

No Forests without Trees: Particulars and Patterns in Visualizing Personal Communication

Ou Jie Zhao
Cornell University
Information Science

301 College Ave., Ithaca NY 14850
1-607-379-9219
oz23@cornell.edu

Tiffany Ng
Cornell University
Information Science

301 College Ave., Ithaca NY 14850
1-607-379-9219
tyn2@cornell.edu

Dan Cosley
Cornell University
Information Science

301 College Ave., Ithaca NY 14850
1-607-379-9219
drc44@cornell.edu

ABSTRACT

When people use visualizations of conversational archives, they typically reflect on particular events, rather than patterns of activity over time. We explore whether this is a fundamental aspect of how people use data to reflect on the past through pieTime, a visualization we developed that focuses on presenting aggregated behavioral data at timescales from hours to months. It builds on work in conversation visualization and lifelogging by focusing on rhythms rather than details, supporting reflection across media, and using phone logs to complement CMC media. A 15-person evaluation supports findings from prior work about the importance of particular details and storytelling in tools that support reflection, even when the design goals emphasize higher-level patterns. Still, aggregate patterns provide additional insight into personal behavior, suggesting that systems that integrate both particulars and patterns may be especially valuable, especially when they also help people build and manage their identities.

Categories and Subject Descriptors

H.1.2 [User/Machine Systems]: Human factors—*human information processing*; H.5.m [Miscellaneous]

General Terms

Design, Human Factors

Keywords

visualization; reflection; temporal rhythms; lifelogging

1. INTRODUCTION

Visualizations and analyses of behavior in social media and computer-mediated communication tools such as email aim to help both researchers and users move from raw records of recorded activity to insights about their communities, relationships, and lives. Harvesting the activities and information people already generate through sources as diverse as Facebook use [14], financial transactions [31], and editing in Wikipedia [39] has proven to be useful for generating insight about both communities and individuals. Analyses of IRC [18] and Facebook

[14] give glimpses into the regular heartbeat of community activity and evolution of these communities over time. For individuals, tools such as Themail [38] allow people to see both details of their email interactions and how they evolve over time, while Snoopy [31] helps people see regular patterns, or “habits”, in their spending data.

In general, when researchers develop tools for this sort of analysis, they focus on temporal patterns and regularities at large scale, as in [14][18]. Individuals, on the other hand, tend to browse to specific times or events and reminisce about them when given tools for visualizing personal archives of conversation, rather than look at patterns [37][38]. Our goal in this paper was to explore whether this focus on details was influenced by the design of these interfaces, or whether it is a more fundamental property of the way people use data to reflect on their pasts. Sociologists recognize time as a major factor in social matters and relationships [35], and our hypothesis was that focusing people’s attention on temporal rhythms of communication would help them learn more about themselves.

To this end, we developed pieTime (Fig. 1), a visualization that

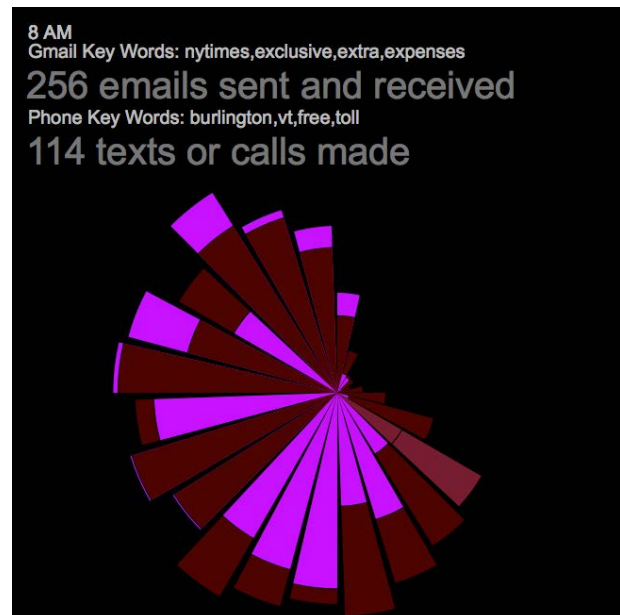


Figure 1. pieTime visualization comparing email and phone activity by hours of the day.

Copyright is held by the author/owner(s).

iConference 2012, February 7-10, 2012, Toronto, Ontario, Canada
Copyright © 2012 ACM 978-1-4503-0782-6/12/02.

helps people compare their behavior across communication media at scales from hours of the day to months of the year. It also chooses characteristic keywords associated with each time period and medium to help ground the data, but our hope was that by focusing on the aggregations rather than the details, pieTime would focus people’s attention on rhythmic patterns of activity. An evaluation of pieTime with 15 people shows that these rhythms have some value, and that people noticed both expected and surprising patterns. Many people, however, still preferred to focus on specific activities and incidents even though pieTime presents few details, suggesting that events and storytelling are a fundamental way that people think about the past. Indeed, pieTime was most successful when people believed that the patterns and keywords reflected their identity, a finding of general interest for visualizations of social media. Finally, although on balance people found the interface engaging, they expressed little desire for ongoing use of pieTime, which raises the question of how these systems might encourage ongoing use and reflection.

