O

Obscenity

▸ Pornography Online

Observation

▸ Sources of Network Data

Old Boys Network

▸ Interlocking Directorate Networks

Online

▸ Mapping Online Social Media Networks
▸ Social Networking in Political Campaigns

Online Collaboration

▸ Wikipedia Collaborative Networks
▸ Wikipedia Knowledge Community Modeling

Online Commerce

▸ E-Commerce and Internet Business

Online Communication

▸ Privacy Issues for SNS and Mobile SNS

Online Communities

Anatoliy Gruzd
School of Information Management, Dalhousie University, Halifax, NS, Canada

Synonyms

Computer-mediated group; Online group; Virtual community; Virtual settlement

Glossary

API	Application programming interface
CMC	Computer-mediated communication
MUD	Multiple-user domain/a multiplayer, real-time, virtual world
SoC	Sense of community

Definition

Online communities (also referred to as virtual communities), like their offline counterparts, can

R. Alhajj, J. Rokne (eds.), *Encyclopedia of Social Network Analysis and Mining*,
DOI 10.1007/978-1-4614-6170-8,
© Springer Science+Business Media New York 2014

be characterized as groups of people who share similar backgrounds, values, or interests and meet regularly. However, the unique characteristic of an online community is that its members form and maintain their relationships via computer-mediated communication (CMC) technologies such as online discussion boards, web blogs, and social media (For consistency, we will use "online communities" throughout this entry instead of using other synonyms.).

An online community can be exclusively "virtual" when members only meet via the Internet or may include both online and face-to-face interactions. Members can interact with each other synchronously (in real time) or asynchronously by posting and replying to others' messages at their convenience. An example of an Internet platform that supports synchronous interactions is a virtual environment called "Second Life," where users interact with each other in a 3D virtual reality through digital representations of themselves called "avatars." An example of an Internet platform that supports asynchronous interactions is an online discussion board where users do not have to be online at the same time to communicate with each other. An online community can utilize more than one CMC tool or more than one mode of communication. For example, on Facebook, a popular social networking web site, users often communicate asynchronously with one another by replying to direct messages or posting messages on each other's virtual profile "walls," but they can also converse in real time by using a text or video chat option if both users are online at the same time.

Introduction

There is a wide variety of online communities that exist on the Internet, ranging from a community of researchers trying to solve the next big problem in their field to a band's fan club, from online communities of political or environmental activists to cancer support groups. Online communities vary based on many different characteristics such as their communal goals, the types of CMC technology they use, the demographic composition of its members, and/or geographic coverage (local versus international groups). Due to the multitude of different factors that characterize an online community, there has been a lot of ongoing work to develop an effective typology of different communities on the Internet. One of the most popular typologies of online communities was proposed by Porter (2004). The first level of Porter's typology differentiates between online communities that were founded by their members and those that were established by organizations. The next level of this typology categorizes communities based on the nature of the relationship among their members. For example, member-initiated communities would be subdivided into social or professional communities and organization-sponsored communities would be categorized based on the types of parent organization: commercial, nonprofit, or government.

Just as there are many different types of communities that exist on the Internet, there are many different methods to study them: ethnography, non-participant observation, survey, interviews, content analysis, social network analysis, etc. Considering the focus of this Encyclopedia, however, the emphasis here will be on studies that relied on social network analysis (SNA) to gain insights into operation of online communities. Since the early 1980s, the number of publications that used SNA has grown drastically (Otte and Rousseau 2002), introducing SNA to researchers from various disciplines, including organizational and community researchers. The inception of the International Network for Social Network Analysis (INSNA) founded by noted sociologist Barry Wellman in 1978 (http://insna.org) furthered SNA proliferation and helped to establish SNA as the leading method to study a wide variety of communities including communities of students, physicians, addicts, cancer survivors, and other populations.

However, until recently, capturing information about connections among community members has been very resource intensive and was primarily done via time-consuming and intrusive social network surveys. The advantage of studying *online* communities is that most (if not all) of their social interactions are recorded and can be used to study the group's underlying social structures automatically. Once a social network of an online community is discovered, it can be used to study group dynamics, active participants and outliers, the formation of shared norms, and other social processes in online communities.

Key Points

This entry discusses different types of online communities that exist on the Internet and applicable automated methods to discover underlying social networks in these communities which then can be studied using SNA. The discussion of the different types of online communities and corresponding network discovery methods will be grouped by the type of CMC platform that a community is based on. Different CMC platforms allow different ways for their users to interact; thus, they facilitate the formation of different types of relations among their users and require different approaches for discovery of social structures. The historical background section that follows will offer a brief review of the evolving notion of "communities" in the era of the Internet and social media.

Historical Background

Community scholars have been debating for decades about what constitutes "a community." Most of the early definitions of "community" stipulated that a community is geographically bounded. However, since the 1970s, fueled by the proliferation of communication technologies and the increasing mobility of the world population, the notion of "community" has been changing to also incorporate long-distance relationships (Wellman and Leighton 1979). The idea that a community does not have to be linked to any particular geographic location and that community members do not have to be physically coterminous was strengthened in the 1990s with the rise of the Internet and, subsequently, the emergence of social media and social networking technologies. Together, these technologies have made it easier for people to form and maintain social connections across space and time.

These technological changes led to the expansion of the traditional definition of a community to include the notion of "online community," referring to communities formed and/or supported over the Internet. To explain users' behavior in online communities, new community-driven concepts emerged such as "networked individualism," proposed by Barry Wellman, which describes the fundamental shift from group think and hierarchal relations to a "networked society," where individuals create their own "personal communities" consisting of multiple, loosely connected (and often overlapping) communities of their family members, friends, coworkers, acquaintances, and others. Although not unique to online communities, the concept of "networked individualism" is especially relevant and aptly describes how Internet users interact in various online communities (Wellman 2001).

Another prominent concept in this area is "media multiplexity," coined by Caroline Haythornthwaite and Barry Wellman, which suggests that networked individuals who are closely connected to each other tend to communicate more frequently, share more kinds of information, and rely on more types of CMC for their information and communication needs (Haythornthwaite and Wellman 1998). Later, concepts such as "local virtuality" (describing people who are close geographically but still rely on CMC for communication and information exchanges between each other) and "hyperconnectivity" (referring to the ability to reach people via CMC anywhere and anytime, were introduced by Quan-Haase and Wellman (2005).

Based on a 2009 bibliometric study, since the early 1990s, the number of published peer-reviewed articles on online communities has been on the rise covering five major disciplines (Information Systems, Psychology, Management, Computer Science, and Sociology), peaking in 2004–2005 (Iriberri and Leroy 2009). Among this body of knowledge are studies that explored what constitutes an online community and how to determine whether a group of people who are interacting via the Internet can be classified as an online community. As part of this ongoing process of defining and redefining online communities, community scholars turned to a number of established theories of physically bounded communities to see to what extent they can be used to also define and identify online communities.

One such theoretical foundation is Benedict Anderson's notion of *Imagined Community* (Anderson 1983). Anderson studied developing countries such as Indonesia where the local governments tried to construct a common identity among various ethnic groups in their country who lived on different islands. In this work, Anderson considered three key elements of community formation: the presence of a common language, temporality (shared history), and "high centers" (community leaders). Interestingly, these three elements have been shown to be applicable to online communities as well. The first element is the development of a common language, which facilitates inter group communication and helps to build a communal identity and make it more difficult for outsiders to participate in a community. Just like in Anderson's communities, Internet researchers found that many online communities develop their own common language (e.g., slang, common terminology). Studies exemplifying this trend include Marvin (1995), who observed the development and use of a specialized communal vocabulary in early MUDs (multiple-user domains), multiplayer, real-time, virtual worlds, and Sen et al. (2006), who examined the evolution of users' vocabulary in social tagging communities. Second, following Anderson's work, a community needs to have a shared history. This holds true for online communities as well (see, e.g., Stanoevska-Slabeva and Schmid 2001). That is why sustained membership over time and the conversion of newcomers to active members are crucial for the stability and success of online communities (Iriberri and Leroy 2009). Finally, the third element discussed by Anderson is "high centers," referring to the idea that a society operates around powerful individuals or "high centers." At first glance, this element may seem to be less relevant to online communities since the Internet offers a decentralized structure where online participants are free to choose what online communities to join, what messages they post, and when. However, based on the previous research, only a very small proportion of online users in a group contribute to the group – a phenomenon known as the "long tail" of user participation. Thus, only a small proportion of users actually influence what is being discussed. For example, only 4 % of members in open-source development communities contributed 66 % of fixes and 88 % of new code (Mockus et al. 2002). Members who participate in online communities by reading other's contributions, but do not contribute themselves, are called "lurkers." For example, for email-based discussion lists, the percentage of lurkers in an online community ranged from 46 % in health-support groups to 82 % software-support groups (Nonnecke and Preece 2000). Similar patterns have been observed in other online communities.

