Building Social Life Networks

Ramesh Jain* Laleh Jalali* Siripen Pongpaichet*
Amarnath Gupta†
* Department of Computer Science, University of California Irvine
{jain, lalehj, spongpai}@ics.uci.edu
† San Diego Supercomputer Center, University of California San Diego
a1gupta@ucsd.edu

Abstract

The availability of enormous volumes of heterogeneous Cyber-Physical-Social (CPS) data streams allow design and implementation of networks to connect people with essential life resources. We call these networks Social Life Networks (SLNs). We are developing concepts, technology, and infrastructure to design and build these networks. SLNs will be helpful in addressing several essential societal problems in everyday life as well as during abnormal situations. A person needs to be connected to appropriate resources under the given situations and her own persona and context. Situations should be detected by using heterogeneous data streams. We are building a software framework for situation recognition and determining persona and personal context to connect people to resources efficiently, effectively, and promptly. We present our research in situation recognition, EventShop, persona building, and making connections using a few example scenarios.

1 Introduction

Like most other technologies, computing was developed and used for meeting human needs. It has been one of the most effective and pervasive technologies to help in diverse human needs. In the last few years its scope has exploded due to the tremendous growth in sensing, mobile, storage, processing, and communication technologies. In the form of mobile phones, armed with various sensors and actuators, computing can now reach more than 80% humans even in the remotest and most underdeveloped parts of the world [12]. In a very practical sense, we are living through one of the most pervasive social transformations being facilitated by computing related technologies. Abraham Maslow developed the famous hierarchy of needs in 1943 [10]. His theory identifies five levels in basic human needs that start from the most basic needs for survival to highest intellectual and sociological drive resulting from unique accomplishments. After physiological demands the needs at the next level are about various forms of security ranging from bodily safety, to economic security through employment, and emotional security through family. The human aspect of serving and being served by social presence starts at this level, and hence basic social norms are part of this level. When these two basic needs are satisfied, a person can live as a part of the social system. The next stage expands the social needs from mere existence to belongingness, and thus creating a personal social

Bulletin of the IEEE Computer Society Technical Committee on Data Engineering

structure that satisfies emotional needs and creates a sense of collective good. At the top of the needs hierarchy is the need for humans to seek recognition from self, and this recognition comes from the feeling that they are doing really well by contributing to the betterment of the society.

Maslow's hierarchy helps us in understanding currently popular social computing and social media. Starting from early days of civilization, new technology has always served the needs of the society. In a very real sense, most of the issues in nature are the results of needs and resources. For meeting every need, a resource is required. Therefore, it is not surprising that much of human knowledge addresses issues that ultimately lead to either creation of resources or proper distribution of resources so that they can be used to connect to proper needs. A very fundamental issue addressed by different fields of studies is how to connect people to resources effectively, efficiently, and promptly based on their need in specific situations.

In the last few years, there is growing interest in social media, crowd-sourcing, participatory sensing and social networks. Social networks are basically an infrastructure developed for connecting people to other people. Like the last few years, social networks will keep evolving to increasingly meet needs of society. Current social networks focus mainly on the third and fourth level of Maslowian hierarchy. We believe that the computing infrastructure is now ready to meet the needs at the first two levels of the hierarchy also. We call such an infrastructure Social Life Networks (SLNs) [1, 2]. Here we discuss an approach for building SLN.

It is essential that SLN bridge across what is commonly called cyber, physical, and social systems. Increasingly, emerging systems have sensors to provide data about physical world, and humans provide data about the social world. All this data is assimilated and processed in cyber world of computers. This has serious implications for data engineers for designing systems to solve new challenging problems. In the following, we first discuss the changing nature of data and computing and then present our approach to building SLN.