2. RELATED WORK

pieTime sits at the intersection of three main streams of work: reflective technologies, visualizations of communication activity, and temporal rhythms in computer-mediated communication. We discuss each in turn.

Reflective Technologies. Technology designed for reflecting and self-forming is an emerging area within HCI. Hallnäs and Redström argue for designs that use technology to encourage reflection rather than efficiency [16], a theme echoed by both designers and users of such technologies [33]. Lifelogging and reminiscing are two emerging application areas for reflection.

Studies have recognized the value of noticing patterns in behavior for fostering self-awareness, self-control, and self-organization [2][27]. *Lifelogging*, or capturing, analyzing, and visualizing daily behavior [19], is one strategy for finding these patterns. Assogba and Donath showed that habitual behavior can be used for self-expression and introspection through their blogging website, Myrococosm [1], which lets users share graphs of their personal statistics with others (e.g. how many people irritated me today). Users enjoyed the novel form of graphical communication and felt connected to their data, which gave them a personalized depiction of the “eccentricities of life” they experienced.

In contrast to the aggregation and pattern-mining that often characterizes lifelogging, *reminiscing* is about reflecting on specific past experiences, both as a therapeutic psychological activity [41] and a way to foster relationships and understand oneself [40]. Building on theories of episodic memory [6] and empirical studies of family reminiscing [26], a number of systems attempt to support this activity. For example, Pensieve motivates everyday reminiscence by emailing “memory triggers”—Flickr photos, tweets and blog posts from your past, and non-personalized text prompts [23]—to your inbox. SenseCam [32] uses a small camera worn around a user’s neck to take hundreds of photos throughout the day that create a “visual journal” that can be used to reflect on one’s behavior [20].

The availability of personal data is a key issue. Li points out the value of using data that is automatically collected, as it is easily obtainable and pre-existing, requiring minimal effort from users to reflect on the collected information [19]. For example, Mint.com [21] leverages financial transactions to facilitate personal budgeting; these data can then be used to visualize everyday habits in activity and travel, as well as the environmental impact of those habits [31].

Communication Visualization. Early visualizations of email focused on tools for managing mail flows (e.g., helping people organize and visualize conversation threads [29]). Recent work, however, has focused on supporting reflection on email activity and relationships. Themail showed columns of keywords associated with a contact along a timeline that might help users see the evolution of a relationship, although in practice people often used it to recall specific events [38]. PostHistory’s calendar-based interface for correspondence logs was also useful in identifying meaningful dates and events in the past, and its display of social networks among contacts prompted users to align the evolution of these relationships with their email archive [37]. Zooming out further, Perer et al. analyzed a 15-year email archive to study the evolution of long-term social relationships [25].

Visualizing group discussions is also a related, well-explored area, especially in the context of workgroups [3][8] and online communities [9]. Some of these visualizations are similar to pieTime in that they concentrate on revealing patterns (e.g., using concentric rings [4], spirals [5] and stacked timelines [10]). However, pieTime’s focus is not on the evolution of conversations, relationships, or workgroups, but the aggregation of communication patterns for individual, personal reflection. By giving people not only long-term views by month and year, but also views that divide activity down to the hour and by focusing on patterns across all of one’s contacts rather than individuals, we hope to help people see their life rhythms reflected at many scales.

Temporal Rhythms. These rhythms, and time more generally, are a major factor in social matters and relationships [35]. Researchers have explored these patterns at the community level, such as Jones et al.’s analysis of temporal patterns in IRC, including session durations and hourly message densities [18], and Golder et al.’s discovery of regular patterns across days and weeks in Facebook activity [14]. Temporal rhythms can also be used in concert with social network data, as with Fisher and Dourish’s visualizing of past interactions among email contacts to infer social roles and relationships [12] and Eagle and Pentland’s use of mobile phones as “wearable sensors” [11].

These technologies use a number of media beyond email to provide insight into behavior: IRC [18], Facebook [14], mobile phone use [11], email [12][25][29][37][38], instant messaging [35], spending data [31], blog entries [1], and so on. In general, these systems focus on one medium at a time.

3. SYSTEM DESIGN

pieTime’s goal is to help people reflect on patterns in their behavior in ways that complement existing approaches to visualization and reflection described above. As pieTime and our understanding of the design space evolved, we arrived at three main design goals that build on prior work:

Focus on aggregate patterns, not specific details. Most evaluations of the tools described above suggest that people focus on specific events in the past. Although reminiscence is useful, our goal was to explicitly focus attention on patterns to see what they might provide. Thus, we minimize the presentation of specific details that are known to support reminiscing to focus people’s attention on higher-level patterns—explicitly contrary to the “details on demand” portion of Shneiderman’s visual information seeking mantra [34].

