Four Data Visualization Heuristics to Facilitate Reflection in Personal Informatics

Andrea Cuttone, Michael Kai Petersen, and Jakob Eg Larsen

Technical University of Denmark (DTU)
Department of Applied Mathematics and Computer Science
Cognitive Systems Section
Matematiktorvet 303B, 2800 Kgs Lyngby, Denmark
{ancu,mkai,jaeg}@dtu.dk

Abstract. In this paper we discuss how to facilitate the process of reflection in Personal Informatics and Quantified Self systems through interactive data visualizations. Four heuristics for the design and evaluation of such systems have been identified through analysis of self-tracking devices and apps. Dashboard interface paradigms in specific self-tracking devices (Fitbit and Basis) are discussed as representative examples of state of the art in feedback and reflection support. By relating to existing work in other domains, such as event related representation of time series multivariate data in financial analytics, it is discussed how the heuristics could guide designs that would further facilitate reflection in self-tracking personal informatics systems.

Keywords: personal informatics, quantified self, self-tracking, information visualization, feedback, reflection, heuristics.

1 Introduction

In recent years self-tracking and lifelogging have received increased interest with the introduction of a wide variety of low-cost mobile apps, wearable computers, and sensors. These devices allow easy collection of data that can describe various aspects of human behavior. However, making sense of the ever increasing amounts of everyday self-tracking data retrieved across multiple domains create new demands for turning data points and trends into affordances for action. The reflection stage is a fundamental component in modeling and using Personal Informatics (PI) systems [18,8,16] to facilitate an understanding of self-tracking data reflecting daily habitual patterns and to make such data actionable for behavioral change [9]. Different solutions have been suggested to facilitate self-reflection, including the usage of charts [18,20,21,6], avatars [22,12,14], notifications [1], narrative [23] and abstract art [7,10]. Although these solutions may facilitate increased awareness due to the fact that behavioural aspects are being observed, the process of turning observations and insights into actions remains a challenge. Even a recent review of activity trackers in *The New York Times*[1]

[1] http://www.nytimes.com/2014/01/30/technology/personaltech/review-the-fitbit-force-activity-tracker.html Last accessed Jan 29, 2014.

C. Stephanidis and M. Antona (Eds.): UAHCI/HCII 2014, Part IV, LNCS 8516, pp. 541–552, 2014.

© Springer International Publishing Switzerland 2014

emphasized that while self-tracking devices enable the user to collect behavioral data, they fall short of assisting the user in learning how to change habits.

One element of personal informatics is the iterative process with self-reflection questions phrased by a user and feedback provided by a self-tracking system to answer those questions. However, we suggest that state-of-the art systems offer fairly limited flexibility in terms of the types of questions that can be phrased, and the possible feedback that can be provided. In a broader perspective, the self-tracking data obtained might be characterized as quantitative time series data which combines behavioral data with associated discrete events. Similar to how financial analytics like those provided by Bloomberg might combine a vertical flow of business related earning reports or corporate news updates, with distinct time stamped markers, outlined horizontally within the continuous timeline fluctuations of stock values. Thus hierarchically adding layers of relevant information embedded both within the chart and in adjacent panels linked to external events, that may facilitate interpretation or be a direct cause of rising or falling trends visualized in the data [26]. The emphasis of integrating distinct interpretable events in continuous flows of quantitative data within financial analytics reflects a need for these interfaces to provide a foundation for taking concrete action related to aspects of optimizing profits or avoiding a loss. We suggest that these advances within financial analytics software for interpretation of complex data may both provide underlying design patterns for long term comparison of trends in multivariate flows, as well as defining thresholds which could likewise turn quantified self generated data into actionable parameters for optimizing lifestyle in personal informatics systems.

