AFRL-OSR-VA-TR-2013-0131

Networks of Memories

Simon Dennis, Mikhail Belkin

Ohio State University

March 2013
Final Report

DISTRIBUTION A: Approved for public release.

AIR FORCE RESEARCH LABORATORY
AF OFFICE OF SCIENTIFIC RESEARCH (AFOSR)
ARLINGTON, VIRGINIA 22203
AIR FORCE MATERIEL COMMAND

REPORT DOCUMENTATION PAGE	*Form Approved* *OMB No. 0704-0188*

1. REPORT DATE *(DD-MM-YYYY)*	**2. REPORT TYPE**	**3. DATES COVERED** *(From - To)*
15-02-2013	FINAL	14-06-2011 - 13-06-2012

4. TITLE AND SUBTITLE
Networks of Memories

5a. CONTRACT NUMBER

5b. GRANT NUMBER
FA9550-09-1-0614

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Simon Dennis
Mikhail Belkin

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Ohio State University
1835 Neil Avenue,
Columbus, OH, 43210

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
AFOSR
875 N Randolph St
Arlington, VA 22203

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
AFRL-OSR-VA-TR-2013-0131

12. DISTRIBUTION/AVAILABILITY STATEMENT
DISTRIBUTION STATEMENT A: Approved for public release

13. SUPPLEMENTARY NOTES

14. ABSTRACT
People naturally divide their everyday experience into a sequence of events and use these representations to organize perception, memory and communication (Zacks & Tversky, 2001; Zacks, Tversky & Iyer, 2001). In this project, we have created a platform for collecting lifelogging data to study context, characterized the distributional structure of context in the real world, empirically investigated people's ability to isolate when events occurred and developed algorithms capable of automatically segmenting and tagging lifelog data.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			Dr. Jay Myung
U	U	U	U		**19b. TELEPHONE NUMBER** *(Include area code)* 703-696-8487

Reset

INSTRUCTIONS FOR COMPLETING SF 298

1. REPORT DATE. Full publication date, including day, month, if available. Must cite at least the year and be Year 2000 compliant, e.g. 30-06-1998; xx-06-1998; xx-xx-1998.

2. REPORT TYPE. State the type of report, such as final, technical, interim, memorandum, master's thesis, progress, quarterly, research, special, group study, etc.

3. DATES COVERED. Indicate the time during which the work was performed and the report was written, e.g., Jun 1997 - Jun 1998; 1-10 Jun 1996; May - Nov 1998; Nov 1998.

4. TITLE. Enter title and subtitle with volume number and part number, if applicable. On classified documents, enter the title classification in parentheses.

5a. CONTRACT NUMBER. Enter all contract numbers as they appear in the report, e.g. F33615-86-C-5169.

5b. GRANT NUMBER. Enter all grant numbers as they appear in the report, e.g. AFOSR-82-1234.

5c. PROGRAM ELEMENT NUMBER. Enter all program element numbers as they appear in the report, e.g. 61101A.

5d. PROJECT NUMBER. Enter all project numbers as they appear in the report, e.g. 1F665702D1257; ILIR.

5e. TASK NUMBER. Enter all task numbers as they appear in the report, e.g. 05; RF0330201; T4112.

5f. WORK UNIT NUMBER. Enter all work unit numbers as they appear in the report, e.g. 001; AFAPL30480105.

6. AUTHOR(S). Enter name(s) of person(s) responsible for writing the report, performing the research, or credited with the content of the report. The form of entry is the last name, first name, middle initial, and additional qualifiers separated by commas, e.g. Smith, Richard, J, Jr.

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES). Self-explanatory.

8. PERFORMING ORGANIZATION REPORT NUMBER. Enter all unique alphanumeric report numbers assigned by the performing organization, e.g. BRL-1234; AFWL-TR-85-4017-Vol-21-PT-2.

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES). Enter the name and address of the organization(s) financially responsible for and monitoring the work.

10. SPONSOR/MONITOR'S ACRONYM(S). Enter, if available, e.g. BRL, ARDEC, NADC.

11. SPONSOR/MONITOR'S REPORT NUMBER(S). Enter report number as assigned by the sponsoring/ monitoring agency, if available, e.g. BRL-TR-829; -215.

12. DISTRIBUTION/AVAILABILITY STATEMENT. Use agency-mandated availability statements to indicate the public availability or distribution limitations of the report. If additional limitations/ restrictions or special markings are indicated, follow agency authorization procedures, e.g. RD/FRD, PROPIN, ITAR, etc. Include copyright information.