2 Recreating Dynamic Real World

A very important aspect of the real world is that it is dynamic. It is always evolving. Most of the new data is collected to capture the dynamic nature of the world. The dynamic data about the past is important because one can model the world using that data and use those models to understand the world, and more importantly, predict future events. Until a few years ago, information systems limited collection and processing of data to very limited aspects of the world. Moreover, many systems were designed assuming that updates were less frequent than retrieval of data. A good example is the popularity of enterprise data warehouses towards the end of the last century. These systems collected data related to limited aspects of operations of an enterprise and tried to gain insights and understanding to create business intelligence. Around the same time one saw beginnings of stream processing starting from streams of stock market data to the so called complex event processing that dealt with limited number of data streams in an organization. Just in less than two decades, the situation has changed dramatically. Collectively, now we are creating a global data warehouse that involves massive number of heterogeneous data streams that must be analyzed in real time for predicting evolving situations and managing them. Some fundamental differences between the data about two decades ago and now are obvious. These differences are essential to building systems that will span cyber, physical, and social aspects of computing. We believe that some important differences are:

1. Most of the data used to be structured and usually human mediated. Now most of the data comes from disparate sensors as data streams.
2. Data used to be related to well-defined operations and expected events. Now the number of operations and events is vast and one needs to deal with the unexpected all the time. Unexpected is the new expected.
3. Data was either unrelated to geo-location or had only a few geo-spatial streams. Now the world is being covered by sensors producing dense data streams. Soon almost every point will be covered by

multiple sensors. Also, the granularity at which data is being collected is becoming finer and finer every day. Location of data has become a fundamental attribute.

4. Most decisions were required for planning the operations of an enterprise. Increasingly, most decisions are communicated to actuators for action. In many cases similar to human-sensing (in Twitter), human actuation (as in flash mobs) may be used in situations like disaster management. Suggestions and actuation are becoming more common than planning.

An obvious fact is that physical world events and situations take place in physical space. Situations are the result of interactions among several related events. Events are the results of some happenings that are due to significant state changes. State changes can be observed by measuring some physical attributes using a sensor. Usually a sensor measures only a specific attribute, which is one of many different independent or correlated attributes required to detect the state change. Norbert Wiener [11] showed that cybernetics is the theory of control and communication in humans as well as machines. His work may be considered first work in building CPS. Systems theory was strongly influenced by his research. Inspired by his approach, we say that the state at a point in space at a given time can be characterized by N attributes. Thus,

$$S = [a_1, a_2, ..., a_N]^T$$

Where S is the state of a point (x, y, z) and a_i is an attribute at the point as measured by a sensor at time t. For brevity we are dropping (x, y, z) and t from each variable and will be assuming until specifically mentioned that all discussion is related to measurements and sensors at location (x, y, z) at time t.

It must be mentioned here that in CPS, many different types of sensors are used. The values obtained directly from sensors may range from measured facts to expressed opinions. This is a challenge for CPS systems to appropriately deal with these values from disparate sensors to reconstruct the state.

The information obtained from sensors varies in many respects. Methods to convert data to information and the reliability of information could be entirely different for different sensors. Humans used as sensors, as in participatory sensing, provide opinions, not measurements. The goal of many research projects with Twitter data is to develop techniques to convert these opinions to one or more attributes above. Physical sensors provide either direct measurements (as in a thermometer) or indirect measurements that are a result of many-to-one mapping. To convert measurements to attributes at a point in space, complex approaches involving inverse mapping are used. Cameras are the best example of such sensors. Thus, one can say that

$$a_i = f_i(m_i)$$

where a_i is the attribute derived from a measurement m_i using the function f_i for this attribute. The function used to convert the measurement to an attribute could be very simple or it could be extremely complex. Good examples of complex functions required for such processing are computer vision systems and tweet analysis systems. As is well known, in a computer vision system, the intensity at a pixel is the result of a 3-dimensional to 2-dimensional projection as well as the light ray's behavior in the environment. Computer vision systems have been so challenging because of this complex many-to-one mapping that should be inverted to find what the pixel really represents. The data value at a pixel is only suggestive and in itself not much important. When analyzed in concert with other values and context, it becomes valuable and useful.

In the last few years, many computer science researchers have been fascinated by analysis and use of Twitter data [3,4]. We believe that this is because a Tweet is text and computing community finds it easy to deal with. Automatic tweet analysis faces formidable challenges. A tweet is the result of a person expressing her opinion on a topic. Though a tweet may have location and time associated with it, the topic of the tweet may be unrelated to its origin. Another challenge is that different people have different language, style, and motivations. To make this more complex, people need to package their thoughts in less than 140 characters so they need to invent suggestive language. All these factors make tweets a weakly related, uncertain, suggestive source of information. One needs to consider this nature of tweets in using it as a source of measurement coming from human sources.