2 Related Work

Several frameworks have been proposed to formalize the reflection process in personal informatics. Li et al. [19] identify six kinds of questions for reflection: *Status* (what is my situation now?), *History* (what was my situation in the past?), *Goals* (what future status should I aim for?), *Discrepancies* (how does my status compare with my goals?), *Context* (what affects my status?), *Factors* (how are different attributes related?). Moreover, two alternating phases are defined: *Maintenance* (known relation between status and behavior) and *Discovery* (not known goals or effect of behavior). Fleck et al. [8] define a multi-layer reflection framework: *Description*, *Reflective Description*, *Dialogic Reflection*, *Transformative Reflection*, *Critical Reflection*. Each layer builds on top of the previous, and corresponds to a deeper understanding of personal data. Guidelines for facilitating reflection are proposed, including supporting questions and providing multiple perspectives on the data. Rivera et al. [24] apply Boud's reflective learning framework to personal informatics, and identifies two levels of reflection: *Triggering* (active notification or passive feedback) and *Recalling* (aggregating, contextualizing, visualizing).

Several feedback schemes have been suggested including avatar-based feedback that employs a virtual object to represent a judgment on behavior. These solutions exploit participants' empathy with the virtual avatars to persuade them in adopting positive behavioral changes. For example, *Fish'n'Steps* [22] provides feedback about daily step count as a virtual fish, *Ubigreen* [12] using virtual trees and polar bears to provide feedback on green transportation habits, and *UbiFit Garden* [5] represents fitness activity as a virtual garden. *Spark* [7] visualizes physical activity as abstract art through an ambient display, whereas *Lifestyle Stories* [23] provides feedback about mobile sensing personal data in form of stories composed by events of various categories. Many commercial self-tracker systems employ a combination of traditional charts, maps and dashboards (see for example *Nike+*[2], *Fitbit*[3], *Basis*[4], *Jawbone UP*[5], *Mint*[6], *DailyBurn*[7], *Moves*[8]).

3 Reflection as Data Analysis

We suggest to treat reflection as data analysis on personal information. What are the crucial questions or answers that analytics should provide? In Bloomberg financial analytics the ability to couple external events to timeline charts appears crucial for interpreting the causality behind the data. One might in a more general context consider online news media like *The Wall Street Journal* or Twitter feeds an expanded version of this paradigm, providing not only the current market data but also a highly curated selection of background material as well as live updates on events, that would provide the necessary foundation for making informed decisions. Even in admirably simple single sensor quantified self apps like Fitbit, limited to measuring the number of steps taken during the day or week, data might provide valuable insights into user behavior. But it would require that annotations are added automatically with calendar events or smartphone location data, thereby enriching the representation beyond the current ability of manually attaching labels. We see a similar potential for advanced self-tracking devices with multiple sensors like Basis, which likewise translates complex patterns of behaviors into singular goal oriented habits to be fulfilled on a daily basis at regular hours. Coupling calendar events for monitoring heart rate related to specific physical tasks, indicating how this sensor data is correlated to differences in sleep patterns, or influenced by levels of exercise across weeks, might provide additional value.

A user may want to retrieve specific information from his own dataset (current status, progress), or explore it for finding interesting patterns. In order to facilitate this analysis, we propose to use data visualization, that is the representation of data using position, size, shape, color, and text [3]. Visualization facilitates analysis by exploiting the human visual system, which is extremely

[2] http://nikeplus.nike.com/plus/
[3] http://www.fitbit.com/
[4] http://www.mybasis.com/
[5] http://jawbone.com/up
[6] http://www.mint.com/
[7] http://dailyburn.com/
[8] http://www.moves-app.com/

good at processing large quantities of information and spotting patterns. Visualization is widely used in Exploratory Data Analysis [28], a statistical technique for exploring datasets, in order to gain insights, obtain a better understanding, spot patterns, trends, correlation, and outliers. In many cases, the data analyst does not know in advance which specific question to ask, so he can explore the data in order to find interesting patterns. This process is highly iterative, as once a question has been posed, its answer often leads to more questions to be asked. Similarly, in PI systems the user may not know which questions to ask, or may not be interested in a specific question but in exploring his own data for curiosity. Indeed, one of the barriers for reflection is not knowing which questions to ask to personal data [18]. We can represent this iterative process as a cycle between questions asked, and feedback provided. We define *question space* the set of all possible questions, and *feedback space* the set of all visualizations. Each feedback type can answer one or more questions, and each question can be answered by one or more type. In data analysis, there are a number of common questions that can be asked: distribution of values (mostly around a central value and gradually less on the sides? Mostly for a value and very rapidly decaying? Multiple peaks?), grouping and outliers (are there values much different than most of the others? Are items clustered into groups?), correlation (what is the relation between x and y? Is there a linear, quadratic, exponential, sinusoidal trend?), geographical (how are values related to locations? Are there locations with similar values?), connectivity (are there items related together? Are there items which are more tightly connected? Are there non-connected items?) We identify the most relevant questions for the goal of self-reflection, and we summarize them into heuristics.

