13). These are valuable as research suggest links between the levels of depression and sociability, and the possibility of measuring sociability from phone based measures, such as phone calls and usage (Saeb et al., 2015). High levels of phone usage have also been linked to be a risk factor in the development of depression and sleep disturbances (Thomée et al., 2011). These visualizations likewise tried to visualize sociability in a daily summary form.

4.3 Results

The results of the visualizations were mixed, based on the visualizers opinions. From the mood visualization (Figure.9) it was easy to see the different values for the mood measures and compare their progression. The sleep sensor status (Figure.10) combined with the actigraph (Figure.11) of the patient made it easy to discern if the patient was

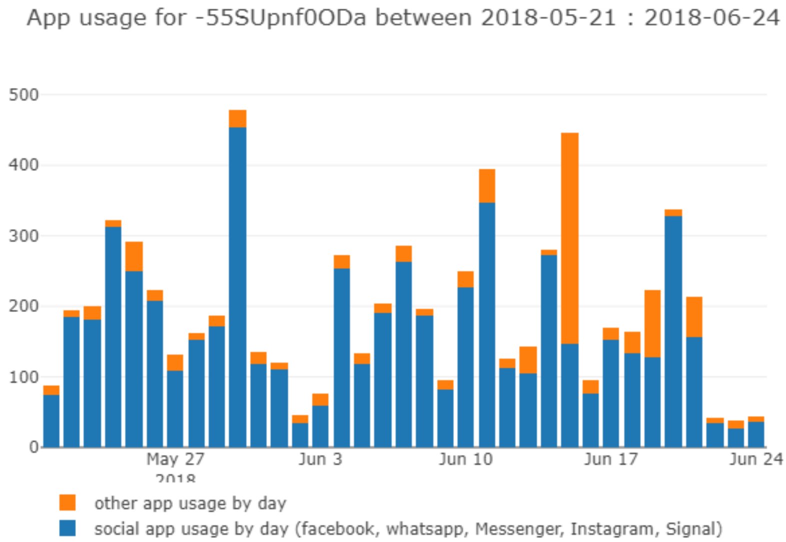

Figure 12: Application usage bar chart

sleeping on the bed and gauge the sleep duration. Assessment of sleep quality from the plot was not attempted, but theory suggest it is clinically proven based on the actigraph data. With the measurements of activity, the daily and momentary movement were crudely gauged as the visualization made no attempt to discern types of activity such as physical exercise from other activities. With sociability (Figure.12 &13) the raw amount of application use including social media, and the amount of calls were visualized daily and the descriptive statistics of sociability were visualized weekly. As the amount of phone usage was gauged from the number of systems calls and not from the minutes application was used the accuracy of the measurement is questionable. The phone call visualizations did not consider the number of missed calls or the patient caller network characteristics.

Assessing the visualization better especially towards how they handle the three big problems of visualization in healthcare would require access to user feedback from possible groups taking advantage of the visualization. Furthermore, the time constraints prevent the development of meaningful evaluation criteria. Still we can speculate how this could be done. User feedback could be used to gather information of qualitative and quantitative insights gained from the visualizations. For complexity the amount actual data in visualization could be measured (Tufte, 1986, p.91), and for analytical validity criteria could be developed, which could include measures of graphical integrity, such as the charts do not distort the data (Tufte, 1986, p.53-87).

Despite the mixed results the case study suggest that the data gathered as part of a digital phenotyping study can be utilized more broadly within the context of delivering