The above analysis of these two different sources shows that one needs to seriously consider the nature of data sources. In data engineering research, until recently the source and nature of data was usually direct symbolic input in structured form. With emerging systems, this is just not doable for most emerging applications. CPS systems present a formidable challenge in "integrating" uncertain heterogeneous data streams with minimal latency. Given that the problems that could be solved using these systems are so important and rewarding; these challenges should not be ignored.

3 General Architecture of the SLN System

In building SLN, we adopt a perspective that there are sensors, databases, and social networks that are observing, storing, and reporting what is happening in the world. A subset of all these information sources is used to characterize the status of available physical or human resources. Similarly, another special subset is used to characterize needs. Thus, we consider all data sources in three classes: Observers, Resources, and Needs. There is overlap and change of roles among these classes. Since the world is dynamic, we consider these classes only in the context of a specific application.

Situation recognition is a central component of the SLN. Situations are the result of interactions among several related events. Events are the results of some happenings that are due to significant state changes. Based on the situation at a place, the system Identifies needs and available resources to satisfy those needs. Situation recognition is always based on the context of a specific application and so are all other operations. The data sources used by the system are also those publicly available or specifically made available in the context of the application. The system closes the loop by sending actuation or action information to appropriate needs and resources as a result of the matching.

An SLN system, as shown in Figure 1, then essentially implements the following loop:

1. Observe multiple real-time data sources
2. Determine interesting atomic events
3. Combine events using complex event definitions
4. Identify situations from these complex events based on personal, social and global context
5. Determine needs using the situations and a set of needs definitions
6. Monitor known resources based on their capabilities, availability, reachability and cost of access
7. Find maximally matching resource for each need detected
8. Communicate decisions and perform actions commensurate with needs by feeding the information back to the social network

Here, not only people, but other objects like mobile applications (e.g. body activity monitors if allowed by the owner), databases, and the Internet of Things (e.g., traffic sensors) also observe, store and report information about the state of entities in the world. In this setting, we conceive of a world where (a) a significant body of information today comes from sensors, (b) the number of sensors is huge and the number of events generated by them is even larger, (c) a large fragment of data, both human and device generated, have associated locational information, (d) most situation and needs assessment decisions are for controlling and managing real time and evolving situations, and (e) keeping pace with the real-time nature of our problem space, planning and decision processes need to be viewed like a real-time control system that interoperates with the publish- subscribe and update-propagation model of standard social networks.

A user update in a social network is analyzed to create a microevent (or a personal event), which is then fed to the situation recognizer. The situation recognizer evaluates this microevent with respect to other events from different sources and creates an action (e.g., a message, a recommendation, an alert) that goes back to the sender or a potential resource that can service the needs of the original message sender.

In the next section we describe two prototypes of situation recognition modules that take two different approaches to detect situations. The first, called EventShop [7, 8], is designed to detect situations that are global and spatio-temporal from sensor streams that observe some portion of a real-world phenomenon. The second, called Personal EventShop, is designed to detect situations about a person's private world. It can serve as a personal agent that recognizes and handles a situation "locally" when possible, or creates a micro-event.