4 Design Heuristics

In this section we introduce four design heuristics that can be applied as a guideline for creating and evaluating interactive visualizations of self-tracking data with the aim to facilitate effective exploration of personal data and make such data actionable for behavior change. Throughout the discussion of the heuristics we relate to existing state of the art self-tracking systems using Fitbit and Basis as examples of personal informatics systems with interactive visualizations. We do not intend to criticize these two systems in particular, but rather consider them as representative and illustrative examples of state of the art in personal self-tracking systems. The scope of the discussion is limited to facilitating the reflection process, while acknowledging that further discussion is needed in terms of providing actionable items as well as other aspects of reflection in personal informatics as mentioned in the Related Work Section.

4.1 Make Data Interpretable at a Glance

Often users want to obtain answers to a question with the minimal effort and time. For this reason, data visualizations optimized for interpretation at a glance are needed, in order to provide a swift overview of personal tracking activities, and to augment and support subjective recollection. Quantified self apps like Fitbit or Basis may aim to simplify visualization of complex patterns by transforming the collected measurements into single activity dashboard dials or progress bars reflecting goal oriented accomplishments. Figure 1 shows the personal dashboard provided by Fitbit, which reduces the collected data to simple indications of (daily) goal fulfillment (percentage) and an overview of daily activity levels. Although it provides an overview of the level of goal fulfillment this division of activities into separate silos makes it a challenge to interpret the data in a larger context. In contrast financial analytics interfaces like those provided by Bloomberg [26] may contain large amounts of data which is nevertheless made interpretable based on established conventions for using dynamically changing font colors to signify up- or downward moving prices, or positive negative outlooks based on earning reports, which when collapsed form independent parallel layers of color coded trend lines that remain surprisingly legible on top of contrasting neutral background screens.

Fig. 1. The Fitbit dashboard show daily goal fulfillment (percentage) and an overview of daily activity

While data visualizations can be very useful they may be complex and difficult to interpret and understand. The complexity of visualizations can range from data-poor, informal infographics for the general audience to complex, rigorous scientific visualization aimed at scholars. Several issues should be considered when designing a visualization, including the technical and domain knowledge of the users, the goals of reflection (exploring data, asking specific questions, testing hypothesis), the time and effort expected from the user (a quick glance or a long interaction?). We suggest to provide a simple visualization as the starting point, and allow advanced users to dynamically increase the level of complexity and details. The usage of explanatory elements, such as text, axes, legends, annotations can greatly facilitate the comprehension. Visualizing data is often compared to storytelling, especially in the fields of journalism and business reporting. Some authors prefer to present the data as raw as possible, with little or no annotations and highlighting, in order to give the reader full freedom in the interpretation of the facts behind the data. Others prefer to editorialize the data to various degrees, by marking samples, comparing with other distributions, providing comments. We argue that a certain editorialization is good for PI system, as it can act as a persuasive force towards positive behavior. For instance a fitness tracker system that encourages the user to be more active by visualizing and forecasting the positive consequences.

4.2 Enable Exploration of Patterns in Time Series Data

There are two fundamental patterns to analyze: global trends (does the variable increase, decrease or remain constant over a period of time?) and periodic patterns (does the variable value change with a repeating pattern, for example hourly, weekly or yearly?).

Although several self-tracking app interfaces emphasize simplified dashboard representations of accumulated data which limits exploration, the recently added Basis sleep monitoring goes far beyond the previous single modality heatmaps by breaking down total sleep duration into the different phases of rapid eye movement (REM) and deep sleep, thereby making it possible to translate these periodic patterns into quantifiable aspects of mind and body refresh, see Figure 2.