13. SUPPLEMENTARY NOTES. Enter information not included elsewhere such as: prepared in cooperation with; translation of; report supersedes; old edition number, etc.

14. ABSTRACT. A brief (approximately 200 words) factual summary of the most significant information.

15. SUBJECT TERMS. Key words or phrases identifying major concepts in the report.

16. SECURITY CLASSIFICATION. Enter security classification in accordance with security classification regulations, e.g. U, C, S, etc. If this form contains classified information, stamp classification level on the top and bottom of this page.

17. LIMITATION OF ABSTRACT. This block must be completed to assign a distribution limitation to the abstract. Enter UU (Unclassified Unlimited) or SAR (Same as Report). An entry in this block is necessary if the abstract is to be limited.

Networks of Memories

FA9550-09-1-0614
Professor Jay Myung

PI: Simon Dennis
Ohio State University
February 15, 2013

Introduction

People naturally divide their everyday experience into a sequence of events and use these representations to organize perception, memory and communication (Zacks & Tversky, 2001; Zacks, Tversky & Iyer, 2001). Even under passive viewing conditions, neural data suggests that people do not perceive time in a continuous stream, but rather spontaneously parse their experience into distinct context representations (Zacks, Braver, Sheridan, Donaldson, Snyder, Ollinger, Buckner, & Raichle, 2001).

Episodic memory refers to the ability to bind item representations to these context representations and subsequently retrieve those bindings (Humphreys, Wiles & Dennis, 1994, Tulving, 1972, 2002[1]). Although context is definitional for the study of episodic memory at this point there is no theory of context. Contexts are typically operationally defined by referring to a study list, aspects of the experimental task or the physical attributes of the laboratory environment (Johnson, Hashtroudi & Lindsay, 1993; Smith & Vela, 2001). However, it is unclear to what extent the contexts used in the laboratory resemble those that people typically employ outside the laboratory (c.f. Conway & Pleydell-Pearce, 2000).

The lack of a theory of context is brought into stark relief when one considers work showing the importance of context noise in paradigms such as single item recognition (Dennis & Humphreys, 2001). In a series of studies, Dennis and colleagues (Dennis & Humphreys, 2001, Dennis, Lee & Kinnell, 2008, Dennis & Chapman, 2009, Kinnell & Dennis, 2011; Kinnell & Dennis, 2012) have shown that interference in recognition paradigms is likely to arise from the occurrence of items in pre-experimental contexts. Consequently, any substantive progress in our understanding of episodic memory awaits a better understanding of episodes in the wild.

Although it has long been argued that memory research that is focused solely on laboratory work is futile (Neisser, 1976), the difficulty has been how to proceed when the experience of the participant before they enter the laboratory cannot be rigorously quantified. One approach is to look for generic proxies to an individual's experience. For example, Anderson and Schooler (1991) conducted analyses on newspaper headlines, corpora of child speech and emails. They observed a remarkable correspondence between the patterns of recurrence in the data and the form of memory retention and practice curves collected in the laboratory. However, these methods require one to make an inference about the individual's experience on the basis of the experience of others. They are suitable for use in discovering strong trends, but given the considerable individual variability in people's life experience, their resolving power is necessarily limited.

Another approach is to have people keep personal diaries or to elicit memories from family members (e.g. Loftus & Pickrell, 1995). While useful for the purposes to which they have been applied, these methods have limitations. The protocols that are collected are products of the memory system - either of the individual being tested or family and friends and so do not represent an objective record of events. As a consequence, their veracity cannot be confirmed and more critically they are selective in nature. They are not comprehensive and cannot be used to characterize the entire experience of the participant.

Today, however, technology provides us with entirely new options. Easy to carry and able to monitor multiple sensor streams, smartphones can provide a convenient and ubiquitous window into the contexts of daily life. In this project, we have conducted psychological work that uses this data to develop a theory of context as well as machine learning work that builds on the psychological insights to create create algorithms capable of automatically segmenting and

[1] In this context, we are interested in the informational requirements of episodic memory, not the neuroanatomical hypothesis or the relationship to consciousness, which was added to the original concept later.

tagging naturally occurring contexts.

Objectives of Research

a) To create a platform for collecting lifelogging data.
b) To characterize the distributional structure of context in the real world.
c) To empirically investigate people's ability to isolate when events occurred.
d) To develop algorithms capable of automatically segmenting and tagging lifelog data.