Help people reflect on patterns at varying levels of granularity. This brings the work on community rhythms in IRC and Facebook

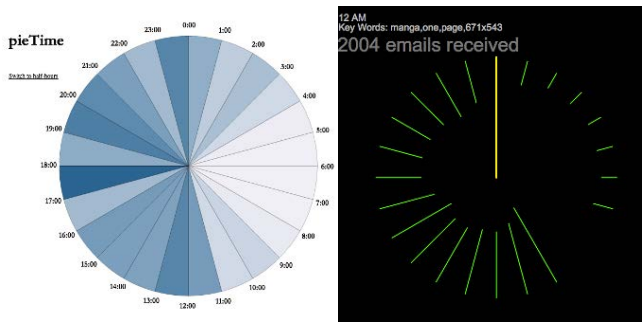


Figure 2. Early prototypes showing hourly activity levels. At left, version 1 shows phone activity indicated by color intensity. At right, version 2 shows email activity indicated by line length and keywords associated with the time interval.

down to the individual scale. Our hope is to help people see both familiar and unexpected patterns at a variety of timescales.

Support comparisons across media. People choose different media for different goals [13], and we hoped that comparing patterns across media might lead to insights. Multi-medium visualizations such as Intel’s Museum of Me [22] also engage users.

Below, we discuss how the major components of the design—the pie, the choice of data source, and the inclusion of keywords—evolved. Our choices and changes were driven by feedback on early iterations as well as by our own reactions and design goals.

3.1 The Pie

Timelines are a common way to visualize temporal data [10]. Spirals are another way to simultaneously display both serial and periodic data [5]. However, we were inspired by more traditional representations of time, such as the analog clock, that suggest repeated, cyclical patterns. These circular forms are familiar, and we used the clock metaphor as a starting point for developing the pie. The goal of the pie was to emphasize sections as parts of a whole—recurring periods, or meaningful “slices” of one’s life.

pieTime aggregates the amount of activity in each time slice by counting the number of timestamps corresponding to each incoming and outgoing communication. Our first prototype was a near-literal clock representation that associated color intensity with the amount of e-mail activity aggregated by hour of the day,

while our second represented activity with lengths of line segments, making it easier to compare activity levels at different time slices with the naked eye (Figure 2).

However, users found color intensities to be hard to compare and line segments hard to interact with (e.g., by hovering over with a mouse to see more details about the time period). The current version provides both an easy way to compare time periods while preserving the idea of “pieces of a whole” through time slices of varying radii (Figure 3). The radii are scaled so that the piece representing maximum activity for that timescale extends to the edge of the invisible circle that encloses the pie. Segment lengths in version 2 were computed the same way, but started from the edge of the pie rather than the center. Each pie slice consumes $1/n$ of the circumference, where n is the number of time divisions. Although these choices create a visual illusion that make periods with higher levels of activity look disproportionately larger in area than pieces displaying less activity, we hoped that increased usability and more aesthetic visuals would outweigh the need for precision in comparing exact levels of activity.

3.2 Timescales

Our initial thought was that since daily rhythms are some of the most basic ones we experience, we should focus on time divisions that give people insight into the structure of their day. Thus, the first prototype showed activity aggregated by either hour or half hour, with the pie divided into 24 or 48 pieces. This led to some expected regularities, such as mealtimes and sleep patterns—and this may not be bad. In recommender systems, suggesting movies that people already know they like can help to build trust in the system [36]. As we will see, people need cues to help them trust visualizations of conversation as well.

People reacting to the first prototype thought they would like to see patterns in changes in activity at larger timescales. Some expected that they were likely to exhibit the same behavior, for instance, after going to the gym on certain days of each week. In the final version of pieTime, we took this suggestion, as well as following prior work that divides time into months and years. Thus, pieTime allows users to segment the pie into 24, 31, 7, and 12 pieces representing hours of the day, days of the month, days of the week, and months of the year respectively (Figure 3). Hourly is the default (the half-hour timescale was removed since it provided little to no additional information over this). We hoped that these scales might each reveal, and prompt reflection on, regular events: daily coffee, weekly meetings, monthly club activities, or the way that fiscal years or seasonal weather all




Figure 3. Final version of pieTime showing aggregated email activity at different timescales.

shape people’s communications and lives. For example, in the phone data shown in Figure 2, pieTime showed an unexpected spike of activity for one of the authors between 11pm and 12am. He had been taking late night ice cream trips with a friend and had not realized how frequently they were occurring. pieTime was able to pick up on this behavior because it accumulates activity over time.

3.3 Data Sources

pieTime was designed to help people see patterns of activity in communication media. To choose appropriate media, we first brainstormed potential sources of communication histories, then narrowed our list based on three main criteria. First, the data had to be readily available and quick to download, because we wanted to be able to rapidly generate the visualization. This imposed constraints, such as using headers rather than the content of communication, that were not so onerous given our focus on temporal rhythms but might be unsuitable for other goals. Second, the sources we chose had to be widely used in order to find potential users. Third, the source needed to be a rich enough representation to suggest meaningful insights.

These criteria ruled out instant messaging, for instance; there are no dominant IM clients, and user interviews suggested that many people don’t store chat history. Blog posts are generally infrequent and are more about self-expression than interpersonal communication. Twitter has tight download limits. Facebook activity logs contain rich information, but are hard to access because of its frequently changing API and awkward privacy model. We even experimented with including web-browsing activity, but surveys suggested that people clean out their browsing history on a regular basis and many people use multiple browsers on different machines, making it hard to gather enough data to represent activity patterns.