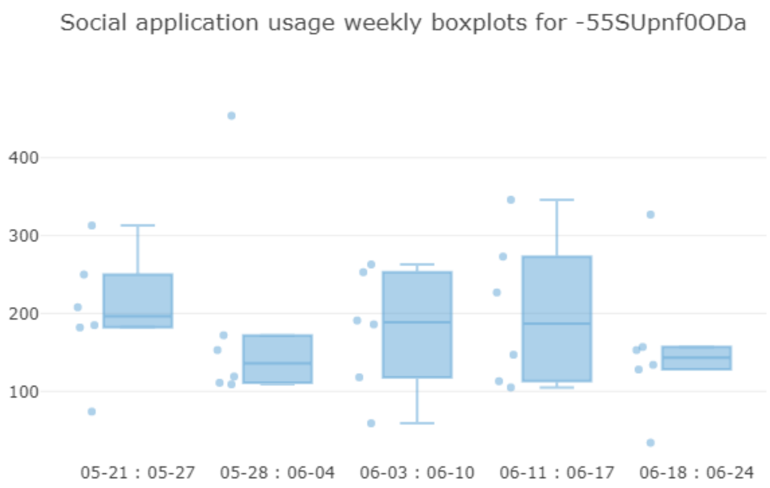

Figure 13: Social application usage boxplot

telemedicine, such as visualizing the results of a patient filled mood questionnaire or the measures of sleep quality to the clinician. The shortcomings in the visualization work were mainly related to the single iteration cycle of the visualization work and to the inexperience of the experimenter, which suggest a more experienced person might be capable off improving these results. This inexperience also led to concerns of analytical validity of the visualizations as the choices made in pre-processing the data did not follow analytical guidelines. As the location data was anonymized the examination of spatial and spatio-temporal characteristics of the patient data and its possible links to depression was partial. This meant the spatial information visualized consisted of indirect information derived from temporal measures, such as is the patient at home or at work or the duration of physical exercise.

The use of Plotly satisfied its expectations. It allowed its user to zoom in the data and inspect each individual data point value. The default styles for different chart types were sensible and easily modifiable. The interactive visualization worked fluidly despite the large amount of data points in the individual plots. Building dashboard was unfortunately not possible in the offline mode of Plotly, which was remedied by building a small website inside jupyter Python notebook cell. It also had problems with handling inconsistent date intervals, falsely extending the time-axis of the plot into the next day. Setting axis time-intervals was likewise difficult. These problems handling time-interval might compromise the analytical validity of the visualization.

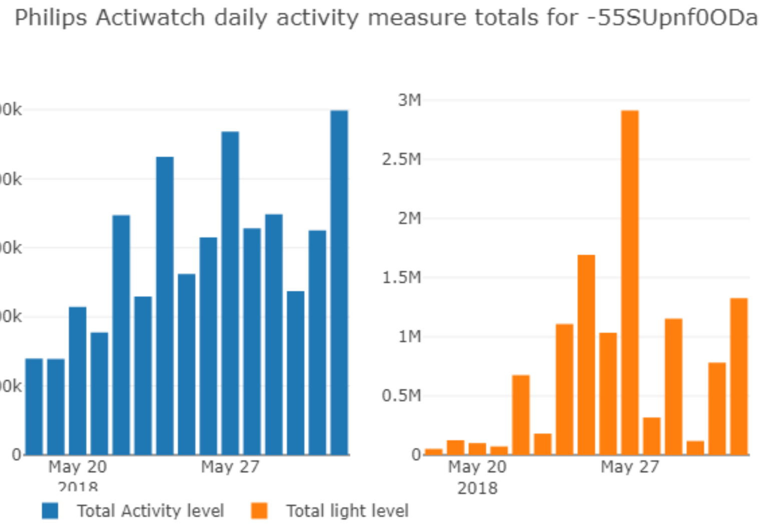

Figure 14: Philips Actiwatch daily activity measure totals bar charts

5 Conclusions

Behaviour visualization provides a large set of benefits in healthcare for both the patient and the clinical personnel responsible for the treatment of the patient. Patients benefit most directly from digital wellbeing and eHealth applications, which support goal formation and behaviour change. Telemedicine applications which are used to provide physical and physiological diagnosis and monitoring could benefit from better treatment adherence and outcomes by incorporating behaviour visualization and tracking to the patient facing side of the application. Clinical personnel benefits directly from the visual analysis of behaviour data, which is stored in EHR systems and is increasingly created by all the eHealth applications. The visual analysis can come in the form of discovering patterns in the collected behaviour data or simply through checking the values of performance indicator of patient health.

This all comes with a set of challenges most of who are inherent to visualizations use in healthcare. Matching the needs of users requires large amount of flexibility for the visualizations or producing completely different visualizations for the different users. Traditionally the patient side of visualization has favoured simple and personalized dashboard visualization, whereas the clinical developments have favoured the use of visualization as a tool for exploratory data analysis of the patient data. On the other side of things clinician benefit of simple dashboard visualizations that display the key metrics of diagnosis and monitoring of the patient as this reduces the time the clinician must work with the individual patient. Conversely a well-informed patient might benefit from deeper exploratory visualization of his data if it does not lead to false conclusion and supports better health outcomes. There also exist possibilities for the visualization to act as a conversation tool as a part of clinician patient communication.

The problem of data complexity and density can be managed through careful use of interactive visualization techniques, such as the Schneiderman information seeking mantra which allow the user to view the data at different granularity levels from simple overviews to viewing individual variables. Using appropriate visualization we can increase the amount of insights gained from the huge quantities of temporal and spatial behaviour data. Line charts, heatmaps and their variations are effective at displaying temporal trends. Box-plot excels at providing a large set of descriptive statistics for temporal data intervals, which can be plotted next to each other revealing possible trends in the statistics. For spatial data spatial heatmaps visualize the density of the behaviour but are often misleading, and by itself often do not support deeper analysis of the data. Human movement can be visualized using spatial trajectories. Often the addition of temporal dimension into spatial data can help in the discovery of new insights. Aggregations and filtering help to combat the large number of spatial data points.