Background and Technical Approach

a) To create a platform for collecting lifelogging data

In the course of this project, we have built a system which consists of an Android app, server infrastructure and user interfaces. The app continuously acquires data, including vision, audio (short sub-second snippets to preserve privacy), location and motion. Users wear the phone around their neck to allow an unobstructed view for the camera and the data is sent automatically to a secure server once a day. The user reviews each day's data and provides context boundaries, descriptions and labels, with the option to delete private portions.

b) To characterize the distributional structure of context in the real world.

Building a theory of context requires an understanding of the nature of episodic experience outside the laboratory. An initial concern might be that people's understanding of what constitutes a context might be so variable as to render forming generalizations impossible. It is certainly the case that people are able to conceive of contexts at different levels of abstraction (Zacks & Tversky, 2001). However, it also seems to be the case that there is a basic level of the event hierarchy that subjects naturally assume is the appropriate one to employ in our studies, much as has been argued to be the case for object categories (Rosch, 1978). Subjects in our experiments do not mark boundaries as consistently as is typically the case in laboratory studies of event segmentation (c.f. Newtson,1976; Speer, Swallow, & Zacks, 2003), but it does not require extensive instruction for subjects to understand what is required and between subject segmentation F scores are around 0.57 which indicates moderate agreement (see Zhuang, Belkin, & Dennis, 2012, for a description of how segmentation agreement is calculated). It seems then that there is a notion of context or event in real world situations that participants can employ reliably.

A number of interesting regularities are evident given the context boundaries, labels and descriptions that our participants have provided. Figure 1 shows histograms of the durations of contexts plotted on log and log-log coordinates. Short durations are more probable and when plotted on log-log axes the function is approximately linear, suggesting that the distribution conforms to a power law (although we are aware that demonstrating this rigorously is nontrivial, c.f. Lee, 2004). The distribution bears a striking resemblance to the pattern found by Anderson and Schooler (1991) when they examined the time between occurrences of words or sources in newspaper headlines, child speech and email - a pattern that they point out resembles the retention function found in laboratory studies of human memory. Finding a similar pattern in real world context durations supports the idea that the retention function is a result of contextual overlap.

The group data does not appear to conform to an exponential distribution (see log plot Figure 1), which suggests that the generating process is not first order Markovian and places constraints on the kinds of machine learning algorithms that can be entertained to predict context boundaries and labels (see Aim D, below). However, more data is required as it is important that the power law pattern is seen at both the group and individual levels. Heathcote, Brown & Mewhort (2000) have demonstrated in the case of the power law of practice, that averaging over exponential patterns can produce a curve that mimics a power law, and the same may be occurring in this case.

We can also look at the distributions of context types by examining the labels that participants provide to describe each context. Figure 2 shows the result of a content analysis of these labels. Activity is the dominant way in which people characterize contexts, with places, day of week, emotion/states, people and objects also contributing significantly to the context concept.

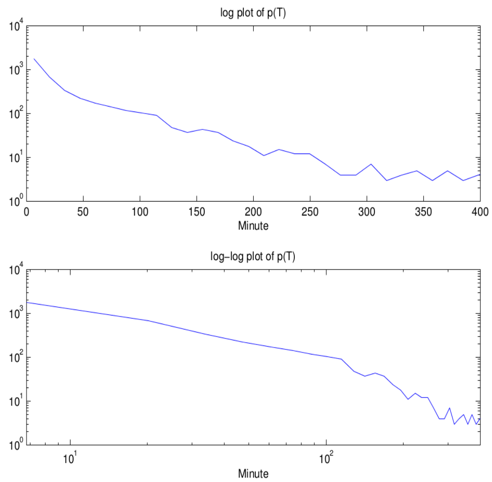

Figure 1: Log and Log-log plot of the histogram of context durations.

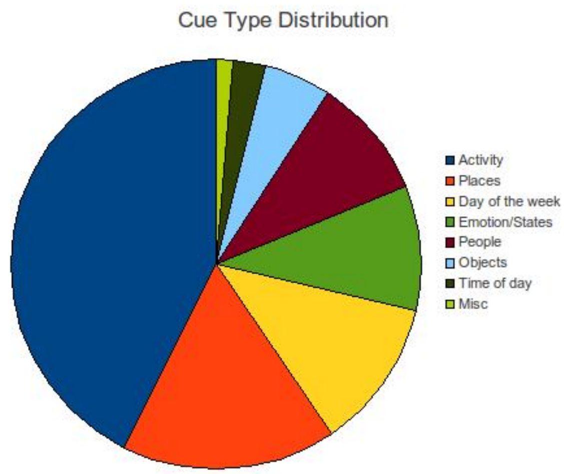

Figure 2: Cue type distribution derived from participants context labels.