In the end, we chose email and mobile phone usage. Email was a natural choice given its ubiquity and the amount of prior research that has used it, and although our informants described email as primarily a work medium, work is still an integral part of life. Cell phone usage is relevant because phone calls coupled with texting facilitate many everyday interactions among today’s youth [15][30]. Our system uses a preprocessor that downloads mail headers from Gmail and phone bills (including phone numbers, timestamps of text messages and calls, and locations) from either AT&T or Verizon.

Users of early prototypes expressed a strong interest in being able to compare their data sets. We first tried presenting two visualizations side-by-side, but users had trouble moving between them. Since concentric circles work well for visual comparisons [28], we used them to establish a common time scale and to extend the “pieces of a whole” metaphor by “stacking” the two data sets. We also extended the hovering feature to display the detailed data for both data sets. Figure 1 shows the stacking in action. We did no scaling to favor either email or phone data as some users may communicate more through either medium. Instead, we combined the raw number of phone calls with the number of text messages into “texts or calls made.” We ignored the length of communication because we were primarily concerned with frequency of communication rather than duration.

3.4 Keyword extraction

Although seeing spikes of activity in the first prototype of pieTime triggered informants to recall or construct reasons why these spikes existed, they sometimes had trouble explaining them.

To help people create explanations and to create the personal connections between the user and the visualization recommended by [37], we decided to extract representative keywords for each time slice. For emails we focused on subject lines; for phone bills we used location data about SMS messages and calls, as this was the only qualitative data we could retrieve.

Fisher linear classifiers exploit probabilities of word occurrence to categorize email as either spam or legitimate. This seemed like a plausible approach for finding distinctive keywords, so for each time slice, timescale, and medium, we trained a Fisher classifier that treated content from that medium and time slice as “ham” and everything else as “spam”. We then returned the four words that had the highest relative probability of appearing in that medium and time slice. For example, in Figure 1, the words nytimes and burlington are most uniquely identified with emails and phone activity between 8am and 9am. Four words is a somewhat arbitrary choice that happened to fit well in the allotted space.

4. EVALUATION

We used the version of pieTime shown in Figures 1 and 3 to evaluate how well it accomplished its primary goal of supporting reflection on patterns of activity, as well as how our design decisions to avoid detail, present data at multiple levels of granularity, and support comparison across media affected the ways people reflected.

We recruited 15 college-aged participants (3 female, 12 male) from a large northeastern university. Each participant was required to have access to electronic AT&T or Verizon phone bills and a Gmail account. The advertisement for the experiment made it clear that none of their data would be kept and that the researchers would only see the visualization if the participant gave permission. Participants were compensated with either course credit or a piece of pie¹.

Each session was conducted by a pair of experimenters. After obtaining consent, we told participants that the goal was to “evaluate a tool that visualizes people’s communication patterns.” They were then asked to enter their Gmail username and password so our preprocessor could prepare their email. This processing took about five minutes, as many participants had thousands of emails. During this processing, we described the experimental procedure and helped them download their phone bills.

We first asked participants to “play around” with the tool while thinking aloud. One experimenter observed while the other used a list of questions to guide participants and help them notice different parts of the system. We probed participants to describe what they thought they were seeing, what they were thinking as they used pieTime, whether any of the data or their interpretations of it surprised them, and whether they felt any connection between what they saw and what has happened in their life.

After participants finished exploring the interface, we asked a series of follow-up questions about their reactions to the system as a whole. These include questions about the usability of the visualization, which features were interesting, useful, or confusing, prior experience with similar tools or visualizations, and so on. Sessions averaged approximately 45 minutes.

To analyze the data, we created affinity diagrams based on our observations in order to find major themes that shaped

¹ Seven chose pie, eight chose credit. Apparently, hunger and grades are equally important to students.

participants' experience and their reactions to features. Quotes below are attributed with that participant's id number (e.g., P#).

5. RESULTS

Our data suggested four main themes. First, although our goal was to focus attention on patterns, people first looked to the keywords to orient themselves and to convince themselves that pieTime reflected their identity. Second, they looked for anomalies, both in the keywords and in levels of activity, and then constructed explanations. Third, people detected and reflected on activity patterns pieTime revealed, both expected and unexpected. Fourth, although the interface for comparing email and phone activity was less usable than we hoped, presenting that information did lead people to reflect on their use of those media.

5.1 Personal connection was crucial

Despite our goal of focusing attention on higher-level patterns, people tended to spend much more time analyzing keywords than noticing patterns. This may be because in the current design, there is more information in the keywords than the frequencies: as one user (P2) pointed out, switching from one hour to the next would produce four new keywords (*livermore_ca, sunnyvale_ca, newyork_ny, camden_nj*), for phone data), whereas only one frequency number (*80 texts or calls made*) would be updated.