Another popular theory that was adopted from studies on "traditional" communities and validated with online communities is McMillan and Chavis' (1986) sense of community (SoC). It states that one of the most important criteria for community existence is the presence of a sense of community and that SoC is present if members feel that (1) they belong to the community, (2) they can influence it, (3) they provide support and are offered support by others (integration and fulfillment of needs), and finally (4) they share emotional connection with others in the group. Examples of studies that used this theory to study online communities include a community of blog readers (Blanchard 2004), a scholarly community on Twitter (Gruzd et al. 2011).

Lastly, in the late 1990s a new framework for studying online communities called "virtual settlement" was proposed by Jones (1997). Rooted in CMC and cyber-archaeology, this framework defines a prerequisite for online communities called "virtual settlement."

Virtual settlement exhibits the following four characteristics: interactivity among community members, presence of more than two communicators, existence of a common-public-place where members can meet and interact, and evidence of sustained membership over time. Some examples of studies that relied on this framework include studies of web blog communities (Efimova and Hendrick 2005) and online communities in Second Life (Harrison 2009).

Studying Online Communities

This section will review studies of online communities that relied on both SNA and automated methods to extract information about networks. The studies highlighted herein are grouped based on the different types of CMCs being used to sustain an online community: online discussion forums, blogs, social networking web sites, microblogging web sites, content sharing web sites, wikis, and virtual worlds.

Communities Based on Online Discussion Forums

Many early studies in this area looked at online communities formed around online forums or discussion boards. For example, Fiore et al. (2002) studied online communities that emerged on Usenet discussion groups. Durant et al. (2010) examined social networks among users of a medical forum. These studies primarily relied on information about "who replies to whom" on a forum, found in the message headers, to discover underlying social structure. However, since messages in online forums are open to all group members, it is often not clear who are the actual addressee(s) of a message from only examining the message header. Recent work in this area explored alternative ways to extract network information from online forum discussions. For example, in addition to relying on who replies to whom information, Gruzd (2009) proposed a method called "name networks" that examines the content of each message to identify any personal names automatically. Discovered personal names have shown to be good indicators of the actual addressee(s) of the message and, thus, helped to reveal social connections between the sender of the message and the addressee(s). Gruzd also proposed to use a variation of the name network method – a "name co-occurrence network" that connects people if their names coappear in close proximity in the same message(s). The name network method was validated on a community of online learners (Gruzd 2009). A major advantage of name network-type methods is that they can be applied to discover social networks of online communities where conversations do not necessarily follow a structured format, such as comments posted in response to a YouTube video.

Communities of Bloggers and Blog Readers

Another group of well-studied online communities are communities of bloggers and blog readers. Communities in the blogosphere have been primarily discovered through the analysis of hyperlinks between blogs found in blog posts, blog comments, and "blogrolls" (a list of related blogs, usually manually compiled by a blogger) (e.g., Ali-Hasan and Adamic 2007). Once interlinks between blogs are found, communities of related blogs are then determined by identifying densely connected groups of blogs. This technique is rooted in early studies of identifying communities of web sites (e.g., Gibson et al. 1998). This method identifies online communities in a loosely connected sense since there are many reasons why one blog links to another – it can indicate agreement, disagreement, a related topic, a friendship or collegial connection, etc. – and since there may be more than one person contributing to a single blog (including comments posted by readers). For example, using hyperlinks between blogs, Gruzd et al. (2012) examined communities among blogs on diabetes. Pikas (2008) studied communities in the science blogosphere. Adamic and Glance (2005) studied communities of blogs in the political blogosphere. A number of studies have also looked at how communities emerge among blog readers on some popular blogs (e.g., Chung et al. 2010), where instead of relying on links between blogs, the researchers used a “name network” method to identify connections among blog commentators.

Communities on Social Networking Web Sites

The number of social networking web sites and their user bases grew exponentially in the early 2000s; their popularity subsequently attracted the attention of many community researchers. To date, Facebook is one of the most studied social networking web sites. This comes as no surprise since it is currently also the largest with over 1 billion users worldwide (Facebook.com). Due to the nature and purpose of social networking web sites, Facebook connections tend to better represent people’s social connections with family members, friends, co workers, and others unlike connections that can be uncovered in the blogosphere. An extended discussion of Facebook-related research can be found in Wilson et al. (2012).

Since users of social networking web sites self-report their “friends” connections, there is usually little need to infer connections between users automatically. So the main challenge for researchers here is how to retrieve social network data (also referred to as “social graph” or “friend of a friend” data) already collected by these sites. To facilitate access to social network data, a social networking site such as Facebook would often provide a data protocol for this purpose. However, due to the private nature of users’ information, prior consent is usually required to retrieve such data automatically. In the case of Facebook, such data can be collected via a Facebook application, developed and distributed via their Developer Platform. To start data collection through this application, a researcher needs to invite a Facebook user to install their application. During the installation process, the application will request access from users to retrieve their social network data. A limitation of this approach is that a user’s consent would only grant access to information about that user’s immediate friends and connections between them; in other words, it only provides information about their personal (or ego) network. This makes it difficult for researchers to study online communities on Facebook, since the retrieved network data only represents a skewed viewpoint of a community (or multiple communities) from the perspective of a single user. To address this concern, researchers usually sample multiple users to reconstruct the whole or partial social network representation of a community. More details about the process of collecting and analyzing Facebook data are discussed in Hogan (2008).

Other approaches to studying communities of users on Facebook from the network perspective include relying on anonymized clickstream data from HTTP headers and the automated harvesting of public Facebook profiles. Although the ethical implications of automated collection of online social network data are outside of the scope of this entry, it is important to note that the later method in particular has raised a number of privacy and ethical concerns.

Communities on Microblogging Web Sites

In recent years, microblogging web sites in general, and Twitter.com in particular, have emerged as a popular platform for sustaining online communities. Like Facebook, Twitter allows its users to establish connections with each other by “following” other users of the system. But, unlike Facebook, it does not require users to “follow” back their followers, leading to the creation of asymmetric relationships among its users. Another interesting characteristic of Twitter is that nearly 40 % of connections on Twitter are between people who reside in the same metropolitan area, suggesting that geography still matters at least in online communities on Twitter (Takhteyev et al. 2012). Twitter users can post their own messages (up to 140 characters long, also known as tweets); they can also favor, reply, or share others’ messages. The action of sharing somebody else’s message on Twitter is called “retweeting.” Each of these interactions on Twitter connects one user to another. To study online communities developed on Twitter, researchers most frequently rely on networks derived from user interactions: who follows whom, who retweets whose messages, who mentions whom, and who replies to whom. Public information about users and their interactions with others can be retrieved using Twitter API (application programming interface). Previous studies that used SNA to analyze Twitter data examined communities such as scholarly communities (Gruzd et al. 2011), communities around political discussions (Conover et al. 2011), and even an online book club (Gruzd and Sedo 2012).

Communities on Content Sharing Web sites

Other web 2.0 web sites such as YouTube, Flickr, and Digg (popular video, photo, and hyperlinks sharing web sites correspondently) also have social networking features that allow people to indicate who is connected to whom. Similar to blogs, connections declared on such web sites are less about social relationships and more about shared interests. A number of studies have examined the nature of online communities formed on these sites. Klausen et al. (2012) studied an online community that disseminates Al-Qaeda's propaganda on YouTube and compared it to a community of YouTube users that support the Texas Tea Party movement in the USA. And Zhu (2010) examined social structures among Digg users, focusing on discovering factors behind users being influential in a community.