We have also employed methods from dynamic system theory to understand the nature of visual and semantic context (Doxas, Dennis & Oliver, 2010; Sreekumar, Zhuang, Dennis & Belkin, 2010). Figure 3 shows a recurrence plot derived by taking the images collected from one subject ordered by time, plotting them against each other and filling in black the coordinates that correspond to pairs of images that are sufficiently similar (see the paper for the details of image representations and distance measures). The off-diagonal structure chronicles when the subject is returning to visually similar contexts. One can see immediately that the subject's life is very regular with repeated visits to the same visual contexts. All subjects show a similar pattern, although there are significant individual differences as well.

Figure 4 shows the correlation dimension plot of the same data. The correlation plot is generated by recording how many pairs of points lie within a given radius and plotting on log-log coordinates. Sreekumar et. al., (2010) found that people's visual experience is consistently two scaled. The lower scale ranges in dimensionality from 4-6 and captures within context variation, while the higher scale ranges between 9 and 13 and captures between context variation. Despite the high dimensional nature of images, visual contexts exist on a low dimensional manifold. Furthermore, we have also conducted similar analyses on a large email corpus, designed to capture the semantic contexts through which an individual traverses. The two scaled structure is also seen there, suggesting that these observations are not just characteristic of visual context (Sreekumar, 2012).

The correlation dimension plot characterizes the geometry of context representations, but is not a statement about the dynamics of context change. To assess the degree to which the dimensionalities we observe are a direct consequence of the contextual time series, we employed Taken's embedding theorem. The theorem states that under fairly general conditions a delayed embedding of the time series of an observation function will retain the properties of the original time series. To employ the theorem, we construct an observation function by applying singular value decomposition to the context vectors and extracting just the first component. Delay embeddings are constructed by running a moving window with fixed delays across this series. The dimensionality of the resultant vectors is then calculated. Figure 5 shows the calculated correlation dimension as a function of the embedding dimension (i.e. the window size). Note that the dimensionality increases as a function of embedding dimension until one reaches the intrinsic dimension of the time series. Sreekumar (2012) showed that there is a strong correspondence between the dimensionality determined using the embedding procedure and that determined on the basis of the correlation plot of the original vectors. By construction, the delayed embedding dimensionality must be a consequence of the dynamics of the system and so we can conclude that the observed correlation dimensions are not just a function of the geometry of context space, but are intrinsic to the dynamics of context.

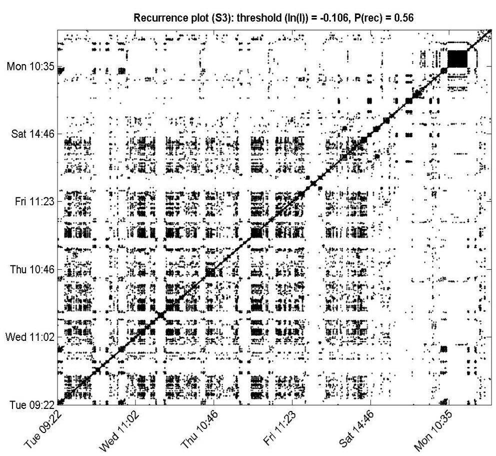

Figure 3: A recurrence plot showing the episodic structure of one subject's daily activity.

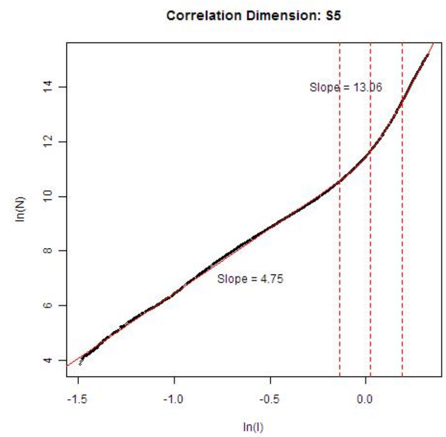

Figure 4: Correlation dimension plot shows the log of the number of images that fall within a given radius by the the log of that radius.

Yet, another way of characterizing the nature of real world context is by looking at the network structure induced by connecting similar images. Figure 6 shows this structure for one subject. Each dot represents a single image with some images expanded to provide a sense of the visual similarity. The cluster structure is apparent, but to quantify we calculated the global clustering coefficient as a function of the similarity threshold. Even for small proportions of total edges (~0.005) the coefficient is above .3 which is very high. Furthermore, the average path length on the graph is about five, the diameter is about nine and the degree histogram falls off exponentially. That is, the episodic network has small world properties like those found in semantic memory networks (Steyvers & Tennenbaum, 2005).