The keywords, however, were often difficult for people to interpret. The Fisher algorithm selects the *most identifying* keywords for a given time slice, relative to other times, and these are often cryptic (as with *671x543* in Figure 2). Many users had a hard time connecting these keywords to their own lives, especially the *nonsensical [ones] bloated by spam* (P14). A few even wondered if we were showing them someone else's data, because their intuition was that the keywords should be the most commonly occurring words:

"I email my mother a lot...I would really expect to see my mom's keywords because that's really important to me." (P11)

"I'm in a frat so I should get a lot of frat emails." (P1)

Thus, activity by itself didn't mean much unless it was associated with recognizable keywords. Recognition led to connection, as with this user who was browsing through a view sorted by days of the month, then stopped in delight at a promising discovery:

"Oh, the 27th day of the month shows up with my girlfriend's last name." (P5)

Suddenly, this data was his. The keyword confirmed his connection to the data as he tried to uncover the relation between the date and his previous correspondences. These findings resonated with similar notions of personalization found in PostHistory [37] and Themail [38]. Users expected the dataset to represent them as individuals, though one said his *"name shouldn't show up as a keyword"* (P12).

5.2 Sensemaking of the Past, Especially the Unusual

Participants were also quick to detect anomalies in the data. These prompted users to search for specific events such as vacations or projects to explain bursts of activity or unusual keywords. All participants linked keywords and patterns in the visualization to life events. For instance, high volumes of calls and texts around certain days of a month were attributed to birthdays, holidays, or travel periods that required more on-the-go coordination: *"During the holidays I texted and called a lot. I was traveling and coordinating with people."* (P7)

A handful of users were especially intrigued by the keywords, scanning through each time slice's set of words and creating stories to explain why the words may have appeared. They were able to jump through a series of quick anecdotes while mumbling through each keyword:

"Evina: some health care thing, I forget what for...Nxpowlite: something at [a company I worked for] people use to reduce jpeg file sizes to make Powerpoints smaller...EZ-pass: at the end of November, my dad pays for EZ-pass, cause I drove home for Thanksgiving and there's a lot of tolls on the way back and he wanted to make sure he had enough money in our account." (P3)

"Freetoll: that's from me trying to figure out many minutes I have left...[Laughs] June keywords: "burning" and "exposed" all have to do with me getting sunburned—I peeled twice." (P8)

Somewhat like the nagging feeling of being unable to remember the answer to a trivia question, participants were annoyed or puzzled when they couldn't explain what they saw: *"I don't think I've ever called Toronto at 3pm"* (P5). One user (P2) kept returning to a mysterious spike in emails and obscure words at 3am, until he finally sighed in relief, realizing it was a daily mailing list for dictionary and SAT words.

In general, users enjoyed finding bits of their past to discover. Whether they were simply impressed that the system was able to unearth these fine details, or just interested in getting a glance of moments in the past, everyone had a story to tell. Further, their focus on specific details and stories echoes findings from other evaluations of visualizations [37][38], even though our design avoided such details in order to focus attention on patterns.

5.3 Finding Patterns in Behavior

However, reactions were mixed about the utility of patterns. Some participants claimed they didn't learn anything that they didn't already know after using the tool, but others made observations about their lives they hadn't explicitly considered before:

"7pm is tiny—this is when I eat. That's interesting. Same thing at 11:30 and noon...And throughout the entire morning. I'm not a morning person." (P7)

"Apparently I text everybody in the world at 4pm...Yeah that sounds about right. It's when I get out of work. It's also right before I go to workouts." (P4)

People's reactions to the timescales varied. The default hourly breakdown gave users a direct, *"nice and well-labeled"* (P15) way of looking at their daily patterns, though one described it as *"too granular"* (P11) to reflect anything in his randomly ordered schedule. Most users spent their time looking at hourly, weekly, and monthly views, though a few spent so much time interpreting each hour of the default timescale that they had little time to explore the others. Users spent the least time on days of the month, which had few discernable patterns. Viewing by months of the year and days of the week tended to make the most sense to people. Nine users observed that months of the year revealed cycles of the academic year, noting distinct drop-offs during breaks and obvious growth during classes and important events like interviews and trips. All but one user noticed a significantly higher amount of emails sent and received from Sunday through Tuesday compared to the rest of the week, and half of them found the discovery surprising. After considering possible causes, most users devised a likely explanation:

"Friday is the day I don't wanna do any work. Saturday is not as busy 'cause you're waiting for Sunday. Thursday—early partying..." (P2)

Ten participants were able to link patterns in their lives to times, using the keywords associated with each time:

“It looks like 6:30 is a contact California [half] hour, whereas 6pm appears to be New Jersey. It’s just interesting that there’s a time of day that seems to be dominated by Californians.” (P7)

And although in the post-interviews, most participants claimed that using a tool such as pieTime would not change anything in their lives, seeing the visualizations still prompted some to reflect on their habits and lifestyles:

“The keywords at 4am—guess that tells me I’ve been up too late working on that project. Definitely the most useful piece of info I’ve seen so far. Working on a project at 4am in the morning is kinda bad.” (P5)

“Maybe I use email too much during the week.” (P4)

One user extended his observations in a pensive manner, reflecting on how a hiatus in email communication with others made him feel:

“I do tend to notice that I get a lot less [e-]mail in the summer. [It] makes me sound pathetic. I feel more connected getting emails. I don’t get even much personal correspondence through email, but it’s depressing not to get any email at all even if I just delete it.” (P8)

Communication is in an integral part of people’s lives. People have an idea of how much correspondence they feel is too little or too much, and providing a way to explore these data might help people become more aware of time-based patterns. Further, these patterns don’t have to be unexpected in order to deliver a message that prompts reflection.