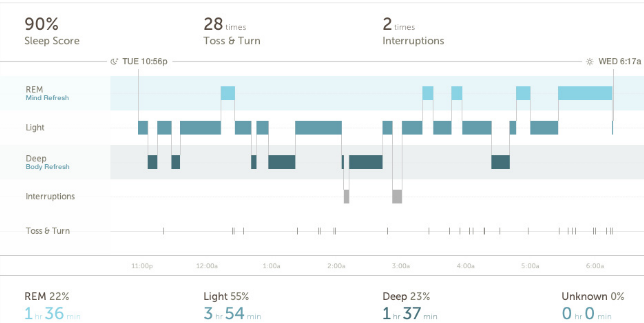

Fig. 2. The Basis sleep visualization interface also break continuous sleep data into discrete sleep phases including rapid eye movement (REM), light, and deep sleep

Time series analysis is a common task for self-trackers, which may be interested in observing changes over time, periodic patterns, rate of change, and time left to reach goals. The most common representation for time series data is *line plots*, which allow to easily see the overall change over time of a variable. Due to the unavoidable noise, it is useful to add *trend lines* such a LOWESS [4] or least squares. This enable the support of forecasting on future status if the current behavior is kept or modified, thus a prediction of future values can be

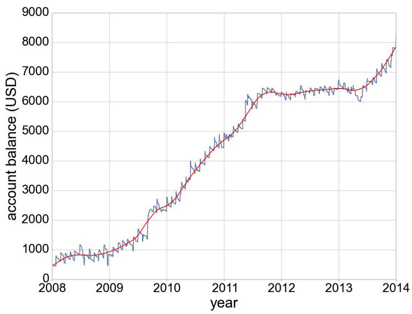

Fig. 3. Amount of money in a bank saving account, visualized as line plot with LOWESS trend line

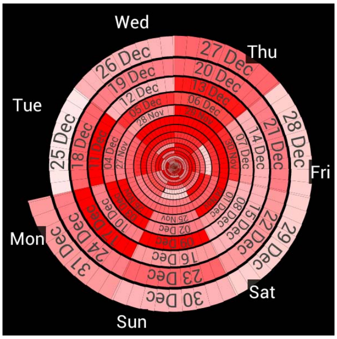

Fig. 4. Number of steps per day represented as a spiral heatmap, with colors from white (low) to red (large)

visualized [15]. As an example Figure 3 shows the amount over time of money in a bank saving account with a clear trend of increase among the individual month-to-month fluctuations. When the duration of time periods is a factor, *timelines* may be used to represent events in a linear layout to captures the temporal sequence.

Line plots and linear timelines do not however facilitate the exploration of periodic patterns, which are characteristic of human behavior. A familiar metaphor for displaying regular patterns is a calendar. A *calendar heatmap* represents each day as a cell, and the variable value as the color shade of the cells. Cells can be aligned for example by day of the week to allow to spot weekly patterns. Spirals have been recently proposed [17] as another representation to facilitate the exploration of periodic patterns in the quantified-self domain. A *spiral heatmap* represents each time unit as an arc in the spiral, and variable value as the color shade of the arcs. By choosing different periods, periodic patterns at various time scale can emerge. Figure 4 show a step count value over time as calendar as spiral heatmaps, with a color scale ranging from white (small count) to red (large count).

4.3 Enable Discovery of Trends in Multiple Data Streams

Financial analytics may also offer inspiration in terms of comparison of key performance indicators in multivariate data. As an example *The Wall Street Journal* allows for extensive personalization when exploring moving averages for smoothing fluctuating trends in time series data, high and low relative to previous values, weighted blends of values and their variation over time, which may be further customized based on choice of display graphics, adjustable time frames or sensitivity of measures.

The Basis activity details visualization only allows the user limited possibilities of exploring the relations between multiple time series data in an adjustable

Fig. 5. The Basis Activity Details visualization allows the user to explore the relations between multiple biometric time series data (heart rate, steps, calories, skin temperature, perspiration, and activities) in an adjustable time interval

time interval, as shown in Figure 5. Multivariate analysis is the process of analyzing multiple variables together, in order to find and understand their relation. In the simplest case, two variables x and y are to be compared, and the following relations can exist: direct correlation (y increases as x increases), inverse correlation (y decreases as x increases), or no correlation. For example, fitness trackers may be interested in how the weight loss is affected by exercise and food intake, or productivity trackers may be interested in how coffee intake, sleep patterns, and exercise affect productivity.