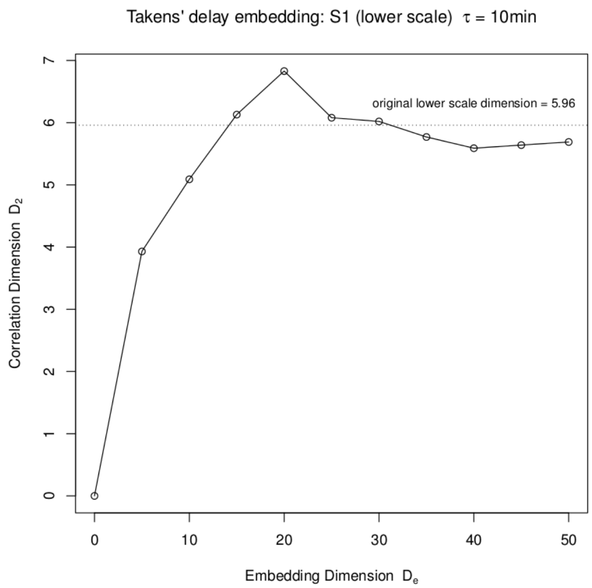

Figure 5: Taken's plot showing the dimensionality of images as a function of embedding dimension.

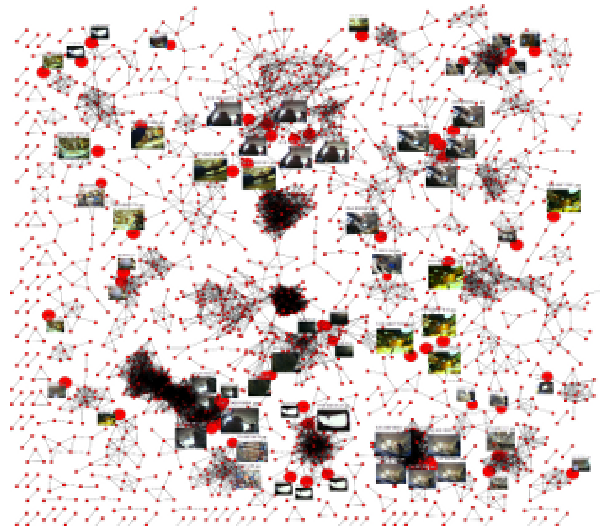

Figure 6: Network of images derived by connecting similar images.

Using multiple methods, we are starting to form a picture of the nature of real world context. The two scaled structure of visual experience and network analysis suggests that context segmentation is not just a psychological abstraction that people apply to experience, but is rather a property of that experience (albeit determined by the choices that people make). Context durations seem to conform to a power law, although more work is required to establish that and activity seems to be the dominant cue that people use to describe contexts followed by place and day of week. Overall the observed degree of regularity emphasizes how recurrent people's lives are (c.f. Song, Qu, Blumm & Barabási, 2010), a fact that has not been fully appreciated in the laboratory-based memory literature and which does not play a significant role in most memory models.

c) To empirically investigate and model people's ability to isolate when events occurred.

To understand how people isolate when events occurred, we presented subjects with images from the last two weeks of their data collection and asked them to indicate in which week the image appeared. Pilot work has indicated that there are features of the lifelogging results that do not appear in similar laboratory paradigms. Hintzman, Block and Summers (1973) found that people are more accurate at the boundaries between study lists. However, we find that accuracy is poor on the second Monday, suggesting that some form of backward telescoping occurs (Hinrichs & Buschke, 1968; see Figure 7). In addition, we tested our subjects on the Thursday following their data collection. Interestingly, there is a substantial decrease in reaction time when subjects are judging images from the preceding Thursday - suggesting a same context advantage. However, this advantage does not appear to extend to the first Thursday (see Figure 8). We are currently running additional subjects to clarify this result and intend to apply a tensor-based model of episodic memory that we have been working on to account for pure laboratory tasks to the lifelogging phenomena.