Communities on Wikis

Wiki communities are communities that formed around "wikis," web sites specifically designed for collaborative writing and editing of interlinked web pages. One of the most studied web sites in this category is Wikipedia, a wiki-based open encyclopedia. Wikipedia researchers are especially interested in studying group dynamics in online communities of Wikipedia contributors. Often many people contribute content to a single Wikipedia page; therefore, it is possible to make an assumption that contributors to the same page are somehow connected. The site tracks any changes made by its contributors. Researchers in this area use the log of changes to a particular page to track how contributors form connections on this site through collaborative editing. To discover possible connections among contributors, researchers first build a so-called "two mode" network, containing two types of nodes: contributors and wiki pages. At first, contributors are not connected directly to each other, but through the wiki pages that they edited. To discover connections between contributors, a two-mode network is converted to a one-mode network (called "coediting") by counting how often any two contributors edit the same page or section of a page. For example, using this approach, Iba et al. (2010) explored different roles of Wikipedia contributors and examined changes in their editing networks over time. Similar two-mode networks were also studied in social tagging systems such as "Delicious" and "CiteULike." For example, Liang et al. (2010) detected communities of users with shared interests based on the similarity of the items they tagged (in their case, primarily for the purpose of improving social recommender systems).

Communities in Virtual Worlds

The final type of Internet-based platforms that will be discussed here is a virtual world where people interact in a 3D virtual reality via self-created avatars. Some virtual worlds are designed primarily for socializing purposes like Second Life and "Habbo Hotel"; others are designed to support massively multiplayer online games (MMOGs), such as World of Warcraft (WoW) and Lineage II (successors of the original text-based and then 2D MMOGs). To discover social networks on these platforms, one approach is to build a so-called "co-location network" (a variation of a "coediting" network in Wikipedia communities and "name co-occurrence network" in online forums). A co-location network connects people based on how often and how long their avatars appear in close proximity within a virtual world (see, e.g., Ducheneaut et al. 2006). Most of these virtual worlds also allow users to chat via voice or text-based channels. Using information about "who talks to whom" and other users' interactions, similar to approaches applied to discussion groups (discussed above), a social network among users can be derived. For example, Suznjevic et al. (2009) tracked and connected "gamers" based on their recorded interactions of in-game activities such as *combat* or *trade*. The interactions were recorded with the help of an add-on built for the WoW.

Future Directions

Below are some of the emerging directions in this research area.

First, since most online communities still rely on text-based interactions, there is a need for new methods and research tools that would better integrate automated text analysis with SNA, either during the community discovery process or during the analysis and visualization of the resulting social networks. Some existing research tools are already moving in this direction. For example, Netlytic.org relies on text analysis to discover connections among community members using the "name network" approach (see subsection on "Communities Based on Online Discussion Forums" above). Also, during the network visualization process, a user can click on a connection between any two nodes in the network to explore the nature of any connection by reading the messages that were used to infer the connection. NodeXL is another popular system designed with online community researchers in mind. It can now automatically label densely connected groups of people (subcommunities) in the network based on the most used topics/hashtags found in their Twitter messages.

Second, most of the current community detection methods are ill-equipped to handle large-size, dynamic online social networks. Future work in this area will need to address issues such as computational complexity, the dynamic nature of online social networks, unknown numbers of communities in advance (and/or unknown sizes of communities), and cases when people belong to multiple communities or to none. Papadopoulos et al. (2011) discuss some recent trends in this area, in particular, how current community detection methods can be expanded to account for the dynamic nature of online communities. For a discussion of other emerging community detection methods, see sections on "Combining Link and Content for Community Detection" and "Dynamic Community Detection" in this Encyclopedia.

Third, the proliferation of mobile technologies such as smartphones and tablets and the boom in location-based online services such as "Foursquare" and "Facebook Places" are beginning to change the way people discover and manage their social connections, share content, and interact with their various offline and online communities. In doing so, these devices and services are also generating a massive amount of geo-tagged, time-stamped user data on an unprecedented scale. This data is a treasure trove for researchers interested in studying online communities, because it allows researcher to gain even better insights into the operation of online communities and especially the interplay between online and offline worlds. As a result of this trend, there is a growing interest in the research community to study location-based networks and compare them to other types of online networks as outlined in the above discussion (see section on "Location-Based Social Networks" in this Encyclopedia).

Finally, as more and more data about online communities and their users are being captured automatically, the issue of data privacy and security will become even more pressing. This is a major looming issue for researchers who rely on access to these data to study online communities. To address these concerns, researchers in this area need to work on developing best practices, revisiting existing policies and protocols, and creating new methods and tools for ethical data collection, analysis, and reporting of results in studies of online communities.

Cross-References

- ▸ Analysis and Mining of Tags, (Micro)Blogs, and Virtual Communities
- ▸ Community Evolution

▶ Community Detection, Current and Future Research Trends

▶ Community Identification in Dynamic and Complex Networks

▶ Connecting Communities

▶ Creating a Space for Collective Problem-Solving

▶ Dynamic Community Detection

▶ Flickr and Twitter Data Analysis

▶ Folksonomies

▶ Gaming and Virtual Worlds

▶ Geography and Web Communities

▶ Inferring Social Ties

▶ Mapping Online Social Media Networks

▶ Mining Trends in the Blogosphere

▶ Virtual Team

▶ Web Communities Versus Physical Communities

▶ Wikipedia Collaborative Networks

▶ Wikipedia Knowledge Community Modeling

References

Adamic L, Glance N (2005) The political blogosphere and the 2004 U.S. election: divided they blog 2nd annual workshop on the weblogging ecosystem: aggregation, analysis and dynamics. In: Proceedings of the 3rd international workshop on Link Discovery (LinkKDD '05). ACM, New York, 36–43. doi:10.1145/1134271.1134277

Ali-Hasan N, Adamic L (2007) Expressing social relationships on the blog through links and comments. In: Proceedings of international conference on weblogs and social media. Boulder

Anderson B (1983) Imagined communities: reflections on the origin and spread of nationalism. Verso, London

Blanchard A (2004) Blogs as virtual communities: identifying a sense of community in the julie/julia project. In: Gurak SAL, Johnson L, Ratliff C, Reyman J (eds) Into the blogosphere: rhetoric, community, and culture of weblogs. University of Minnesota, Minneapolis

Chung CJ, Gruzd A, Park HW (2010) Developing an e-research tool for humanities and social sciences: korean internet network miner on blogosphere. J Humanit (인문연구) 60(12):429–446. ISSN 1598–2211

Conover MD, Ratkiewicz J, Goncalves B, Francisco M, Flammini A, Menczer F (2011) Political polarization on twitter. In: Proceedings of the 5th international conference on weblogs and social media (ICWSM 2011), Barcelona, p 8

Ducheneaut N, Yee N, Nickell E, Moore RJ (2006) Alone together? In: Proceedings of the SIGCHI conference on human factors in computing systems. ACM, pp 407–416. doi:10.1145/1124772.1124834

Durant KT, McCray AT, Safran C (2010) Social network analysis of an online melanoma discussion group. AMIA Summits Transl Sci Proc Amia Summit Transl Sci 6–10

Efimova L, Hendrick S (2005) In search for a virtual settlement: an exploration of weblog community boundaries. In: Proceedings of the communities and technologies conference, Milano

Fiore AT, Tiernan SL, Smith MA (2002) Observed behavior and perceived value of authors in usenet newsgroups: bridging the gap. In: Proceedings of the SIGCHI conference on human factors in computing systems: changing our world, changing ourselves, Minneapolis. ACM, New York, pp 323–330

Gibson D, Kleinberg J, Raghavan P (1998) Inferring web communities from link topology. In: Proceedings of the 9th ACM conference on hypertext and hypermedia. ACM, New York, pp 225–234. doi:10.1145/276627.276652

Gruzd A (2009) Studying collaborative learning using name networks. J Educ Libr Inf Sci 50(4):243–253

Gruzd A, Sedo DR (2012) #1b1t: investigating reading practices at the turn of the twenty-first century. J Stud Book Cult (Special issue on New Studies in the History of Reading) 3(2):34–42 doi:10.7202/1009347ar