The relation between two variable can be visualized as a *scatterplot*, where each variable is represented on one axis. Scatterplots allow to easily spot trends and outliers. If more than two variables are to be compared, a *scatterplot matrix* allows to inspect all possible combinations. A scatterplot matrix is a array of $n \times n$ scatterplots, where the scatterplot S_{ij} displays the relation between variables X_i and X_j. Scatterplot matrices are a very powerful tool, but can also be intimidating for users. A simpler version of multivariate visualization is the *Corrgram* [11] which distinguishes positively and negatively correlation between pairs of variables using color-coding. As a constructed example Figure 6 and 7 show the relation between coffee intake, hours at work, hours of sleep, and steps. The coffee intake appears directly correlated with working hours, and inversely correlated with sleep hours, and no correlation seem to be present with number of steps.

Often users do not need to explore relations of variables between each others, but they are interested in the change of multiple variables over time. A *Streamgraph* [2] can be used when it is important to show the contribution of each variable to a total. Each variable generates a section of different height, and the resulting areas are stacked to form a stream. *Small multiples* [27] allow to display multiple facets of a dataset, often in comparison to time. Each variable is displayed separately in his own subview, and subviews are layed out side-by-side to facilitate

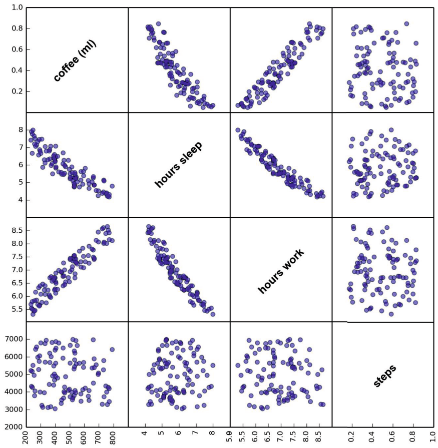

Fig. 6. Scatterplot matrix: all pairs of 4 variables are plotted. The alignment of the data points helps to identify direct, inverse or no correlations.

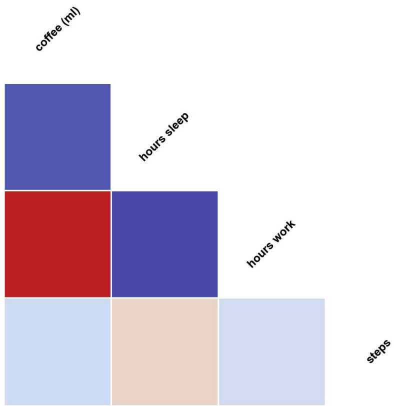

Fig. 7. Corrgram: the correlation between all pair of 4 variables is represented in a color scale from blue (inverse correlation), grey (no correlation) and red (direct correlation)

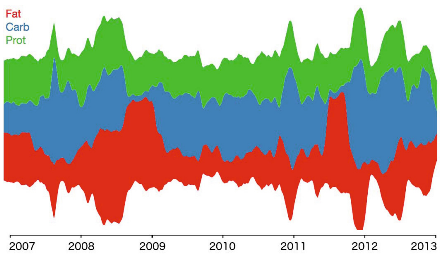

Fig. 8. The composition of dietary intake over time as a streamgraph

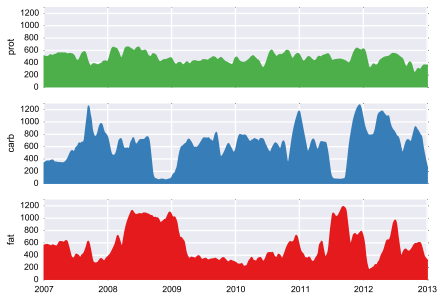

Fig. 9. Small multiples: the three variables of dietary intake are shown side-to-side in separate sub-views

comparison. Figures 8 and 9 show the composition of dietary intake over time. The small multiples enable multivariate comparison, and the streamgraph facilitates the understanding of the total caloric intake.