The problem of analytical validity in visualization can be remedied through auditing the whole process from the data collection to the final visualization itself. First the devices, which collect the data are checked for data quality as they often are not approved for medical use. With data requiring active user collection it should be checked for relevance and completeness, as the user often logs unimportant data for the clinician or simply does not record it. Second, the visualizations should follow the standards of data presentation in healthcare what are often lacking in self-tracking applications. Moreover, the visualization could be audited on how intuitive they are to use as they are interpreted by many different audiences as part of clinical practice. Appropriate use of interactive visualization techniques and visualizations contributes to solving to this problem by helping to infer the right values.

The case study displayed in practice the visualization through which patient and clinicians might access this behaviour data as part of healthcare, but the visualization itself were not evaluated by their possible users on how well they match their user group's needs, handle complexity and analytical validity. Still not making use of the visualization opportunities that rise up in behaviour tracking in healthcare seems wasteful as the numerous applications have demonstrated. Furthermore, the failures demonstrated how important it is to have domain experience when building the appropriate visualizations and the need for multiple iteration cycles.

Producing an appropriate visualization is often a highly individual process as this thesis demonstrated. In the whole cycle of gathering user needs, gathering, reprocessing and visualizing the data, and collecting feedback for future iterations there exist countless opportunities for future research. In the front of user needs we did not directly cover how researchers benefit from behaviour data visualization. The visualization of the data was covered mainly from the perspective of individual level and focused on personal health leaving out population level quantitative analysis and its contribution to healthcare. Inferring possible relationships in the data was limited to relationships between variables and time neglecting the visualization of intervariable relationships. The network nature of human behaviour and the accompanying network visualizations were also not covered.

One promising research direction, which was not widely covered in this thesis is how behaviour data visualization can be used in the context of physical patient clinician encounter. The collected behaviour data could be used to converse about the disorder or the disease with the patient. It is also careless to base treatment and diagnosis blindly to data gathered without face to face encounters as the data can be often unreliable because of its nature or simply through carelessness in how the patient gathers it. Another direction comes from the use of spatial data for healthcare, although it is widely used in the self-tracking and research applications it seems not to be utilized by the clinician in understanding his patients' relationships towards his environment.

References

Nancy E Adler and William W Stead. Patients in Context — EHR Capture of Social and Behavioral Determinants of Health. *New England Journal of Medicine*, 372(8): 698–701, 2 2015. ISSN 0028-4793. doi: 10.1056/NEJMp1413945. URL `https://doi.org/10.1056/NEJMp1413945`.

N Adrienko and G Adrienko. Spatial Generalization and Aggregation of Massive Movement Data. *IEEE Transactions on Visualization and Computer Graphics*, 17 (2):205–219, 2011. ISSN 1077-2626. doi: 10.1109/TVCG.2010.44.

Talayeh Aledavood, Ana Maria Triana Hoyos, Tuomas Alakörkkö, Kimmo Kaski, Jari Saramäki, Erkki Isometsä and Richard K Darst. Data collection for mental health studies through digital platforms: requirements and design of a prototype. *JMIR research protocols*, 6(6), 2017.

Geoff Appelboom, Melissa LoPresti, Jean-Yves Reginster, E Sander Connolly and Emmanuel P L Dumont. The quantified patient: a patient participatory culture. *Current Medical Research and Opinion*, 30(12):2585–2587, 12 2014. ISSN 0300-7995. doi: 10.1185/03007995.2014.954032. URL `https://doi.org/10.1185/03007995.2014.954032`.

Ragnar Bade, Stefan Schlechtweg and Silvia Miksch. Connecting Time-oriented Data and Information to a Coherent Interactive Visualization. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '04, pages 105–112, New York, NY, USA, 2004. ACM. ISBN 1-58113-702-8. doi: 10.1145/985692.985706. URL `http://doi.acm.org/10.1145/985692.985706`.

Agnieszka (Aga) Bojko. Informative or Misleading? Heatmaps Deconstructed. *Human-Computer Interaction. New Trends*, Julie A Jacko, editor, pages 30–39, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg. ISBN 978-3-642-02574-7.

Maged N Kamel Boulos, Ann C Brewer, Chante Karimkhani, David B Buller and Robert P Dellavalle. Mobile medical and health apps: state of the art, concerns, regulatory control and certification. *Online journal of public health informatics*, 5(3):229, 2 2014. ISSN 1947-2579. doi: 10.5210/ojphi.v5i3.4814. URL `https://www.ncbi.nlm.nih.gov/pubmed/24683442https://www.ncbi.nlm.nih.gov/pmc/PMC3959919/`.