5.4 Comparing Email and Phone

In general, users found phone keywords, which were simply the locations of the other caller, to be more meaningful than keywords from the email data set. These locations had a much more limited range than the keywords and were often less idiosyncratic, allowing users to associate activity with specific friends and family with little ambiguity.

Many users found stacking email and phone data together to be unintuitive because of multiple issues, including normalization of activity levels and color. Still others felt it was too much information to look at all at once, though some were interested in the patterns it revealed:

“Obviously I get a lot more emails than phone calls. I’m just laughing at how small phone is compared to email. It’s almost opposite [on the weekend], because people like to do things on Thursday/Friday and they’re not emailing me as much, and we’re trying to do stuff.” (P8)

Some users were able to find congruencies between their phone and email:

“Emailing and phone drops right about 3am and comes back at about 7. They’re sort of similar.” (P5)

Another user linked a location from the list of phone keywords to a friend’s name from the email keywords, both occurring within the same time interval:

“I’m seeing Texas and [a friend’s name] together, which is cool. I was seeing [him] with a friend from Texas. This was happening at the same time. Interesting.” (P7)

Participants also noted that media served different purposes in their lives, with email primarily pertaining to work, while phone contacts were more social: *“I talk to people mostly on*

Fridays...because I try to go out.” (P10) *“I personally don’t send any emails usually, except to teachers.”* (P6) *“I use [email] for school. I prefer to call if I have something...If I have something short to say I would just text.”* (P2)

5.5 Overall Reactions

All users found the tool interesting—a *“cool new concept”* (P1). Only two participants had looked at a similar tool in the past or regularly checked their phone logs, while nine users *“didn’t care about this stuff before”* (P2) or *“hadn’t thought about it”* (P7). Most users said they enjoyed playing with the visualization and looking at the pretty shapes and colors. Four were deeply curious about their statistics down to each hour, with one claiming he could *“stare at this all day”* (P11).

Users suggested a number of more interactive features such as the ability to zoom into a specific date or time, or a play button that animates the accumulation of communication activity over time as a way to give a glimpse into how temporal patterns form. One user’s suggestion of a geographic map displaying origins of phone calls aligns with work around SenseCam [32] and Google’s MyMaps [24] showing that location data can trigger memories.

Five users saw some use for the tool, ranging from repeated use *“once a week”* (P1) to more practical goals around self-monitoring. One user found his phone usage *“alarming”* (P5) after realizing it was roughly even throughout the day, citing it as a good reminder to watch his use of limited free daytime minutes before 9pm. Another user (P8) wanted to cancel his texting services after seeing how little he actually texted each month while paying for an unlimited texting plan. Three users claimed to see no practical application for the tool (*“don’t need it to survive”* (P6)), while seven others said it wasn’t useful but *“good for reminiscing”* (P3) or playing with out of curiosity.

6. DISCUSSION

Our experience suggests that, although pieTime was interesting and supported reflection, it did not accomplish its primary goal of focusing attention on patterns of behavior. Here, we would like to take a step back from pieTime and talk about why this is, and how our results apply to the larger collage of tools for reflection on communication and behavior.

Patterns complement particulars. Our work gives further evidence to the importance of helping people reflect on particular events as a way to understand their pasts. We had hoped to focus attention on rhythms by downplaying individual events in our design in order to encourage people to process the past in the aggregate and in the large. However, as with other such tools (e.g., [37][38]), people often used the data to reflect on specific experiences. This suggests that people use specific incidents and events as a fundamental way of reminiscing, reflection, and understanding the past, and that even if the goal is seeing a “bigger picture”, details are important [34].

The question then becomes how to encourage people to look for patterns. Systems that provide a balance of aggregate and detailed behavioral information might be a sweet spot in this design space. Consider a lifelogging tool that supports exercise. Patterns over time could help people see a steady increase in activity, while providing access to specific moments such as a particularly long or fast run, or the time they broke through a milestone, might be useful for reflection and motivation.

Support storytelling. pieTime users were happier when they could make personal connections with the data. Users were interested in the stories that they could tell from seeing the keywords and

patterns provided by pieTime, even if the data did not always immediately make sense. These snippets of the past motivate people to tell stories based on what they believe the snippets are about, and though theories of autobiographical memory suggest that these stories are not perfect recallings of the exact event [6], these stories may help people more deeply understand themselves. Further, people commonly wish they captured more meaningful information about their past [7][26].

This suggests that visualization and lifelogging tools should encourage interpretation and storytelling about the past, to complement the system's strengths around capture and representation. For instance, a visualization of Facebook data might allow people to pull together a variety of media and activities to tell stories about their relationships—a manually curated version of its existing “See Friendship” feature.

Personalization and identity matter. We call out identity management because another main outcome of our study was that people were much more convinced about, and engaged with, the system when they felt that it reflected their identity. Users felt proud of keywords or trends they thought characterized them: “*You can tell I’m a geek*” (P11). Evidence that assured people that data they saw was their own was satisfying, and helped them tell stories about who they were:

“*Yeah, [the visualization] seems pretty accurate to me, especially just the time of day. Even though I get up early in the morning, I try not to send e-mails that early cause I don’t want people to think I’m crazy. I receive e-mails from other crazies at 2am and 1am. That’s so not me.*” (P10)

This idea of seeing data that reflects one’s identity parallels the idea of public visualization of behavior as a tool for creating public identity or “social capital” [37]. Such values of finding “personal relevance” and “identity presentation” within visualizations have found support in collaborative visualization systems as well [17], while identity creation, maintenance, and management are a fundamental psychological concern.