4.4 Turn Key Metrics into Affordances for Action

The emphasis on interaction in data visualization [13] is reflected in a typical analyst workflow including the generation of a data view, exploration of the result, adjustment of parameters or creation of a completely different visualization in order to further explore the data. This process may lead to new insights or the identification of key metrics. As an example, financial portfolio management often requires a specific action in response to events that cause values to get out of bounds due to regulatory terms that trigger alarms for reevaluation. Likewise in currency trading applications one may need to quickly buy or sell when prices transcend previously set values indicating pain or gain thresholds. In a similar fashion we suggest that the self-tracking workflow involves feedback provided by a personal informatics system which may generate insights about personal data. This may lead to phrasing new questions or directly imposing threshold values that proactively trigger responses to be considered, based on general monitoring of health issues known to be of general concern. This iterative process may be one approach to identify key metrics that can be turned into affordances for actions related to changing behavior.

With the complexity of multi channel self-tracking data sets it may have limited utility to try to visualize all data at once. A user may want to slice the data in various ways, such as by time or category. A user may also want to select a specific set of elements that match a given criteria, such as points inside a geographical region, or values between some thresholds. *Filtering* allows to focus on a specific subset of the data. One of the recommended interaction pattern is "Overview first, zoom and filter, then details-on-demand" [25]. *Navigation* may be supported by allowing scroll and zoom views. When focusing on a subset of the data, the context for the current details could hold valuable information to understand behavior.

The long sequence of interactions with a visualization system, such as filtering, zooming, transforming can be recorded in form of *history*. This log helps the user remember the steps he took, navigate in his interaction sequence, undo eventual mistakes and facilitate a trial-and-error exploration. Providing a visual representation of this history (such a timeline or snapshots of the views) can help the user to orient in his own workflow. In the process of reflection, the user may want to document his findings, write down questions to be investigated, add notes to self. To this end, a visualization can support *annotation* with text and sketches.

These interaction techniques can be readily applied in data visualization in personal systems in order to facilitate the reflection process. Imagine a fitness tracking system, where the user is provided with an overview of his activities. He filters the activity log to a specific part of the year. He views his activity both as a timeline of step counter and as a breakdown of his caloric intake and annotated activities.

5 Conclusions

In this paper we have discussed the support of reflection in state of the art personal informatics systems arguing that it is limited in terms of making observations and insights obtained from interactive visualizations of self-tracking data actionable. We have proposed four heuristic principles for the design and evaluation of interactive data visualization feedback that could further facilitate the process of reflection in self-tracking personal informatics systems. Each design heuristic has been discussed on the basis of an analysis of visualization feedback available in state of the art personal informatics systems.