Gruzd A, Wellman B, Takhteyev Y (2011) Imagining twitter as an imagined community. Am Behav Sci, special issue on imagined communities 55(10): 1294–1318. doi:10.1177/0002764211409378

Gruzd A, Black FA, Le Y, Amos K (2012) Investigating biomedical research literature in the blogosphere: a case study of diabetes and hba1c. J Med Libr Assoc 100(1). doi:10.3163/1536-5050.100.1.007

Harrison R (2009) Excavating second life: cyber-archaeologies, heritage and virtual communities. J Mater Cult 14(1):75–106. doi:10.1177/1359183508100009

Haythornthwaite C, Wellman B (1998) Work, friendship, and media use for information exchange in a networked organization. J Am Soc Inf Sci 49(12):1101–1114. doi:10.1002/(SICI)1097-4571(1998) 49:12<1101::AID-ASI6>3.0.CO;2-Z

Hogan B (2008) Analyzing social networks via the internet, In: Fielding N, Lee RM, Blank G (eds) Sage handbook of online research methods. Sage, Thousand Oaks, pp 141–160

Iba T, Nemoto K, Peters B, Gloor PA (2010) Analyzing the creative editing behavior of wikipedia editors. Procedia 2(4):6441–6456. doi:10.1016/j.sbspro.2010.04.054

Iriberri A, Leroy G (2009) A life-cycle perspective on online community success. ACM Comput Surv 41(2):1–29. doi:10.1145/1459352.1459356

Jones Q (1997) Virtual communities, virtual settlements and cyber-archaeology. J Comput Mediat Commun 3(3)

Klausen J, Barbieri E, Reichlin-Melnick A, Zelin A (2012) The youtube jihadists: a social network analysis of al-muhajiroun's propaganda campaign. Perspect Terror 6(1). Retrieved from http://www.terrorismanalysts.com/pt/index.php/pot/article/view/klausen-et-al-youtube-jihadists

Liang H, Xu Y, Li Y, Nayak R, Tao X (2010) Connecting users and items with weighted tags for personalized item recommendations. In: Proceedings of the 21st ACM conference on hypertext and hypermedia. ACM, pp 51–60. doi:10.1145/1810617.1810628

Marvin L (1995) Spoof, spam, lurk and lag: the aesthetics of text-based virtual realities. J Comput-Mediat Commun 1(2)

McMillan DW, Chavis DM (1986) Sense of community: a definition and theory. J Community Psychol 14(1): 6–23

Mockus A, Fielding RT, Andersen H (2002) Two case studies of open source software development: apache and mozilla. ACM Trans Softw Eng Methodol 11(3):309–346

Nonnecke B, Preece J (2000) Lurker demographics. In: Proceedings of the SIGCHI conference on human factors in computing systems. ACM, pp 73–80. doi:10.1145/332040.332409

Otte E, Rousseau R (2002) Social network analysis: a powerful strategy, also for the information sciences. J Inf Sci 28(6):441–453. doi:10.1177/016555150202800601

Papadopoulos S, Kompatsiaris Y, Vakali A, Spyridonos P (2011) community detection in social media. Data Min Knowl Discov 24(3):515–554. doi:10.1007/s10618-011-0224-z

Pikas CK (2008) Detecting communities in science blogs. In: Proceedings of the 2008 4th IEEE international conference on e-science. IEEE, pp 95–102. doi:10.1109/eScience.2008.30

Porter CE (2004) A typology of virtual communities: a multi-disciplinary foundation for future research. J Comput-Mediat Commun 10(1)

Quan-Haase A, Wellman B (2005) Local virtuality in an organization: implications for community of practice. In: Besselaar P, Michelis G, Preece J, Simone C (eds) Communities and technologies 2005, Milano. Springer, Heidelberg/Berlin, pp 215–238

Sen S, Lam SK, Rashid AM, Cosley D, Frankowski D, Osterhouse J, Harper FM et al (2006) Tagging, communities, vocabulary, evolution. In: Proceedings of the 2006 20th anniversary conference on computer supported cooperative work (CSCW '06). ACM, New York, pp 181–190. doi:10.1145/1180875.1180904

Stanoevska-Slabeva K, Schmid BF (2001) A typology of online communities and community supporting platforms. In: Proceedings of the 34th annual hawaii international conference on system sciences. IEEE Comput Soc 10. doi:10.1109/HICSS.2001.927041

Suznjevic M, Dobrijevic O, Matijasevic M (2009) Hack, slash, and chat: a study of players' behavior and communication in MMORPGs. In: Proceedings of the 8th annual workshop on network and systems support for games, Paris. IEEE, Piscataway, Article 2, p 6

Takhteyev Y, Gruzd A, Wellman B (2012) Geography of twitter networks. Soc Netw 34(1):73–81. doi:10.1016/j.socnet.2011.05.006

Wellman B (2001) Physical place and cyberplace: the rise of personalized networking. Int J Urban Reg Res 25(2):227–252. doi:10.1111/1468-2427.00309

Wellman B, Leighton B (1979) Networks, neighborhoods and communities. Urban Aff Quart 14:363–90

Wilson RE, Gosling SD, Graham LT (2012) A review of facebook research in the social sciences. Perspect Psychol Sci 7(3):203–220. doi:10.1177/1745691612442904

Zhu Y (2010) Measurement and analysis of an online content voting network. In: Proceedings of the 19th international conference on World Wide Web. ACM, p 1039. doi:10.1145/1772690.1772796

Recommended Reading

Adamic LA, Zhang J, Bakshy E, Ackerman MS (2008) Knowledge sharing and yahoo answers. In: Proceeding of the 17th international conference on world wide web. ACM, p 665. doi:10.1145/1367497.1367587

Blanchard A, Markus L (2004) The experienced "sense" of a virtual community: characteristics and processes. SIGMIS Database 35(1):64–79

Boyd DM, Ellison NB (2007) Social network sites: definition, history, and scholarship. J Comput-Mediat Commun 13(1)

Brooks B, Welser HT, Hogan B, Titsworth S (2011) Socioeconomic status updates. Inf Commun Soc 14(4):529–549. doi:10.1080/1369118X.2011.562221

Butler BS (2001) Membership size, communication activity, and sustainability: a resource-based model of online social structures. Inf Syst Res 12(4):346–362. doi:10.1287/isre.12.4.346.9703

Chin A, Chignell M (2007) Identifying communities in blogs: roles for social network analysis and survey instruments. Int J Web Based Communities 3(3): 343–365

Cho H, Gay G, Davidson B, Ingraffea A (2007) Social networks, communication styles, and learning performance in a CSCL community. Comput Educ 49(2):309–329. doi:10.1016/j.compedu.2005.07.003

Dean J, Henzinger MR (1999) Finding related pages in the world wide web. Comput Netw 31(11–16):1467–1479. doi:10.1016/S1389-1286(99)00022-5

Ferrari L, Rosi A, Mamei M, Zambonelli F (2011) Extracting urban patterns from location-based social networks. In: Proceedings of the 3rd ACM SIGSPATIAL international workshop on location-based social networks. ACM, New York, pp 9–16. doi:10.1145/2063212.2063226

Gjoka M, Kurant M, Butts CT, Markopoulou A (2011) Practical recommendations on crawling online social networks. IEEE J Sel Areas Commun 29(9):1872–1892. doi:10.1109/JSAC.2011.111011

Golder SA, Bernardo A (2006) Usage patterns of collaborative tagging systems. J Inf Sci 32(2):198–208. doi:10.1177/0165551506062337

Gopal S (2007) The evolving social geography of blogs. In: Miller HJ (ed) Societies and cities in the age of instant access, vol 88. Springer, Dordrecht, pp 275–293

Gruzd A (2009b) Automated discovery of emerging online communities among blog readers: a case study of a canadian real estate blog. In: Proceedings of the internet research 10.0 conference, Milwaukee, 7–11 Oct 2009

Haythornthwaite C (2002) Strong, weak, and latent ties and the impact of new media. Inf Soc 18(5):385–401. doi:10.1080/01972240290108195

Haythornthwaite C, Gruzd A (2007) A noun phrase analysis tool for mining online community. In: Steinfield C, Pentland B, Ackerman M, Contractor N (eds) Communities and technologies 2007: proceedings of the third communities and technologies conference, Michigan State University 2007. Springer, pp 67–86. doi:10.1007/978-1-84628-905-7_4