Rhythms do support awareness. That said, pieTime users did find value in awareness of their rhythms, noticing both mundane and unexpected patterns. While our study did not directly compare the use (or lack of use) of other data display tools with pieTime, some data suggests the value of revealing rhythms that might seem obvious—as with the user who regularly checked his phone bill, yet found motivation to change his service plan when he saw his usage habits visualized, rather than as a list of numbers. Awareness is valuable in various realms, including improvement of personal and mental health. Supporting awareness of rhythms and routines, for instance, can improve the life of the elderly [28]. These patterns, as with habit-finding [31] and lifelogging [19] systems, can support understanding and reflection.

Thus, designers of tools for reflection should consider features that help people see these aggregated rhythms while still having access to the particulars that appear to be necessary. One strategy would be to allow filtering by communication partner, by time period, topic, or location, to allow people to see aggregations that might be more meaningful to them. Another would be to use a spiral layout (e.g., [5]) that gives access to both high-level and detail views, although in pieTime’s case that went against the aesthetic and conceptual goals of presenting parts of a whole.

Limitations. In retrospect, filtering of spam and mailing lists would have made rhythms and keywords more apparent and meaningful; further, we did not attempt to manage timezone and character encoding differences. Better handling of these issues

might have improved the experience. Our participants were also relatively homogeneous. Our biggest limitation, however, is that we used a short term, lab-based deployment to evaluate pieTime. We did this consciously, to protect people’s privacy, to guard against bugs, and to respond to limitations in our resources for conducting experiments. This limitation is not unique to our work. Almost all of the tools for visualizing conversation are evaluated in similar short-term, one-time use scenarios, and given our participants’ general skepticism toward the everyday utility of these tools, such one-shot evaluations might overstate their value. It is an open and interesting question how these kinds of tools might function in the long term to support repeated use and continued reflection.

7. CONCLUSION

We designed pieTime to visualize personal communication data in an effort to reveal temporal rhythms in behavior and encourage reflection. It extends prior work around reflecting on conversation in several ways: focusing on aggregate behavior, providing views at multiple timescales, using multiple data sources, and choosing phone logs as an unusual, but revealing, source of behavioral data. Putting our results into the context of related prior efforts suggests that themes around personalization, reminiscence, detail, explanation, and storytelling are fundamental to how people think about their pasts, and that these aspects and activities need to be supported across a wide range of designs and goals.

Our hope is that future tools that support reflection on personal communication archives go beyond single-use evaluations and a focus toward the past. Instead, they could use people’s patterns and stories to help people bring the present into focus, and to encourage ongoing reflection. Adding a cyclical element to the visualization of archived data is one way to help people to not only reflect on past events, or relationships, but also think about their behaviors and relationships over time: from their past, into their present, and perhaps with consequences for their future. By making tools about the past relevant to people’s present lives, we hope we can create a world where visualizations that support reflection and self-understanding are less like children’s toys that are played with once and forgotten, and instead become artifacts that integrate into our lives, much like the clocks that we hang on our walls.

8. ACKNOWLEDGMENTS

We would like to acknowledge Erica Horowitz and Mitchell Davis for their work in helping to evaluate pieTime, and the NSF for support through grant IIS 08-35451.

9. REFERENCES

- [1] Assogba, Y. and Donath, J. 2009. Myrococosm: Visual microblogging. *Proc. HICSS*.
- [2] Barton, S. 1994. Chaos, self-organization, and psychology. *The American Psychologist*, 49, 5–14.
- [3] Begole, J., Tang, J., Smith, R., and Yankelovich, N. 2002. Work rhythms: Analyzing visualizations of awareness histories of distributed groups. *Proc. CSCW*, 334–343.
- [4] Bergstrom, T. and Karahalios, K. 2007. Seeing more: visualizing audio cues. *Proc. INTERACT*, 29–42.
- [5] Carlis, J. and Konstan, J. 1998. Interactive visualization of serial periodic data. *Proc. UIST*, 29–38.
- [6] Conway, M. A. and Pleydell-Pearce, C. W. 2000. The construction of autobiographical memories in the self-memory system. *Psychol. Rev.*, 107(2): 261–288.