Jesus J Caban and David Gotz. Visual analytics in healthcare – opportunities and research challenges. *Journal of the American Medical Informatics Association*, 22 (2):260–262, 2015. doi: 10.1093/jamia/ocv006. URL `http://dx.doi.org/10.1093/jamia/ocv006`.

E K Choe, B Lee and m. c. schraefel. Characterizing Visualization Insights from Quantified Selfers' Personal Data Presentations. *IEEE Computer Graphics and Applications*, 35(4):28–37, 2015. ISSN 0272-1716. doi: 10.1109/MCG.2015.51.

Kate Crawford, Jessa Lingel and Tero Karppi. Our metrics, ourselves: A hundred years of self-tracking from the weight scale to the wrist wearable device. *European Journal of Cultural Studies*, 18(4-5):479–496, 6 2015. ISSN 1367-5494. doi: 10.1177/1367549415584857. URL `https://doi.org/10.1177/1367549415584857`.

Andrea Cuttone, Michael Kai Petersen and Jakob Eg Larsen. Four data visualization heuristics to facilitate reflection in personal informatics. *International Conference on Universal Access in Human-Computer Interaction*, pages 541–552, 2014.

CxAalto. Niimpy - Python module for analysis of behavorial data, 2018. URL `https://github.com/CxAalto/niimpy`.

Dawn Dowding, Rebecca Randell, Peter Gardner, Geraldine Fitzpatrick, Patricia Dykes, Jesus Favela, Susan Hamer, Zac Whitewood-Moores, Nicholas Hardiker, Elizabeth Borycki and Leanne Currie. Dashboards for improving patient care: Review of the literature. *International Journal of Medical Informatics*, 84(2):87–100, 2015. ISSN 1386-5056. doi: https://doi.org/10.1016/j.ijmedinf.2014.10.001. URL `http://www.sciencedirect.com/science/article/pii/S1386505614001890`.

F Du, B Shneiderman, C Plaisant, S Malik and A Perer. Coping with Volume and Variety in Temporal Event Sequences: Strategies for Sharpening Analytic Focus. *IEEE Transactions on Visualization and Computer Graphics*, 23(6):1636–1649, 2017. ISSN 1077-2626. doi: 10.1109/TVCG.2016.2539960.

Andrew T Duchowski, Margaux M Price, Miriah Meyer and Pilar Orero. Aggregate Gaze Visualization with Real-time Heatmaps. *Proceedings of the Symposium on Eye Tracking Research and Applications*, ETRA '12, pages 13–20, New York, NY, USA, 2012. ACM. ISBN 978-1-4503-1221-9. doi: 10.1145/2168556.2168558. URL `http://doi.acm.org/10.1145/2168556.2168558`.

J A Dykes and D M Mountain. Seeking structure in records of spatio-temporal behaviour: visualization issues, efforts and applications. *Computational Statistics & Data Analysis*, 43(4):581–603, 2003. ISSN 0167-9473. doi: https://doi.org/10.1016/S0167-9473(02)00294-3. URL `http://www.sciencedirect.com/science/article/pii/S0167947302002943`.

Empatica. Embrace by Empatica is the world's first smart watch to be cleared by FDA for use in Neurology, 2018. URL `https://www.prnewswire.com/news-releases/`

embrace-by-empatica-is-the-worlds-first-smart-watch-to-be-cleared-by-fda-for-use-i
html.

Daniel Epstein, Felicia Cordeiro, Elizabeth Bales, James Fogarty and Sean Munson. Taming Data Complexity in Lifelogs: Exploring Visual Cuts of Personal Informatics Data. *Proceedings of the 2014 Conference on Designing Interactive Systems*, DIS '14, pages 667–676, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-2902-6. doi: 10.1145/2598510.2598558. URL http://doi.acm.org/10.1145/2598510.2598558.

Paul A Estabrooks, Maureen Boyle, Karen M Emmons, Russell E Glasgow, Bradford W Hesse, Robert M Kaplan, Alexander H Krist, Richard P Moser and Martina V Taylor. Harmonized patient-reported data elements in the electronic health record: supporting meaningful use by primary care action on health behaviors and key psychosocial factors. *Journal of the American Medical Informatics Association*, 19(4):575–582, 7 2012. ISSN 1067-5027. URL http://dx.doi.org/10.1136/amiajnl-2011-000576.

Sarah Faisal, Ann Blandford and Henry W W Potts. Making sense of personal health information: Challenges for information visualization. *Health Informatics Journal*, 19 (3):198–217, 8 2013. ISSN 1460-4582. doi: 10.1177/1460458212465213. URL https://doi.org/10.1177/1460458212465213.