Haythornthwaite C, Gruzd A (2008) Analyzing networked learning texts. In: Proceedings of networked learning conference, Halkidiki, 5–6 May 2008, pp 136–143

Hogan B (2010) Analyzing facebook networks. In: Hansen D, Smith M, Shneiderman B (eds) Analyzing social media networks with NodeXL. Morgan Kaufman, New York

Koh J, Kim YG (2004) Knowledge sharing in virtual communities: an e-business perspective. Expert Syst Appl 26(2):155–166

Korfiatis NT, Poulos M, Bokos G (2006) Evaluating authoritative sources using social networks: an insight from wikipedia. Online Inf Rev 30(3):252–262. doi:10.1108/14684520610675780

Kumar R, Novak J, Tomkins A (2010) Structure and evolution of online social networks. In: Yu PS, Han J, Faloutsos C (eds) Link mining: models, algorithms, and applications. Springer, New York, pp 337–357. doi:10.1007/978-1-4419-6515-8_13

Nazir A, Raza S, Chuah CN (2008) Unveiling facebook. In: Proceedings of the 8th ACM SIGCOMM conference on internet measurement. ACM, New York. doi:10.1145/1452520.1452527

Priedhorsky R, Chen J, Lam S, Tony K, Panciera K, Terveen L, Riedl J (2007) Creating, destroying, and restoring value in wikipedia. In: Proceedings of the 2007 international ACM conference on supporting group work (GROUP '07). ACM, New York, pp 259–268. doi:10.1145/1316624.1316663

Reyes P, Tchounikine P (2003) Supporting emergence of threaded learning conversations through augmenting interactional and sequential coherence. In: Proceedings of the international conference on computer supported collaborative learning, Bergen, pp 83–92

Scellato S, Noulas A, Mascolo C (2011) Exploiting place features in link prediction on location-based social networks. In: Proceedings of the 17th ACM SIGKDD international conference on knowledge discovery and data mining. ACM, pp 1046–1054. doi:10.1145/2020408.2020575

Schenkel R, Crecelius T, Kacimi M, Michel S, Neumann T, Parreira JX, Weikum G (2008) Efficient top-k querying over social-tagging networks. In: Proceedings of the 31st annual international ACM SIGIR conference on research and development in information retrieval. ACM, pp 523–530. doi:10.1145/1390334.1390424

Schneider F, Feldmann A, Krishnamurthy B, Willinger W (2009) Proceedings of the 9th ACM SIGCOMM conference on internet measurement conference (IMC '09). ACM, New York, pp 35–48. doi:10.1145/1644893.1644899

Warmbrodt J, Sheng H, Hall R (2008) Social network analysis of video bloggers' community. In: Proceedings of the 41st annual hawaii international conference on system sciences. IEEE, pp 291–291. doi:10.1109/HICSS.2008.402

Whittaker S, Terveen L, Hill W, Cherny L (1998) The dynamics of mass interaction. In: Proceedings of the 1998 ACM conference on computer supported cooperative work. ACM, pp 257–264. doi:10.1145/289444.289500

Zimmer M (2010) But the data is already public': on the ethics of research in facebook. Ethics Inf Technol 12(4):313–325. doi:10.1007/s10676-010-9227-5

Online Communities, Constellation of Communities

▸ Networks of Practice

Online Community

▸ Gaming and Virtual Worlds
▸ Wikipedia Knowledge Community Modeling

Online Forums

▶ Topology of Online Social Networks

Online Group

▶ Online Communities

Online Health

▶ Online Healthcare Management

Online Healthcare Management

Prasanna Desikan[1], Aarti Sathyanarayana[2], and Jaideep Srivastava[2]
[1]Division of Applied Research, Allina Health Hospitals and Clinics, Minneapolis, MN, USA
[2]Department of Computer Science and Engineering, University of Minnesota, Minneapolis, MN, USA

Synonyms

Consumer health informatics; Cyber medicine; e-Health; Online health

Glossary

eICU Electronic intensive care units

Definition

Online healthcare management refers to all activities that facilitate healthcare delivery and management by means of electronic data transmission, storage, and retrieval, primarily through the Internet in a private and secure fashion.

Introduction

Over the last couple of decades, the Internet has created a paradigm shift in how businesses operate. It has enabled the development of myriad technologies, such as Web services and enterprise server applications, and as a result industries increasingly use the Internet as a medium of secure data transmission and communication – it has become an essential component of most businesses. Even as most industries leverage the benefits of the Internet, a majority of healthcare delivery today is still carried out in the traditional brick-and-mortar infrastructure. This is primarily because healthcare delivery largely relies on face-to-face interaction between a patient and a clinician. The healthcare ecosystem, however, has been adopting more ways of creating, using, and sharing digital health data, with the Internet increasingly offering a reliable means of private, secure, and reliable data transmission. The advancements in Internet and Web technologies have also made it possible to share data remotely across geographical regions. This has facilitated care delivery using Internet-based technologies. Various terminologies for such healthcare practice have been proposed and studied in the past decade (Anderson et al. 2003; Eysenbach 2001; Ball et al. 2007; Forkner-Dunn 2003). Examples include eHealth, online health, cyber medicine, and consumer health informatics. The different terminologies encompass a wide variety of systems, from any form of electronic healthcare data storage, processing, and communication to a very specific focus, such as eICUs.

These systems have proven to enhance patient care practices, cut costs, and improve workflow (Halamka 2006). There are a number of options for practices that wish to acquire an online healthcare management system. Larger institutions often build their own systems to allow for extensive customization, but there are a number of other options available. Institutions can choose from on-premise solutions that require the purchase and licensing of software or a software as a service (SaaS) solution that pays a third party to handle data storage and software maintenance (Intuit 2013). There are a number of open-source options available as well, such as OpenEMR. Most healthcare systems that use online technologies serve two purposes – knowledge sharing and care delivery. The impact of Internet and online technologies on the health of various populations (Hesse et al. 2005; Kummervold et al. 2008) has shown the promise for success of these systems in the future.

This entry focuses on online healthcare delivery and online social interactions among clinicians and patients. Online healthcare systems that deal with other applications, such as pharmacy benefit management, provider education systems, and health insurance applications, are beyond the scope of this entry. The goal is to understand how online technologies have been used for advancing care delivery and how knowledge sharing improves provider care delivery and enhances patient satisfaction. In the following sections, we provide an overview, followed by a classification of online systems based on type of interaction. We then illustrate the concepts by describing key applications, which have been chosen purely to illustrate the concepts and not by any preferential criteria. We then describe the current challenges in these systems and point out how data mining and social network analysis can aid the future directions for such systems. We refer the readers to Tan et al. (2006) for an introduction to data mining and to Valente (2010) for an overview of how social network analysis can be applied to healthcare.

Overview

The success of these systems is due to the ease with which information can be stored, accessed, and shared securely. The lack of need for physical resources also drastically reduces the cost, making this an attractive option. In a typical care scenario, information is exchanged between clinician and patient. This exchange takes two forms – verbal communication and physical examination. The communication allows the clinician to gather the required information, process it, and then apply his or her medical expertise to diagnose and treat the patient. Typically, for most conditions, the best care is delivered when the clinician also physically examines the patient. However, for certain acute conditions such as the common cold, sinusitis, and fever, the initial screening and treatment do not require that the patient be physically present, and the basic information provided can help the physician diagnose the condition and provide appropriate treatment.

The Internet is a communication medium that minimizes cost because it doesn’t require the people communicating to physically be in each other’s presence. It allows for synchronous as well as asynchronous communication. Efficient integration of knowledge across various sources is possible. Streaming of audio and video and interactive flash and HTML5 technologies have become commonplace. These enabling technologies help to avoid the necessity of a patient being physically present for any condition that can be diagnosed by verbal communication. Technology advancements, such as EHR systems, allow data about a patient to be collected and presented in a fashion that enables the clinician to make clinical decisions. The clinician can either synchronously or asynchronously interpret the data and provide appropriate care advice. Such online communication and knowledge sharing provides an opportunity for remote patient care through the Web. Currently the advancements in medical device technologies also enable the measurement of patient data such as blood pressure, weight, and heart rate and securely transmit the data to a remote server, which then enables the clinician to monitor the patient (LATITUDE). As more such technologies are developed, the need for a patient to be physically present decreases. The use of online healthcare delivery and management is thus growing in scope and popularity.