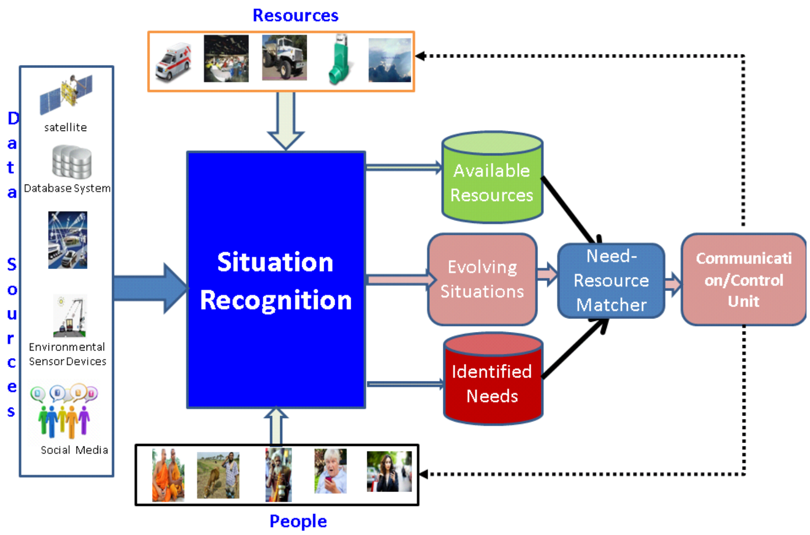

Figure 1: A Social Life Network connects people to resources based on their need in specific situations. Here we show a high level architecture of a SLN system that identifies needs and available resources in a specific situation and then connects needs to resources.

4 Recognizing Situations

4.1 EventShop: The Prototype of a Global Situation Recognizer

Inspired by programming models for distributed systems, we have developed an open source situation modeling and recognition platform called EventShop (ES). The ES framework hides the complex implementation in the background allowing users to simply create massive distributed systems for situation recognition as well as allowing users to focus solely on their application logic and the semantics of their systems. The ES framework has a) a programming model that allows for complex event streaming applications to be created without distributed systems expertise, and b) a prototype implementation of the ES framework that proves its capability to handle wide variety of heterogeneous data streams observing real world events on the internet scale in real-time.

Unlike other CEP systems, ES uses a spatial grid stream as its data model because it is naturally suitable for representing various geo-spatial data; each cell of the grid stores value of certain measure taken from the corresponding geo-location. We adopt the grid structure, and call it E-mage (an event data based analog of image) [6]. We believe that this generic data model can be used to integrate heterogeneous data coming from spatio-temporally distributed web streams. For handling data streams, a special attention needs to be paid to the semantics of aggregated data. Since the data streams are unbounded, combining such data to find simple aggregated value such as summation and average is unclear. This problem is normally resolved by introducing windows which transform an unlimited sequence of data streams into windows of data. In this work, we use tumbling [5] window that splits data streams into non-overlapped contiguous windows.

There are four main components in the ES framework which are data ingestor, stream processing, internal storage, and situation output. In the data ingestor component, original raw spatio-temporal data, either real-time or near real-time streams from the Web are translated into unified STT (Space-Time-Theme) format along with their numeric values using an appropriate data wrapper. A data wrapper for a sensor converts a measurement in a semantic attribute as shown in equation 2 above. Based on users' defined spatio-temporal resolutions mapper, the system aggregates each STT stream to form an E-mage stream. This E-mage stream is then pulled by the stream processing and/or transferred to internal storage. Based on the situation recognition model determined by the domain expert using something akin to equation 1 above,