Fitbit Inc. Fitbit charge3, 2018. URL https://www.fitbit.com/fi/charge3.

Nils Gehlenborg and Bang Wong. Heat maps. *Nature Methods*, 9:213, 2 2012. URL http://dx.doi.org/10.1038/nmeth.190210.1038/nmeth.1902.

Fosca Giannotti, Mirco Nanni, Dino Pedreschi, Fabio Pinelli, Chiara Renso, Salvatore Rinzivillo and Roberto Trasarti. Unveiling the Complexity of Human Mobility by Querying and Mining Massive Trajectory Data. *The VLDB Journal*, 20(5):695–719, 10 2011. ISSN 1066-8888. doi: 10.1007/s00778-011-0244-8. URL http://dx.doi.org/10.1007/s00778-011-0244-8.

D Gotz and D Borland. Data-Driven Healthcare: Challenges and Opportunities for Interactive Visualization. *IEEE Computer Graphics and Applications*, 36(3):90–96, 2016. ISSN 0272-1716 VO - 36. doi: 10.1109/MCG.2016.59.

Stanford Visualization Group. Adapted from image : Horizon Graphs, 2018a. URL http://mbostock.github.io/protovis/ex/horizon.html.

Stanford Visualization Group. Index Charts, 2018b. URL http://mbostock.github.io/protovis/ex/index-chart.html.

A Grünerbl, A Muaremi, V Osmani, G Bahle, S Öhler, G Tröster, O Mayora, C Haring and P Lukowicz. Smartphone-Based Recognition of States and State Changes in

Bipolar Disorder Patients. *IEEE Journal of Biomedical and Health Informatics*, 19 (1):140–148, 2015. ISSN 2168-2194. doi: 10.1109/JBHI.2014.2343154.

Torsten Hägerstraand. What about people in regional science? *Papers in regional science*, 24(1):7–24, 1970.

Jeffrey Heer, Nicholas Kong and Maneesh Agrawala. Sizing the Horizon: The Effects of Chart Size and Layering on the Graphical Perception of Time Series Visualizations. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '09, pages 1303–1312, New York, NY, USA, 2009. ACM. ISBN 978-1-60558-246-7. doi: 10.1145/1518701.1518897. URL `http://doi.acm.org/10.1145/1518701.1518897`.

Jeffrey Heer, Michael Bostock, Vadim Ogievetsky and others. A tour through the visualization zoo. *Commun. Acm*, 53(6):59–67, 2010.

Hello Code. feature_mood_full.png, 2018. URL `https://exist.io/static/feature_mood.html`.

John P Higgins. Smartphone Applications for Patients' Health and Fitness. *The American Journal of Medicine*, 129(1):11–19, 1 2016. ISSN 0002-9343. doi: 10.1016/j.amjmed.2015.05.038. URL `https://doi.org/10.1016/j.amjmed.2015.05.038`.

Jerry L Hintze and Ray D Nelson. Violin Plots: A Box Plot-Density Trace Synergism. *The American Statistician*, 52(2):181–184, 5 1998. ISSN 0003-1305. doi: 10.1080/00031305.1998.10480559. URL `https://www.tandfonline.com/doi/abs/10.1080/00031305.1998.10480559`.

Jeff Huang, Ryen W White and Susan Dumais. No Clicks, No Problem: Using Cursor Movements to Understand and Improve Search. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, pages 1225–1234, New York, NY, USA, 2011. ACM. ISBN 978-1-4503-0228-9. doi: 10.1145/1978942.1979125. URL `http://doi.acm.org/10.1145/1978942.1979125`.

T R Insel. Digital phenotyping: Technology for a new science of behavior. *JAMA*, 318 (13):1215–1216, 10 2017. ISSN 0098-7484. URL `http://dx.doi.org/10.1001/jama.2017.11295`.

Peter Kampstra and others. Beanplot: A boxplot alternative for visual comparison of distributions. 2008.

Denis Klimov and Yuval Shahar. A framework for intelligent visualization of multiple time-oriented medical records. *AMIA ... Annual Symposium proceedings. AMIA Symposium*, 2005:405–409, 2005. ISSN 1942-597X. URL `https://www.ncbi.nlm.nih.gov/pubmed/16779071https://www.ncbi.nlm.nih.gov/pmc/PMC1560450/`.

Emma Komulainen, Katarina Meskanen, Jari Lipsanen, Jari Marko Lahti, Pekka Jylhä, Tarja Melartin, Marieke Wichers, Erkki Isometsä and Jesper Ekelund. The Effect of Personality on Daily Life Emotional Processes. *PLOS ONE*, 9(10):e110907–, 10 2014. URL `https://doi.org/10.1371/journal.pone.0110907`.