- [7] Cosley, D., Akey, K., Alson, B., Baxter, J., Broomfield, M., Lee, S., and Sarabu, C. 2009. Using Technologies to Support Reminiscence. *BCS HCI*, 480–484.
- [8] DiMicco, J., Hollenbach, K., and Bender, W. 2006. Using visualizations to review a group’s interaction dynamics. *Proc. CHI*, 706–711.
- [9] Donath, J., Viégas, F., and Karahalios, K. 1999. Visualizing conversation. *Proc. HICSS*, 2023–2031.
- [10] Dörk, M., Gruen, D., Williamson, C., and Carpendale S. 2010. A visual backchannel for large-scale events. *Proc. InfoVis*, (16)6, 1129–1138.
- [11] Eagle, N. and Pentland, A. 2006. Reality mining: sensing complex social systems. *Pers. Ubiqu. Comp.*, 255–268.
- [12] Fisher, D. and Dourish, P. 2004. Social and temporal structures in everyday collaboration. *Proc. CHI*, 551–558.
- [13] Flanagin, A. and Metzger, M. 2001. Internet use in the contemporary media environment. *Human Comm. Research*, 27(1), 153–181.
- [14] Golder, S., Wilkinson, D., and Huberman, B. A. 2007. Rhythms of social interaction: Messaging within a massive online network. *Proc. C&T*, 41–66.
- [15] Grinter, R. E. and Eldridge, M. A. 2003. Wan2tlk?: Everyday text messaging. *Proc. CHI*, 441–448.
- [16] Hallnäs, L. and Redström, J. 2001. Slow Technology: Designing for reflection. *Pers. Ubiqu. Comp.*, 5(3), 201–212.
- [17] Heer, J. and Agrawala, M. 2008. Design considerations for collaborative visual analytics. *Proc. InfoVis*, 171–178.
- [18] Jones, Q., Moldovan, M., Raban, D., and Butler, B. 2008. Empirical evidence of information overload constraining chat channel community interactions. *Proc. CSCW*, 323–332.
- [19] Li, I., Dey, A. K., and Forlizzi, J. F. 2010. A stage-based model of personal informatics. *Proc. CHI*, 1549–1552.
- [20] Lindley, S. E., Randall, D., Harper, R., Glancy, M., and Smyth, N. 2009. Reflecting on oneself and on others: Multiple perspectives via SenseCam. *CHI workshop on Designing for Reflection on Experience*.
- [21] Mint. <http://www.mint.com>
- [22] Museum of Me. <http://www.intel.com/museumofme>
- [23] Peesapati, S. T., Schwanda, V., Schultz, J., Lepage, M., Jeong, S., and Cosley, D. 2010. Pensieve: Supporting everyday reminiscence. *Proc. CHI*, 2027–2036.
- [24] Peesapati, S. T., Schwanda, V., Schultz, J., and Cosley, D. 2010. Triggering memories with online maps. *Proc. ASIST*, 1–4.
- [25] Perer, A., Shneiderman, B., and Oard, D. W. 2005. Using rhythms of relationships to understand email archives. *JASIST*, 57, 1936–1948.
- [26] Petrelli, D., Whittaker, S., and Brockmeier, J. 2008. AutoTopography: What can physical mementos tell us about digital memories? *Proc. CHI*, 53–62.
- [27] Rachlin, H. 1995. The value of temporal patterns in behavior. *Curr. Directions in Psych. Sci.*, 4(6), 188–192.
- [28] Riche, Y. and Mackay, W. 2010. PeerCare: supporting awareness of rhythms and routines for better aging in place. *JCSCW*, 19(1), 73–104.
- [29] Samiei, M., Dill, J., and Kirkpatrick, A. 2004. EzMail: Using information visualization techniques to help manage email. *Proc. InfoVis*, 477–482.
- [30] Schiano, D. J., Chen, C. P., Isaacs, E., Ginsberg, J., Gretarsdottir, U., and Huddleston, M. 2002. Teen use of messaging media. *Proc. CHI*, 594–595.
- [31] Schwarz, J., Mankoff, J., and Matthews, H. S. 2009. Reflections of everyday activities in spending data. *Proc. CHI*, 1737–1740.
- [32] Sellen, A., Fogg, A., Aitken, M., Hodges, S., Rother, C., and Wood, K. 2007. Do life-logging technologies support memory for the past? An experimental study using SenseCam. *Proc. CHI*, 81–90.
- [33] Sengers, P., Boehner, K., David, S., and Kaye, J. 2005. Reflective design. *Critical Computing*, 49–58.
- [34] Shneiderman, B. 1996. The eyes have it: a task by data type taxonomy of information visualizations. *Proc. IEEE Symposium on Visual Languages '96*, 336–343.
- [35] Sorokin, P. and Merton, R. 1937. Social time: a methodological and functional analysis. *American Journal of Sociology*, 42(5), 615–629.
- [36] Swearingen, K. and Sinha, R. 2002. Interaction design for recommender systems. *Proc. DIS*.
- [37] Viégas, F., boyd, d., Nguyen, D. H., Potter, J., and Donath, J. 2004. Digital artifacts for remembering and storytelling: PostHistory and social network fragments. *Proc. HICSS*.
- [38] Viégas, F., Golder, S., and Donath, J. 2006. Visualizing email content: Portraying relationships from conversational histories. *Proc. CHI*, 979–988.
- [39] Viégas, F., Wattenberg, M., and Dave, K. 2004. Studying cooperation and conflict between authors with history flow visualizations. *Proc. CHI*, 575–582.
- [40] Webster, J. D. and McCall, M. E. 1999. Reminiscence functions across adulthood: A replication and extension. *J Adult Dev.* 6(1), 73–85.
- [41] Wong, P. T. B. 1995. The process of adaptive reminiscence. In *The art and science of reminiscing: Theory, research, methods and application*, 23–35. Taylor & Francis.