Another growing area of online healthcare management is healthcare-based social interaction among patients and clinicians. As the popularity of Internet-based social networks and professional networks increases, topic-specific networks are becoming popular (Greene et al. 2011; Lorig et al. 2010). Many educational healthcare websites have traditionally provided forums to allow users to interact. However, social networking has allowed more focused groups to interact and share information. Networks geared towards, respectively, patients and clinicians allow these players to share information within their specific groups.

Types of Online healthcare Systems

Online healthcare systems that are primarily focused on care involve two key players – the clinician and the patient. These systems rely for their effectiveness on the underlying communication between these two players. These systems can thus be categorized into three main categories, according to the type of interaction: Clinician to Patient, Clinician to Clinician, and Patient to Patient.

Clinician to Patient

Among the incentives for medical practices to invest in online healthcare management systems is the move towards patient-centric care. Patient-centric care encourages a collaborative approach between physician and patient to make clinical decisions together.

There has been a strong push in the medical field towards this approach, because when patients are given a more active role in their healthcare decisions, they feel more accountable for monitoring their self-care, which ultimately improves their overall health and breaks down barriers between patient and caregiver (Tang et al. 2006). Studies have also shown that establishing a more patient-centric approach can cut costs as well as optimize a physician's time (Bertakis and Azari 2011).

An online patient portal allows patients to log in and see their own personal health records (PHRs). Information such as allergies, immunizations, treatment recommendations, and lab test results can all be easily stored electronically (Zaroukian 2008).

The ease of information exchange online will likely decrease the number of calls made to a healthcare provider's support center. Patients can directly access their personal healthcare records, which reduces the risk of miscommunication errors over the phone. Additionally, instead of patients calling the office to set up an appointment, they can do so online by checking their physician's availability. There are also financial management resources that can be a component of the system. This allows patients to pay their bills online, which is quicker and more convenient for both healthcare provider and patient (Ball et al. 2007). Giving patients direct access to their healthcare records also gives them instant access to their lab test results. This reduces communication time, which is particularly important for critical tests. Most such systems focus on delivering information to the patient after a face-to-face clinic visit.

The natural next step is to reduce the amount of face-to-face interaction and help move some of the diagnostic activities to the online system. It has been shown that this online interactivity between clinician and patient can greatly increase physician productivity, allowing doctors to see 11.1 % more patients and saving almost $100 a day (Wallwiener et al. 2009). Many physicians are concerned that by opening up to online communication between doctors and patients, doctors will be flooded with questions that would not otherwise be asked. Studies, however, have shown that only 57 % of the communication sent by patients required a physician's response (Zaroukian 2008). Administrative staff can respond to many of the communications, resulting in a large time-saving. Rather than discrete, episodic interaction between physician and patient, the online portal creates a continuous channel of communication (Tang et al. 2006). The channel provides continuity of care regardless of location and time of day, because all the information can be accessed from anywhere with an Internet connection (Nalari Health).

A patient's overall health can improve from the utilization of an online healthcare management system. When patients have greater access to their own records, they feel more accountable for their well-being and are more conscious of their health (Tang et al. 2006). In turn, patients make more informed decisions, and patient compliance with physicians' orders increases (Ball et al. 2007). For example, if a patient's glucometer directly interfaces with her online healthcare record, she (he) can be warned when her (his) levels are too high and action can be taken. Moreover, doctors can be alerted of her level fluctuation and can adjust her diabetes control (Ball et al. 2007). This can promote earlier interventions (Tang et al. 2006).

By establishing an online interactive relationship between clinicians and their patients, hospitals and clinics can deliver general information and health education so that it is readily accessible to and reusable by multiple patients, for example, posting answers to frequently asked questions for all patients to see or information about new treatment options for patients with the same disease. This is particularly useful for communicating with patients who suffer from chronic illnesses (Ball et al. 2007; Tang et al. 2006). Doctors can save time by providing patients with instant access to information at all hours, day or night. Both parties benefit.

Many doctor-to-patient systems allow patients to give access to their family members or caregivers (Ball et al. 2007). This also improves continuity of care. Patients can print out their own records and give them to any new doctor they are seeing (Medicare 2013). Family members can access critical information in emergency situations, such as allergies or current medications (Medicare).

Another advantage is that a patient's medical history is securely stored in one place and is readily accessible. As a result, patients are more likely to have complete records over time, which will help physicians to make more accurate clinical decisions. This large amount of data can easily be searched online, which provides a huge advantage over handwritten histories (Ball et al. 2007).

While patient portals are a key component of online healthcare management systems, there are other aspects that are equally beneficial. All patient medication is tracked within the system, allowing alerts to be sent to both patient and doctor when a prescription needs to be refilled, thus reducing the chance of a patient being without necessary medication for a period of time. Moreover, the system will allow physicians to send prescriptions online. This saves time because it removes the need for superfluous appointments to be made for the sole purpose of refilling a prescription (Zaroukian 2008). It also allows patients to have their prescriptions refilled more quickly since they don't have to wait for an appointment. Many online healthcare management systems, such as the one provided by Medicare, allow for doctors to send the prescription directly to the pharmacist (Medicare). This simplifies the process for the doctor and avoids errors that occur from illegible handwriting. It also allows pharmacists extra time to check patients' insurance contributions and research cheaper, generic alternatives. It also saves time for the patient, who doesn't have to wait while the prescription is filled.

Similarly, referrals can be handled online. This has many of the same benefits as online prescriptions. Additionally, all the relevant information about a patient that a specialist needs is readily available online (Medicare). Specialists within the same hospital can access the information directly, and specialists at other institutions can receive a copy from the patient. This saves time and money for both parties, since redundant tests are avoided.

Patient to Patient

Another approach to providing patient-centric care is creating a forum for patients to communicate with other patients who have similar illnesses. The forum allows patients to share their experiences about the progression of the disease, their symptoms, and their overall health status (PatientsLikeMe).

By connecting with others who are experiencing the same conditions, patients can be more informed and have a better understanding of their health. They can get a realistic picture of what it's like to live with their disease, based on the experiences of others. They can also gain a clearer picture of the disease progression timeline and know what symptoms to look out for as warning signs of escalation.

Patients can also track their own progress by monitoring their test results and charting the progression. Often the system makes recommendations of questions to ask the patient's doctor. This allows the patient to feel more in control during a time when they may be feeling overwhelmed and powerless.

In online forums, patients have the opportunity to make friends with others who are experiencing the same lifestyle changes and overcoming the same difficulties. This is a particularly huge benefit for patients in more isolated areas or with more unique diseases, because they would otherwise not have the chance to interact with others with the same condition. These new friendships give patients the opportunity to freely discuss their concerns and experiences without judgement and with others who understand what they are going through. This gives patients a sense of community and a sense of belonging. As a result, they are more likely to open up about their disease. The more that patients discuss the details of their experience, the more they will be able to accept it, and the less likely they will become depressed. Knowing other patients who have endured through the difficult times helps newly diagnosed individuals to see the light at the end of the tunnel and keep their illness in perspective.

Another benefit of systems that connect similar patients to each other is that patients can find out about other treatment options. They can find out how popular a particular medication is, how successful it has been, whether other patients recommend it, and what the side effects are. They can also get recommendations from other patients about which specialists to see.

Finding information on a disease used to be a troublesome task for many patients. They would have to troll the Internet for hours before finding what they were looking for. Online healthcare systems allow patients to find all their information in one space that is tailored for easy access. Research articles are often posted to keep patients up to date, including clinical trials that a patient might be eligible for.

These systems also provide an overall benefit to the healthcare research community. All of the data provided by members is stored in a repository. This information could be essential in assisting future research (PatientsLikeMe).

Clinician to Clinician

While systems for clinician-to-patient interaction are gaining in popularity, the clinician-to-clinician online communication channels have not yet been as widely adopted. However, there are currently a number of social networking options available for physicians to interact with each other.