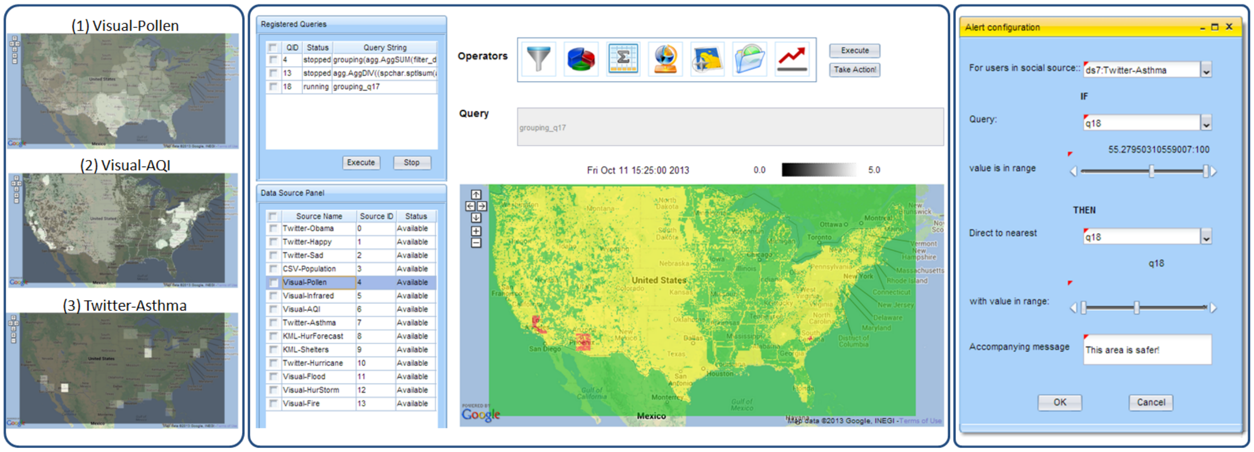

Figure 2: Asthma Relief Application (from left to right): Data Sources E-mages, ES Web UI, and Situation Action Rule.

appropriate operators are applied on the E-mage streams to detect situation. In most of the cases, the final step is a segmentation operation that uses domain knowledge to assign appropriate class to each pixel on the E-mage. This classification results in a segmentation of an E-mage into areas characterized by the situation there. Once we know the situation, appropriate actions can be taken depending on application action control rules in the situation output component. From the application developers' point of view, a workflow of moving from heterogeneous raw data streams to actionable situations consists of three simple steps: 1) register or select appropriate data stream sources, 2) register complex event model by combining a rich set of built-in operators and 3) define event condition action rules to send personalized alerts to relevant people.

The operators of EventShop are used to detect interesting E-mages based on value-patterns on the spatial grid stream. We analyzed situation recognition models across multiple domains and defined the initial set of operators, which are generic enough to capture most of the common requirements. These operators are: selection, segmentation, aggregation, spatial characterization, spatial pattern matching, temporal characterization, and temporal pattern matching. We developed the EventShop system so that a situation modeler can define an algebraic plan using the operators above to specify a situation. The output of such a situation detection plan is an E-mage that can subsequently be transformed to a convenient form and sent directly to a client, to a social network system, or to an end application used by the client.

4.2 Use Case: Asthma Relief Application

For the past couple of years, several applications have been designed and developed using the EventShop framework to recognize real world situations. for example, hurricane detection and migration, demand hotspots of business products identification, flu outbreak, wildfires detection, flood migration, and allergy risk and recommendation [7, 8]. Here, we demonstrate asthma relief application, the most motivating one. The goal is to suggest safer areas to people who are in asthma risk zones. From asthma study, the severeness of an environment to an asthma patient is related to the pollen count and air quality in that area. From crowd sourcing aspect, if many people from a specific area discuss about asthma, it usually suggests that the situation in that area is not very friendly to asthma patients. We combine data from these three data sources, and then segment the aggregated data over entire US into three danger zones based on the values in aggregated E-mages as shown in Figure 2. Then, a situation action rule is created to broadcast a safe area to an individual. Although, these messages is partially personalized based on end-userâĂŹs location, we could provide more specific and useful recommendations if more detailed information about end-user's context was available.

4.3 Personal EventShop

To close the loop in a SLN, the system should send personalized action instructions/requests to users. This requires knowing persona and personal context. We realized that the persona can be computed by using several data streams related to a person. For this task many processing operations are similar in concepts to EventShop. We are building Personal EventShop (PES) for meeting this need. PES is designed with the assumption that the states, events and situations recognized by it come from observing personal information streams. It uses an individual's life events from sensors and mobile devices that capture fitness data, personal events, eating habits, sleep patterns, and every day activities.These information comes from more structured information sources like personal calendars, as well as less structured information such as activities on her social networks. All these data sources represent data streams related to the person's life. The persona is built by collecting and analyzing these streams in a long term data warehouse for the person. For converting different data streams to related event and situation streams, mathematical and computational operators similar to those used in EventShop are required. The event data from Personal EventShop will be stored in a personal data warehouse where all data and important results of intermediate computations will be stored. Key aggregates and results required to build persona may also be saved for rapid computations of requested information and insights when needed.

The major difference between the EventShop and PES is that PES deals with an individual related stream so is basically one dimensional stream, while ES deals with physical space so is either 2-dimensional or 3-dimensional grid. All activities of a person in her persona should be related to her life events.