Ian Li, Anind Dey and Jodi Forlizzi. A Stage-based Model of Personal Informatics Systems. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '10, pages 557–566, New York, NY, USA, 2010. ACM. ISBN 978-1-60558-929-9. doi: 10.1145/1753326.1753409. URL `http://doi.acm.org/10.1145/1753326.1753409`.

A Madan, M Cebrian, S Moturu, K Farrahi and A ". Pentland. Sensing the "Health State" of a Community. *IEEE Pervasive Computing*, 11(4):36–45, 2012. ISSN 1536-1268. doi: 10.1109/MPRV.2011.79.

Alessandro Marcengo and Amon Rapp. *Visualization of Human Behavior Data: The Quantified Self*. IGI Global, 7 2013. ISBN 9781466643093. doi: 10.4018/978-1-4666-4309-3.ch012. URL `https://www.igi-global.com/chapter/visualization-of-human-behavior-data/150231`.

Tom May. Day / Hour Heatmap, 2018. URL `http://bl.ocks.org/tjdecke/5558084`.

David C Mohr, Michelle Nicole Burns, Stephen M Schueller, Gregory Clarke and Michael Klinkman. Behavioral Intervention Technologies: Evidence review and recommendations for future research in mental health. *General Hospital Psychiatry*, 35(4):332–338, 2013. ISSN 0163-8343. doi: https://doi.org/10.1016/j.genhosppsych.2013.03.008. URL `http://www.sciencedirect.com/science/article/pii/S0163834313000698`.

Timothy Morgenthaler, Cathy Alessi, Leah Friedman, Judith Owens, Vishesh Kapur, Brian Boehlecke, Terry Brown, Jr. Chesson Andrew, Jack Coleman, Teofilo Lee-Chiong, Jeffrey Pancer and Todd J Swick. Practice Parameters for the Use of Actigraphy in the Assessment of Sleep and Sleep Disorders: An Update for 2007. *Sleep*, 30(4):519–529, 4 2007. ISSN 0161-8105. URL `http://dx.doi.org/10.1093/sleep/30.4.519`.

MSKTC. Effective Use of Line Charts, 2018. URL `https://msktc.org/lib/docs/KT_Toolkit/Charts_and_Graphs/Charts_Tool_line_508c.pdf`.

T B Murdoch and A S Detsky. The inevitable application of big data to health care. *JAMA*, 309(13):1351–1352, 4 2013. ISSN 0098-7484. URL `http://dx.doi.org/10.1001/jama.2013.393`.

Mark A Musen, Blackford Middleton and Robert A Greenes. Clinical Decision-Support Systems. In *Biomedical Informatics: Computer Applications in Health Care and Biomedicine*, Edward H Shortliffe and James J Cimino, editors, pages 643–674. Springer London, London, 2014. ISBN 978-1-4471-4474-8. doi: 10.1007/978-1-4471-4474-8{_}22. URL https://doi.org/10.1007/978-1-4471-4474-8_22.

C North. Toward measuring visualization insight. *IEEE Computer Graphics and Applications*, 26(3):6–9, 2006. ISSN 0272-1716. doi: 10.1109/MCG.2006.70.

David Nutt, Sue Wilson and Louise Paterson. Sleep disorders as core symptoms of depression. *Dialogues in clinical neuroscience*, 10(3):329–336, 9 2008. ISSN 1294-8322. URL https://www.ncbi.nlm.nih.gov/pubmed/18979946https://www.ncbi.nlm.nih.gov/pmc/PMC3181883/.

Jukka-Pekka Onnela and Scott L Rauch. Harnessing Smartphone-Based Digital Phenotyping to Enhance Behavioral and Mental Health. *Neuropsychopharmacology*, 41:1691, 1 2016. URL https://doi.org/10.1038/npp.2016.710.1038/npp.2016.7.

Catherine Plaisant, Brett Milash, Anne Rose, Seth Widoff and Ben Shneiderman. kaiserlifelines.png, 1996. URL http://www.cs.umd.edu/projects/hcil/lifelines/kaiserlifelines.jpg.

Catherine Plaisant, Richard Mushlin, Aaron Snyder, Jia Li, Dan Heller and Ben Shneiderman. LifeLines: Using Visualization to Enhance Navigation and Analysis of Patient Records. In *The Craft of Information Visualization*, BENJAMIN B BEDERSON and B E N SHNEIDERMAN, editors, pages 308–312. Morgan Kaufmann, San Francisco, 2003. ISBN 978-1-55860-915-0. doi: https://doi.org/10.1016/B978-155860915-0/50038-X. URL http://www.sciencedirect.com/science/article/pii/B978155860915050038X.