The purpose of these systems is to expand the professional networks of doctors. For example, many of these sites encourage primary care physicians to connect with specialists. This allows primary care physicians to find specialists who are qualified to treat their patients. It also allows specialists to open themselves to new referrals and increase the number of patients they are seeing. Physicians can post their specialties on a profile in a similar fashion to the way that people post their career history on LinkedIn (LinkedIn). This also allows the site to recommend literature based on a clinician's specialty.

The systems encourage physicians to post blog entries about their current cases, ongoing research, new approaches, and treatment plans they've tried. This supplements the news that the site itself posts. In addition to creating a pool of current and updated shared knowledge, the sites also post educational tools to keep doctors up to date. Some sites provide access to original medical content to their members.

Many of these sites contain forums for doctors to post questions to be answered by their peers. The forums can be open to all healthcare professionals or limited so that only other physicians can access them. Some even offer the opportunity for doctors to pose their questions directly to an expert. There's also the capability to allow public as well as private messaging. All communication, including images, is HIPAA compliant. This is particularly useful because often doctors can only fax information because email is not secure.

Key Applications

Clinician to Patient

As Web technologies have enabled a secure and private means of exchanging medical information between provider and patient, the need for a visit to a clinic for certain medical conditions has been reduced. Many organizations have started to provide online care delivery for minor conditions (E-Visits; Virtuwell; Zipnosis). These online systems help in diagnosing, treating, and providing prescriptions for common ailments such as the cold and flu, diarrhea, heartburn, pink eye, sinus problems, or urinary tract infections. Systems vary as to the set of conditions they cover.

Figure 1 illustrates the typical flow of interaction between a clinician and a patient in an online setting. The key idea is that patients can answer a set of questions via an online form or an interactive session and leave the provider a message about their symptoms. Using this information, a provider will reply within a specified period of time. Treatment options typically include one or all of instructions for self-care, over-the-counter medications or prescriptions (if required), and instructions for an observational period to determine if the patient needs to pay an actual visit to the clinic. Most online systems charge a fee. These asynchronous systems allow for greater coverage of care as the necessity to match the patient's schedule to the physician's is decreased or eliminated.

Patient to Patient

A key archetype of a patient-to-patient forum is the website PatientsLikeMe.com (PatientsLikeMe).

The goal behind the creation of PatientsLikeMe.com is to create open and honest communities where patients share their experiences, advice, data, and stories. Members can learn about their disease, understand it better, and learn how to manage their symptoms so the disease has the smallest possible impact on their life. Most importantly, they can see the real-life outcomes of other patients and lose the sense of isolation and loneliness that living with a disease can bring. The strength of this type of forum is that patients who are in a similar situation have a stronger connection and communicate better than those who are dealing with different issues (See Fig. 2).

The site hopes to answer patients' questions, such as "Given my status, what is the best outcome I can hope to achieve, and how do I get there?" Each patient is asked about their ability to walk, breathe, use their hands, speak, feel happiness, and be productive. The website then converts this information into a structure that is computable. This data can contribute in two ways. It can contribute to answering questions that patients have about others with their illness, and it can be analyzed to give key insights that can lead to beneficial changes in the way healthcare is provided.

Clinician to Clinician

While systems for doctor-to-patient interaction are gaining in popularity, the doctor-to-doctor communication channels have not yet been as heavily adopted. There are currently a number of social networking options available for doctors to interact with each other. A few popular options are Medscape Connect, from the creators of WebMD; doc2doc, from the BMJ group; and Doximity, from the creators of LinkedIn (Medscape, doc2doc, Doximity). Figure 3 describes the typical elements and flow of interactions between clinicians in an online setting.

Medscape Connect claims to be the largest physician-only discussion community. It is a product of WebMD's health-professional network and allows health professionals, physicians, and specialists access to original medical content, educational tools, and a weekly specialty newsletter. The site also has an "Experts Corner" where physicians can pose questions to experts and a discussion forum where colleagues can communicate freely.

doc2doc is a global professional networking community that connects health professionals. It has many of the same features as Medscape Connect, but does not publish its own journal or newsletter, and does not have a venue for

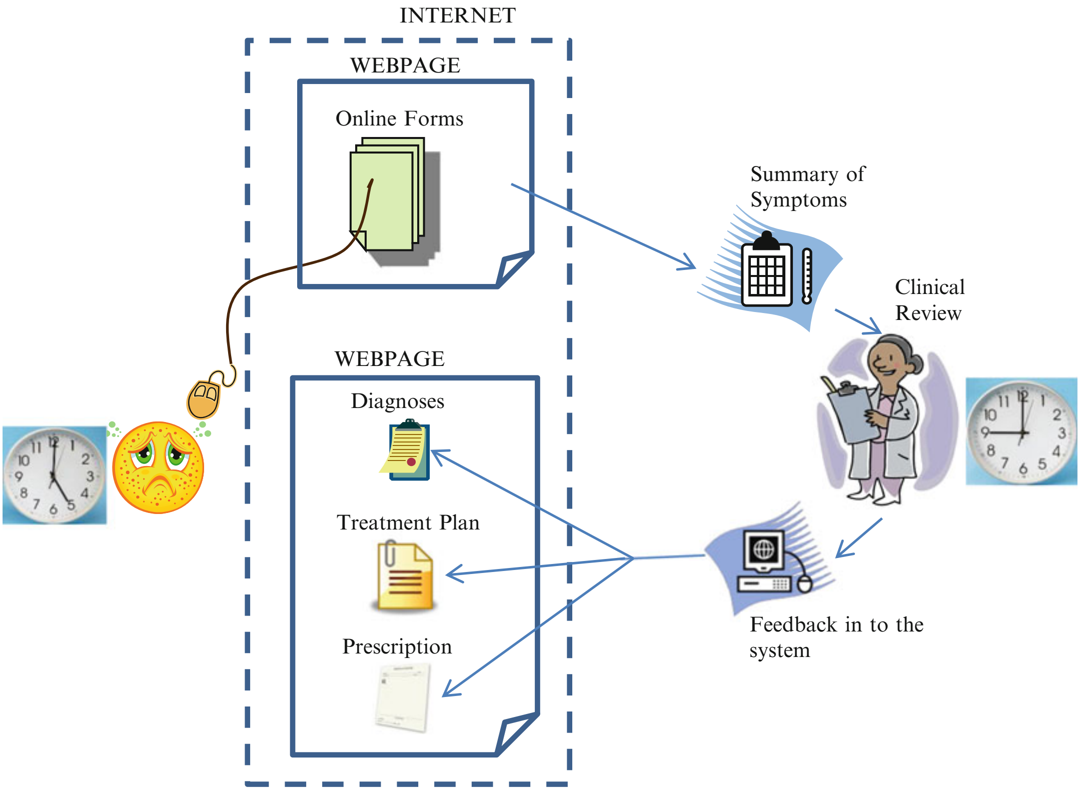

Online Healthcare Management, Fig. 1 Patient-to-clinician interaction in an online system

physicians to pose questions directly to an expert. It does have forums, including special closed clinical forums for physician-only interaction. It also allows public and private messaging between members and allows for any user to create blogs for others to read and follow.

Doximity currently has a membership of over 707,000 physicians in the USA, who work in over 87 specialties. It also has similar features to Medscape Connect and doc2doc. Doximity offers forums, including physician-only discussions, and recommends blog posts and other literature according to what other physicians have read. The site is created by the founders of LinkedIn, and a number of its unique features are closely related to the business networking site. Physicians create a profile that contains information drawn from their curriculum vitae. They can specify their clinical and research interests and network with others to explore job opportunities. All the information transferred, including images, is HIPAA compliant.

A feature that Doximity specifically encourages, and which future up-and-coming physician-centric social networks will likely include, is the expansion of a member's social network. The site offers free tools for doctors to communicate with each other easily. It also sends out a monthly update of new primary care physicians or specialists who are in the area.