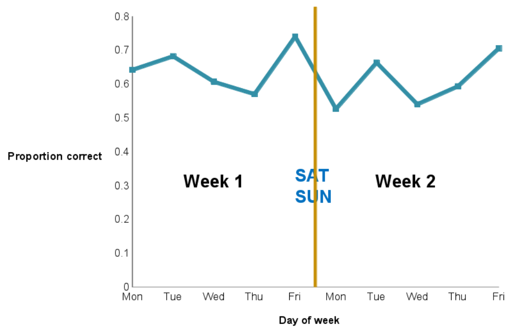

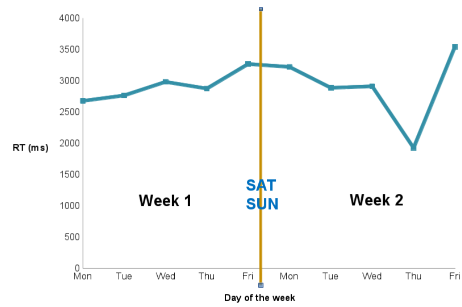

Figure 7: Preliminary accuracy data from week discrimination task.

Figure 8: Preliminary reaction time data from week discrimination task.

d) To develop algorithms capable of automatically segmenting and tagging lifelog data.

Using our initial data, we have constructed a metric-based context segmentation algorithm which relies only on accelerometer data to detect context boundaries (Zhuang, Belkin & Dennis, 2012). We defined a metric for dissimilarity of FFT features from two time windows - before and after a time point - and applied smoothing and peak-selection to detect those time points that indicate the change of contexts. We showed in the paper that the propose method outperformed similar state-of-the-art segmentation methods.

We have also developed a context tagging algorithm that uses only multisensory data to infer the status of multiple unknown tags over time which included places, activities, and people (Hamm, Stone, Belkin, & Dennis, 2012). In the paper, we proposed multisensory bag-of-words representations of data that can be combined with various state-of-the-art learning algorithms, and we performed systematic comparisons of representative classifiers from generative and discriminative models as well as temporal and non-temporal models. In particular, temporal models considered both the dependence of sensory data on the tags and the temporal dependence of tags over time. Figure 8 is an example of true vs predicted tags from the results. Among those algorithms, a large-margin based classifier for structured output (Altun, Tsochantaridis, & Hofmann, 2003) showed superior classification accuracy, achieving >0.9 accuracy in recognizing the majority of 19 tags, and achieving >0.95 accuracy in recognizing “walk”, “drive”, “chores”, “tend to baby”, “restaurant”, and “outdoor” in particular, by leave-one-day-out cross validation.

Several future challenges for the project were identified during our research on automatic segmentation and tagging. Among those, the unreliability of ground truth tags from users and the difficulty of cross-subject generalization became prominent. We are currently conducting experiments with unsupervised hierarchical models (e.g., Duong, Bui, Phung &Venkatesh, 2004) to make our approach feasible for a large collection of weakly- or unsupervised data from a larger population.

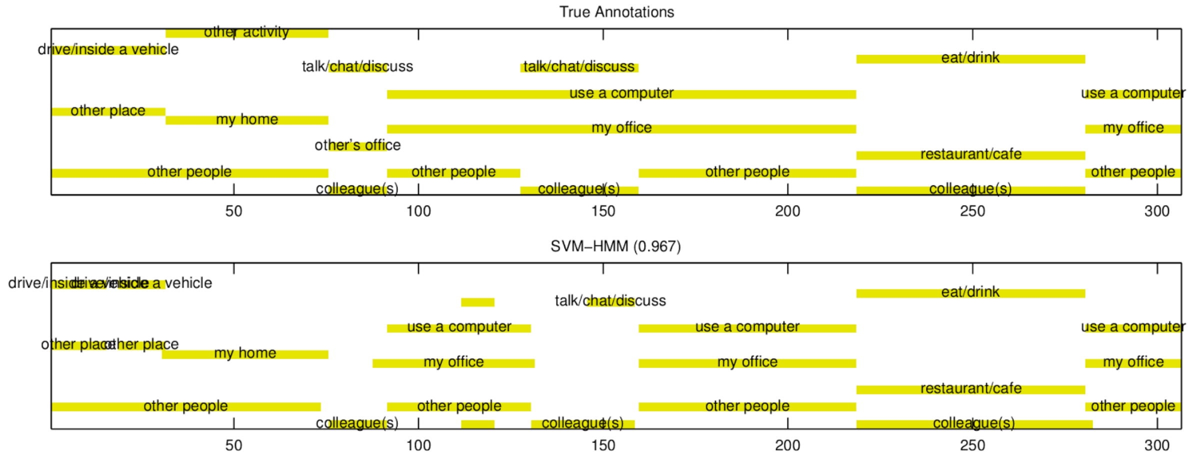

Figure 8: True vs predicted tags. The x-axis is the time in units of minutes, and the yellow bars indicate the presence/absence of each tag over time. The duration of recording was ~5 hours.