5 Research Challenges

We believe that there are several challenges in building a SLN platform that will address many emerging applications. We consider the following as particularly interesting and important research challenges.

1. Massive Heterogeneous Geo-spatial Stream Processing: Traditional data processing techniques considered data streams as a sequence of data items. Increasingly geo-spatial heterogeneous data streams at different granularity must be combined to detect emerging situations. This requires a different perspective on processing and managing data.
2. Situation Recognition: Complex Event Processing was satisfactory when we had a few data streams and the result of CEP was of interest to an organization. Evolving situations affect large number of people and resources and must be determined using complex spatio-temporal semantics of streams. This is a new research area.
3. Persona and Personal Context: Most search engines create very narrow persona for users to serve them advertisements or for recommending them products. Emerging applications have large number of personal data streams that could be used to build more sophisticated persona that could be used in diverse applications considering more specific user context.
4. Chronicle Analytics and Visualization: Enterprise data warehouses started making analytics and visualization of data popular. Big data has created significant interest, some say hype, around analytics and visualization. Now one should consider chronicles at different level and in different situations and develop techniques for analysis, prediction, and visualization with as little latency as possible.
5. Dynamic Need-Resource Optimization: Now that we are developing techniques for situation dependent need identification as well as resource availability, we need to extend techniques for optimal satisfaction of needs in given situations with as little cost as possible. This becomes a challenging task.

6 Conclusion

We are living in a very exciting time. In the last few years data storage, processing, and communication technologies have transformed the nature and volume of data that will be routinely used in emerging applications. The new data is massive number of geo-spatial heterogeneous data streams. This forces us to rethink not only about the data, but also about the processing, communication, storage, analysis, and visualization. This also opens up unprecedented opportunities for addressing serious societal problems that we could not consider earlier. Very challenging, exciting, and rewarding time is ahead for people interested in data.

References

[1] A. Gupta and R. Jain. Social life networks: A multimedia problem? In *Proc. of ACM Conf. on Multimedia (SIGMM'13)*, 2013.

[2] R. Jain and D. Sonnen. Social life networks. *IT Professional*, 13(5):8–11, 2011.

[3] H. Purohita, A. Hamptonb, V. L. Shalinb, A. P. Shetha, J. Flachb, S. Bhatta. What kind of conversation is Twitter? Mining psycholinguistic cues for emergency coordination. *Computers in Human Behavior*, 29(6):2438Ð-2447, 2013.

[4] C. H. Lee. Unsupervised and supervised learning to evaluate event relatedness based on content mining from social-media streams. *Expert Systems with Applications*, 39(18):13338Ð-13356, 2012.

[5] J. Li, D. Maier, K. Tufte, V. Papadimos, and P. A. Tucker. No pane, no gain: efficient evaluation of sliding-window aggregates over data streams. *ACM SIGMOD Record*, 34(1):39–44, 2005.

[6] V. K. Singh, M. Gao, and R. Jain. Social pixels: genesis and evaluation. In *Proc. of the ACM Int. Conf. on Multimedia (MM'10)*, pages 481–490, 2010.

[7] M. Gao, V. K. Singh, Singh, and R. Jain. EventShop: From heterogeneous web streams to personalized situation detection an control. In *Proc. of the 3rd Annual ACM Web Science Conf. (WebSci'12)*, pages 105–108, 2012.

[8] S. Pongpaichet, V.K. Singh, M. Gao, and R. Jain. EventShop: Recognizing situations in web data streams. In *Proc. of the 1st WWW workshop on Web Oservatories (WOW'13)*, 2013.

[9] L. Jalali, and R. Jain. Building health persona from personal data streams. In *Proc. of the 1st ACM Workshop on Personal Data Meets Distributed Multimedia (PDM'13)* , 2013.

[10] A.H. Maslow, A theory of human motivation. *Psychological Review*, 50(4), 370Ð96, 1943.

[11] N. Wiener. Cybernetics or the control and communication in the animal and the machine. *Second Edition, MIT Press*, 1948.

[12] R. Scoble, S. Israel. Age of context: mobile, sensors, data and future privacy. *First Edition, CreateSpace Independent Publishing Platform*, September 5, 2013.