Plotly. Plotly Python Open Source Graphing Library, 2018. URL https://plot.ly/python/.

Kristin Potter, Hans Hagen, Andreas Kerren and Peter Dannenmann. Methods for presenting statistical information: The box plot. *Visualization of large and unstructured data sets*, 4:97–106, 2006.

Abhishek Pratap, David C Atkins, Brenna N Renn, Michael J Tanana, Sean D Mooney, Joaquin A Anguera and Patricia A Areán. The accuracy of passive phone sensors in predicting daily mood. *Depression and Anxiety*, 0(0), 8 2018. ISSN 1091-4269. doi: 10.1002/da.22822. URL https://doi.org/10.1002/da.22822.

Propeller Health. Propeller, 2018. URL https://www.propellerhealth.com/how-it-works/.

Daniel Rainham, Ian McDowell, Daniel Krewski and Mike Sawada. Conceptualizing the healthscape: Contributions of time geography, location technologies and spatial ecology to place and health research. *Social Science & Medicine*, 70(5):668–676, 2010. ISSN 0277-9536. doi: https://doi.org/10.1016/j.socscimed.2009.10.035. URL http://www.sciencedirect.com/science/article/pii/S0277953609007254.

Amon Rapp and Federica Cena. Self-monitoring and technology: challenges and open issues in personal informatics. *International Conference on Universal Access in Human-Computer Interaction*, pages 613–622, 2014.

Amon Rapp and Federica Cena. Personal informatics for everyday life: How users without prior self-tracking experience engage with personal data. *International Journal of Human-Computer Studies*, 94:1–17, 2016. ISSN 1071-5819. doi: https://doi.org/10.1016/j.ijhcs.2016.05.006. URL http://www.sciencedirect.com/science/article/pii/S107158191630060X.

William T Riley, Daniel E Rivera, Audie A Atienza, Wendy Nilsen, Susannah M Allison and Robin Mermelstein. Health behavior models in the age of mobile interventions: are our theories up to the task? *Translational Behavioral Medicine*, 1(1):53–71, 3 2011. ISSN 1869-6716. URL http://dx.doi.org/10.1007/s13142-011-0021-7.

Alexander Rind, Taowei David Wang, Wolfgang Aigner, Silvia Miksch, Krist Wongsuphasawat, Catherine Plaisant and Ben Shneiderman. Interactive Information Visualization to Explore and Query Electronic Health Records. *Foundations and Trends® in Human–Computer Interaction*, 5(3):207–298, 2013. ISSN 1551-3955. doi: 10.1561/1100000039. URL http://dx.doi.org/10.1561/1100000039.

Avi Sadeh and Christine Acebo. The role of actigraphy in sleep medicine. *Sleep Medicine Reviews*, 6(2):113–124, 2002. ISSN 1087-0792. doi: https://doi.org/10.1053/smrv.2001.0182. URL http://www.sciencedirect.com/science/article/pii/S1087079201901820.

Sohrab Saeb, Mi Zhang, Christopher J Karr, Stephen M Schueller, Marya E Corden, Konrad P Kording and David C Mohr. Mobile Phone Sensor Correlates of Depressive Symptom Severity in Daily-Life Behavior: An Exploratory Study. *Journal of Medical Internet Research*, 17(7):e175, 7 2015. ISSN 1439-4456. doi: 10.2196/jmir.4273. URL http://www.ncbi.nlm.nih.gov/pmc/articles/PMC4526997/.

B Shneiderman. The eyes have it: a task by data type taxonomy for information visualizations. *Proceedings 1996 IEEE Symposium on Visual Languages*, pages 336–343, 1996. ISBN 1049-2615. doi: 10.1109/VL.1996.545307.

B Shneiderman, C Plaisant and B W Hesse. Improving Healthcare with Interactive Visualization. *Computer*, 46(5):58–66, 2013. ISSN 0018-9162 VO - 46. doi: 10.1109/MC.2013.38.

Sports Tracking Technologies. Sports Tracker - Power your passion for sports, 2018. URL `https://www.sports-tracker.com/`.

Jennifer G Stadler, Kipp Donlon, Jordan D Siewert, Tessa Franken and Nathaniel E Lewis. Improving the Efficiency and Ease of Healthcare Analysis Through Use of Data Visualization Dashboards. *Big Data*, 4(2):129–135, 2016. doi: 10.1089/big.2015.0059. URL `https://doi.org/10.1089/big.2015.0059`.

Andreas Ströhle. Physical activity, exercise, depression and anxiety disorders. *Journal of Neural Transmission*, 116(6):777, 2008. ISSN 1435-1463. doi: 10.1007/s00702-008-0092-x. URL `https://doi.org/10.1007/s00702-008-0092-x`.