Significance of work and impact on science

Subjects do not enter our laboratories with a clean slate. In many areas of cognition, and particularly in the area of memory, the experience of the subject prior to beginning our experiments has a profound impact on their performance. Our current methods for characterizing that experience are primitive. Consequently, most researchers either ignore the problem, or try to work in domains where they expect the impact of prior experience will be minimal. If we hope to build a science of memory which is both robust and demonstrably applicable to the kinds of memory tasks people experience on a daily basis, we cannot ignore pre-experimental experience. We must do a better job of characterizing it, and in this project we have developed necessary enabling technologies and began the task of constructing a theory of context in the wild.

Furthermore, people forget stuff ($p<0.05$). While forgetting might be optimal in the face of restricted computational resources (Anderson, 1990), in general, forgetting is problematic because it prevents access to the information that would allow us to make informed decisions. Our biological memories for diet, exercise, relationship events, disease symptoms etc are far from perfect. The long term objective of this project is to eliminate forgetting. The development of writing, diaries and the personal digital assistant have all been milestones in this project. However, all of these methods have the disadvantage that they require effort on the part of the user at encoding. When they fail it is often because the information was never recorded in the first instance. We would like to produce a memory prosthesis that makes the encoding of personal information seemless. Imagine being able to search your life the same way you search the internet. We have built a prototype of such a system during the course of this project.

An ability to track context will also have implications for a broad range of areas across the social sciences. Collecting big data is not sufficient and creating mechanisms for summarizing and visualizing the raw data is only the first step. The technologies we are developing fuse multimodal data to allow us to identify what people are doing, when they are doing it, and who they are doing it with. Furthermore, we have developed methods that allow us to share that data despite its personal nature. These are critical enabling technologies for building a science of people's everyday experience.

Publications Arising from the Grant

Sreekumar, V., Kemper, R., & Dennis, S., (in prep). using lifelogging to assess memory for when.

Sreekumar, V., Dennis, S., Doxas, I., Zhuang, Y., & Belkin, M. (in prep). The geometry and dynamics of context.

Zhuang, Y., Sreekumar, V., Dennis, S., & Belkin, M. (in prep). Networks of memories.

Hamm, J.,Belkin, M. & Dennis, S. (2012). Recognizing Daily Contexts from Multisensory Data: A Codebook Approach. NIPS Workshop. Lake Tahoe, NV.

Zhuang, Y., Belkin, M. & Dennis, S. (2012). Metric Based Automatic Event Segmentation. MobiCASE'12: The Fourth International Conference on Mobile Computing, Applications and Services. Seattle, WA.

Hamm, J., Stone, B., Belkin, M., & Dennis, S. (2012). Automatic Annotation of Daily Activity from Smartphone-based Multisensory Streams. MobiCASE'12: The Fourth International Conference on Mobile Computing, Applications and Services. Seattle, WA.

Zhang, Y., Sreekumar, V., Belkin, M. and Dennis, S. (August 2010). The network properties of episodic graphs. The Proceedings of the Thirty Second Conference of the Cognitive Science Society, Portland, OR.

Sreekumar, V., Zhuang, Y., Dennis, S., & Belkin, M. (2010). The dimensionality of episodic images. The Proceedings of the Thirty Second Conference of the Cognitive Science Society. , Portland, OR.

Doxas, I., Dennis, S. & Oliver, W. L. (2010). The dimensionality of discourse. Proceedings of the National Academy of Sciences, 107 (11), 4866–4871.

References

Altun, Y., Tsochantaridis, I., & Hofmann, T., (2003). Hidden markov support vector machines, In the Proceedings of the Twentieth International Conference on Machine Learning, AAAI Press, 3–10.

Conway, M.A. & Pleydell-Pearce, C. W. (2000). The construction of autobiographical memories in the self-memory system. *Psychological Review*, 107(2), 261-288.

Dennis, S., & Chapman, A. (2010). The Inverse List Length Effect: A challenge for pure exemplar models of recognition memory. *Journal of Memory and Language.* 63 (3). 416-424.

Dennis, S., & Humphreys, M. S. (2001). A context noise model of episodic word recognition. *Psychological Review*. 108(2). 452-477.