References

1. Bentley, F., Tollmar, K.: The power of mobile notifications to increase wellbeing logging behavior. In: Proc. of the SIGCHI Conf. on Human Factors in Computing Systems, pp. 1095–1098 (2013)
2. Byron, L., Wattenberg, M.: Stacked graphs-geometry & aesthetics. IEEE Transactions on Visualization and Computer Graphics 14(6), 1245–1252 (2008)
3. Card, S.K., Mackinlay, J.D., Schneiderman, B.: Readings in Information Visualization: Using Vision to Think. Morgan Kaufmann (1999)
4. Cleveland, W.S.: Robust locally weighted regression and smoothing scatterplots. Journal of the American Statistical Association 74(368), 829–836 (1979)
5. Consolvo, S., McDonald, D.W., Toscos, T., Chen, M.Y., Froehlich, J., Harrison, B., Landay, J.A.: Activity sensing in the wild: a field trial of ubifit garden. In: Proc. of the SIGCHI Conf. on Human Factors in Computing Systems, pp. 1797–1806 (2008)
6. Cuttone, A., Lehmann, S., Larsen, J.E.: A mobile personal informatics system with interactive visualizations of mobility and social interactions. In: Proc. of the 1st ACM Int. Workshop on Personal Data Meets Distributed Multimedia (2013)
7. Fan, C., Forlizzi, J., Dey, A.A.: A spark of activity: Exploring informative art as visualization for physical activity. In: Proc. of the 2012 ACM Conf. on Ubiquitous Computing, pp. 81–84 (2012)
8. Fleck, R., Fitzpatrick, G.: Reflecting on reflection: framing a design landscape. In: Proc. of the 22nd Conf. of the Computer-Human Interaction SIG of Australia on Computer-Human Interaction, pp. 216–223 (2010)
9. Fogg, B.J.: Persuasive technology: Using computers to change what we think and do. Ubiquity (5) (2002)
10. Frick, L.: Experiments in self tracking, `http://www.lauriefrick.com/category/work/` (last accessed December 13, 2013)
11. Friendly, M.: Corrgrams: Exploratory displays for correlation matrices. The American Statistician 56(4), 316–324 (2002)
12. Froehlich, J., Dillahunt, T., Klasnja, P., Mankoff, J., Consolvo, S., Harrison, B., Landay, J.A.: UbiGreen: Investigating a mobile tool for tracking and supporting green transportation habits. In: Proc. of the SIGCHI Conf. on Human Factors in Computing Systems, pp. 1043–1052 (2009)
13. Heer, J., Shneiderman, B.: Interactive dynamics for visual analysis. Queue 10(2), 30 (2012)
14. Hekler, E.B., King, A.C., Banerjee, B., Robinson, T., Alonso, M.: A case study of BSUED: Behavioral Science-informed User Experience Design. In: Proc. of the CHI 2011 Workshop on Personal Informatics (2011)
15. Koeman, L., Rogers, Y.: Enabling Foresight and Reflection: Interactive Simulations to Support Behaviour Change. In: Proc. of the CHI 2013 Personal Informatics in the Wild: Hacking Habits for Health & Happiness (2013)
16. Konrad, A., Whittaker, S., Isaacs, E.: Short and long-term benefits of reflective technologies. In: Proc. of the CHI 2013 Workshop on Personal Informatics (2013)
17. Larsen, J.E., Cuttone, A., Lehmann, S.: QS Spiral: Visualizing Periodic Quantified Self Data. In: Proc. of the CHI 2013 Personal Informatics in the Wild: Hacking Habits for Health & Happiness (2013)

18. Li, I., Dey, A., Forlizzi, J.: A stage-based model of personal informatics systems. In: Proc. of the 28th Int. Conf. on Human Factors in Computing Systems, pp. 557–566 (2010)
19. Li, I., Dey, A., Forlizzi, J.: Understanding my data, myself: supporting self-reflection with ubicomp technologies. In: Proc. of UbiComp 2011, pp. 405–414 (2011)
20. Li, I., Medynskiy, Y., Froehlich, J., Larsen, J.E.: Personal informatics in practice: improving quality of life through data. In: CHI 2012 Extended Abstracts on Human Factors in Computing Systems, CHI EA 2012, pp. 2799–2802 (2012)
21. Li, I., Froehlich, J., Larsen, J.E., Grevet, C., Ramirez, E.: Personal informatics in the wild: hacking habits for health & happiness. In: CHI 2013 Extended Abstracts on Human Factors in Computing Systems, pp. 3179–3182 (2013)
22. Lin, J.J., Mamykina, L., Lindtner, S., Delajoux, G., Strub, H.B.: Fish'n'Steps: Encouraging physical activity with an interactive computer game. In: Dourish, P., Friday, A., et al. (eds.) UbiComp 2006. LNCS, vol. 4206, pp. 261–278. Springer, Heidelberg (2006)
23. Pavel, D., Trossen, D., Holweg, M., Callaghan, V.: Lifestyle stories: correlating user information through a story-inspired paradigm. In: Pervasive Computing Technologies for Healthcare, pp. 412–415 (2013)
24. Rivera-Pelayo, V., Zacharias, V., Müller, L., Braun, S.: A framework for applying quantified self approaches to support reflective learning. In: Mobile Learning (2012)
25. Shneiderman, B.: The eyes have it: A task by data type taxonomy for information visualizations. In: Proc. of IEEE Symposium on Visual Languages, pp. 336–343 (1996)
26. Sorenson, E., Brath, R.: Financial Visualization Case Study: Correlating Financial Timeseries and Discrete Events to Support Investment Decisions. IEEE Information Visualisation (2013)
27. Tufte, E.R., Graves-Morris, P.R.: The Visual Display of Quantitative Information, vol. 2. Graphics Press, Cheshire (1983)
28. Tukey, J.W.: Exploratory data analysis, p. 231. Reading, Ma (1977)