Daniel Sui. Looking through Hägerstrand's dual vistas: towards a unifying framework for time geography. *Journal of Transport Geography*, 23:5–16, 2012. ISSN 0966-6923. doi: https://doi.org/10.1016/j.jtrangeo.2012.03.020. URL `http://www.sciencedirect.com/science/article/pii/S096669231200083X`.

James J Thomas and Kristin A Cook. A Visual Analytics Agenda. *IEEE Comput. Graph. Appl.*, 26(1):10–13, 1 2006. ISSN 0272-1716. doi: 10.1109/MCG.2006.5. URL `http://dx.doi.org/10.1109/MCG.2006.5`.

Sara Thomée, Annika Härenstam and Mats Hagberg. Mobile phone use and stress, sleep disturbances, and symptoms of depression among young adults - a prospective cohort study. *BMC Public Health*, 11(1):66, 2011. ISSN 1471-2458. doi: 10.1186/1471-2458-11-66. URL `https://doi.org/10.1186/1471-2458-11-66`.

Edward R Tufte. *The Visual Display of Quantitative Information.* Graphics Press, Cheshire, CT, USA, 1986. ISBN 0-9613921-0-X.

Julie Vallée, Emmanuelle Cadot, Christelle Roustit, Isabelle Parizot and Pierre Chauvin. The role of daily mobility in mental health inequalities: The interactive influence of activity space and neighbourhood of residence on depression. *Social Science & Medicine*, 73(8):1133–1144, 2011. ISSN 0277-9536. doi: https://doi.org/10.1016/j.socscimed.2011.08.009. URL `http://www.sciencedirect.com/science/article/pii/S0277953611004977`.

Taowei David Wang, Catherine Plaisant, Alexander J Quinn, Roman Stanchak, Shawn Murphy and Ben Shneiderman. Aligning Temporal Data by Sentinel Events: Discovering Patterns in Electronic Health Records. *Proceedings of the SIGCHI*

Conference on Human Factors in Computing Systems, CHI '08, pages 457–466, New York, NY, USA, 2008. ACM. ISBN 978-1-60558-011-1. doi: 10.1145/1357054.1357129. URL http://doi.acm.org/10.1145/1357054.1357129.

Taowei David Wang, Krist Wongsuphasawat, Catherine Plaisant and Ben Shneiderman. Visual Information Seeking in Multiple Electronic Health Records: Design Recommendations and a Process Model. *Proceedings of the 1st ACM International Health Informatics Symposium*, IHI '10, pages 46–55, New York, NY, USA, 2010. ACM. ISBN 978-1-4503-0030-8. doi: 10.1145/1882992.1883001. URL http://doi.acm.org/10.1145/1882992.1883001.

Colin Ware. *Information Visualization: Perception for Design*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2004. ISBN 1558608192.

Peter West, Richard Giordano, Max Van Kleek and Nigel Shadbolt. The Quantified Patient in the Doctor's Office: Challenges & Opportunities. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, CHI '16, pages 3066–3078, New York, NY, USA, 2016. ACM. ISBN 978-1-4503-3362-7. doi: 10.1145/2858036.2858445. URL http://doi.acm.org/10.1145/2858036.2858445.

Vivian L West, David Borland and W Ed Hammond. Innovative information visualization of electronic health record data: a systematic review. *Journal of the American Medical Informatics Association*, 22(2):330–339, 3 2015. ISSN 1067-5027. URL http://dx.doi.org/10.1136/amiajnl-2014-002955.

Blair Wheaton and Philippa Clarke. Space Meets Time: Integrating Temporal and Contextual Influences on Mental Health in Early Adulthood. *American Sociological Review*, 68(5):680–706, 2003. ISSN 00031224. doi: 10.2307/1519758. URL http://www.jstor.org/stable/1519758.

D F Williamson, R A Parker and J S Kendrick. The box plot: A simple visual method to interpret data. *Annals of Internal Medicine*, 110(11):916–921, 6 1989. ISSN 0003-4819. URL http://dx.doi.org/10.7326/0003-4819-110-11-916.

Shannon N Zenk, Amy J Schulz, Stephen A Matthews, Angela Odoms-Young, JoEllen Wilbur, Lani Wegrzyn, Kevin Gibbs, Carol Braunschweig and Carmen Stokes. Activity space environment and dietary and physical activity behaviors: A pilot study. *Health & Place*, 17(5):1150–1161, 2011. ISSN 1353-8292. doi: https://doi.org/10.1016/j.healthplace.2011.05.001. URL http://www.sciencedirect.com/science/article/pii/S1353829211000797.