Dennis, S., M. D. Lee, M. D., & Kinnell, A. (2008). Bayesian analysis of recognition memory: The case of the list length effect. *Journal of Memory and Language*, 59, 361-376.

Duong, T.V., Bui, H.H., Phung, D.Q., & Venkatesh, S. (2005). Activity recognition and abnormality detection with the switching hidden semi-Markov model. *International Conference on Computer Vision and Pattern Recognition*, 838-845.

Hamm, J., Stone, B., Belkin, M., & Dennis S. (2012). Automatic annotation of daily activity from smartphone-based multisensory streams. *The Fourth International Conference on Mobile Computing, Applications and Services*, Seattle, WA.

Heathcote, A., Brown, S., & Mewhort, D. J. K. (2000). The power law repealed: The case for an exponential law of practice. Psychonomic Bulletin & Review, 7(2), 185-207.

Hinrichs, J. V., & Buschke, H. (1968). Judgment of recency under steady-state conditions. *Journal of Experimental Psychology*, 78, 574-579.

Hintzman, D. L., Block, R. A., & Summers, J. J. (1973). Contextual associations and memory for serial position. *Journal of Experimental Psychology*, *97*, 220-229.

Humphreys, M. S., Wiles, J., & Dennis, S. (1994). Towards a theory of human memory: Data structures and access processes. *Behavioral and Brain Sciences*. 17(4), 655-666.

Johnson, M. K., Hashtroudi, S., & Lindsay, D. S. 1993. Source monitoring. Psychological Bulletin, 114, 3–28.

Kinnell, A., & Dennis, S., (2012). The role of stimulus type in list length effects in recognition memory. *Memory and Cognition,* 40(3). 311-325.

Kinnell, A., & Dennis, S. (2011). The list length effect in recognition memory: An analysis of potential confounds. *Memory and Cognition*. 39(2), 348–363.

Lee, M.D. (2004). A Bayesian analysis of retention functions. Journal of Mathematical Psychology,
48, 310-321.

Loftus, E.F. & Pickrell, J. E. (1995). The formation of false memories. *Psychiatric Annals* 25: 720–725.

Neisser, U. (1976). Cognition and reality: principles and implications of cognitive psychology. W H Freeman.

Newtson, D. (1976). Foundations of attribution: The perception of ongoing behavior. In J.H. Harvey, W.J. Ickes, & R.F. Kidd (Eds.), New directions in attribution research (pp. 223–248). Hillsdale, New Jersey: Erlbaum.

Rosch, E. (1978). Principles of Categorization. In Rosch, E. & Lloyd, B.B. (Eds), Cognition and Categorization (pp. 27–48), Lawrence Erlbaum Associates Publishers, Hillsdale.

Smith, S. M., & Vela, E. (2001). Environmental context-dependent memory: A review and meta-analysis. Psychonomic Bulletin & Review, 8 (2), 203-220.

Speer, N.K., Swallow, K.M., & Zacks, J.M. (2003). Activation of human motion processing areas during event perception. Cognitive, Affective & Behavioral Neuroscience, 3, 335–345.

Sreekumar, V., Zhuang, Y., Dennis, S., & Belkin, M. (2010). The dimensionality of episodic images. In R. Catrambone & S. Ohlsson (Eds.), *Proceedings of the Thirty-Second Annual Meeting of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.

Steyvers, M., & Tenenbaum, J. B. (2005). The large-scale structure of semantic networks: statistical analyses and a model of semantic growth. Cognitive Science, 29(1).

Tulving E. (1972). Episodic and semantic memory. In Organization of Memory, Ed. E Tulving, W Donaldson, pp. 381–403. New York: Academic.

Tulving, E. (2002). Episodic memory: From mind to brain. Annual Review of Psychology, 53 , 1-25.

Zacks, J. M., & Tversky, B. (2001). Event structure in perception and conception. *Psychological Bulletin,* 127, 3-21.

Zacks, J. M., Tversky, B., & Iyer, G. (2001). Perceiving, remembering, and communicating structure in events. *Journal of Experimental Psychology: General*, 130, 29-58.

Zacks, J.M., Braver, T.S., Sheridan, M.A., Donaldson, D.I., Snyder, A.Z., Ollinger, J.M., Buckner, R.L., & Raichle, M.E. (2001). Human brain activity time-locked to perceptual event boundaries. *Nature Neuroscience, 4*, 651-655.

Zhuang, Y., Belkin, M., & Dennis, S. (2012). Metric Based Automatic Event Segmentation, *The Fourth International Conference on Mobile Computing, Applications and Services*, Seattle, WA.