Key Challenges

Launching and maintaining an online healthcare management system comes with many challenges. We highlight some of the key challenges that are currently being faced. A large concern for patients is the security and privacy of their personal healthcare information. Integrating clinical information with online systems is a

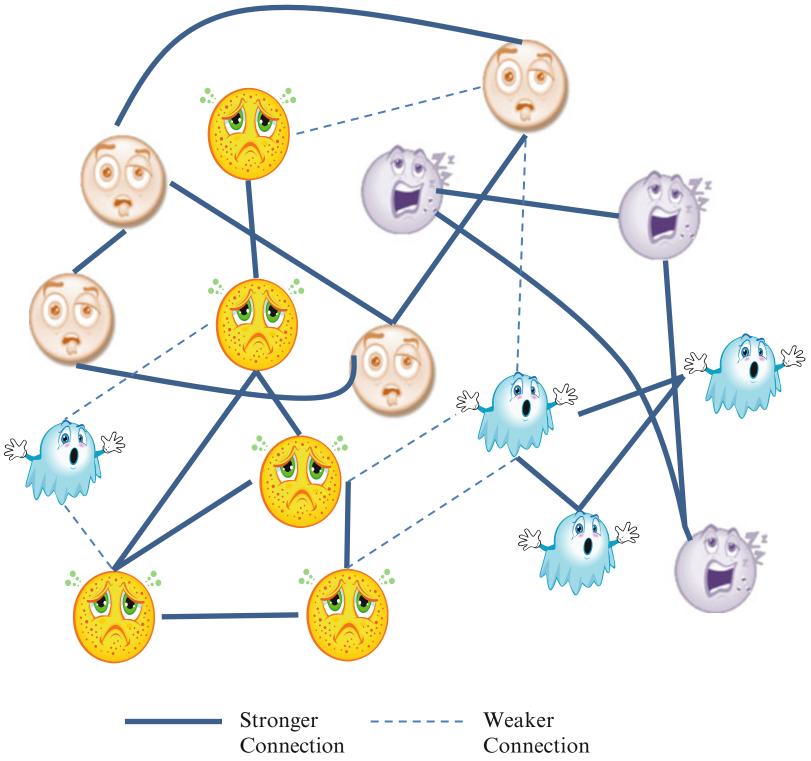

Online Healthcare Management, Fig. 2 Illustration of the formation of patient communities in patient-to-patient interaction

big challenge due to privacy concerns and the necessity for HIPAA compliance. Although many types of online healthcare systems are HIPAA compliant, not all are. Also, creating an effective online healthcare management system depends entirely on how well it is adopted. A system needs to capture all the information that physicians find important, without slowing the workflow process. It also needs to have a user-friendly portal that patients can and want to use. While privacy will always be a big concern when data is transmitted over the Internet, we point out some of the gaps in existing technologies.

Lack of Personalized Care: There are a large number of clinical factors which are monitored for a wide variety of patients. However, each individual patient is different and is likely to be concerned only with a small subset of these factors. Critical information for one patient might be matched with information from other patients with similar profile. Current online systems primarily focus on asynchronous communication, starting with a generic questionnaire. However, the problem with form-filling and asynchronous communication is the assumption that forms contain all the appropriate questions to ask. The personalization of care is lost. An intelligent interactive application can truly simulate a face-to-face appointment to narrow down the information gathering to personalized questions. There is a scarcity of intelligent online systems. Most systems currently lack systematic data collection and advanced analysis to look for patterns to enhance care delivery.

Lack of Social Connection: Another challenge of switching to a predominantly online relationship between clinicians and patients is the lack of personal face time. A face-to-face conversation provides a greater opportunity to socially connect than current online systems and usually translates to a strong and trusting relationship between patient and doctor. There are a number of

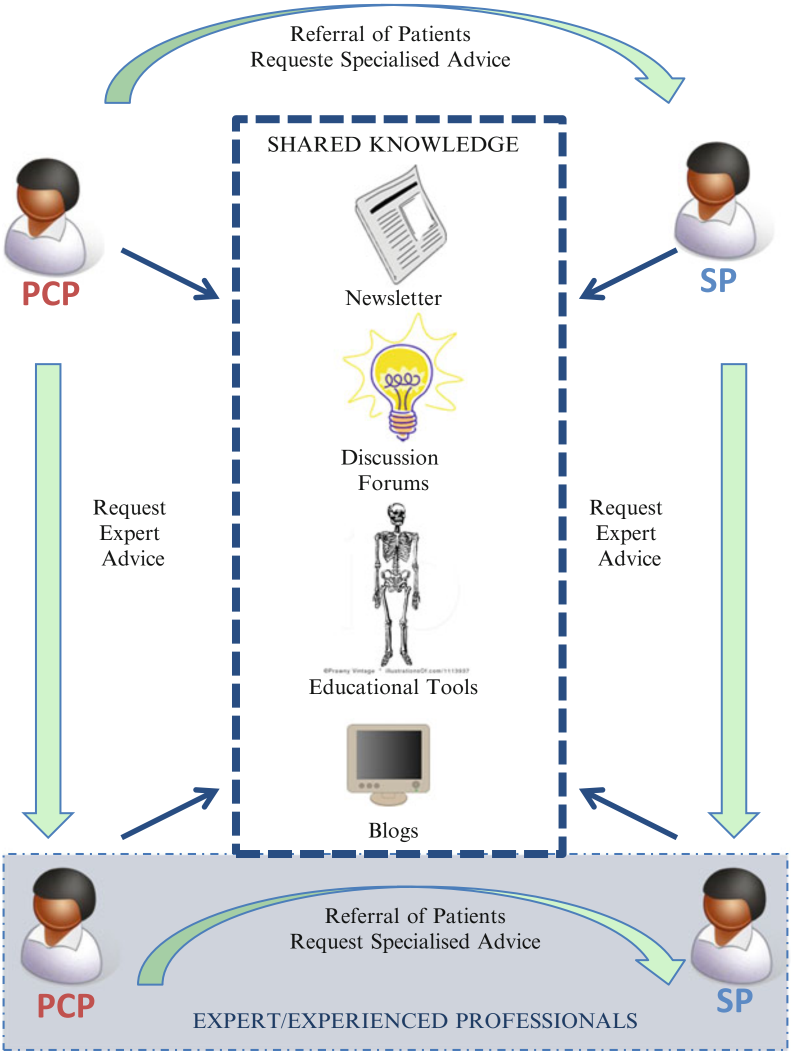

Online Healthcare Management, Fig. 3 Illustration of the interactions between clinicians in a clinician-to-clinician online healthcare management system

patient-to-patient networks and clinician-to-clinician professional networks. There also exist forums where clinicians answer patient questions. However, there is a lack of online networks that aid in the social connection between provider and patient. The social connection plays an important role in longitudinal consistency of care for patients and increases patient satisfaction.

Existing as Silos: There are a huge number of innovative online systems that exist today. Most of these systems focus on a certain specific problem and a certain specific market. However, it is important to consider the complex nature of the healthcare ecosystem and the impact of the whole ecosystem on an individual's health. As providers, patients, payors, and other players build their online systems, it is important to provide a mechanism by which these can talk to each other within the limits of patient privacy. Lack of data sharing and knowledge sharing does not allow these systems to be as productive or efficient as they potentially can be.

Looking Ahead: The future of Online Healthcare

As online healthcare systems expand in scope with newer enabling technologies, they also serve as a great data collection framework, not just from a clinical data perspective but also from a behavior-analysis perspective. Because the technology is in its relative infancy, there is huge scope for growth in a multitude of directions, such as expansion of online services, more intuitive systems, and intelligent systems that focus on core goals such as improved care, continuity of care, personalized care, and affordable care. Potentially, online healthcare technology could be driven towards achieving the "triple aim." While covering the entire range of future directions is beyond the scope of this entry, we present three directions which in the short term will lay a path to open up a wider variety of options. HIPAA and other privacy policies will greatly impact the feasibility of some of the proposed directions. This is something that is best handled on a case-by-case basis, as technologies move forward.

Intelligent Systems: Online technologies are great enablers of explicit and implicit data collection from users. Online healthcare systems provide this unique opportunity to collect data, analyze it, and provide recommendations to users. For example, collecting data about patient preferences through explicit online surveys or implicit browsing patterns can help provide better insights into the kind of care best suited to the patient, especially in cases where there is no clear evidence that one treatment is better than the alternatives and the patient has a choice of treatment, such as in the case of treatments for uterine fibroids. Another example is CogCubed (Roots et al. 2013), which develops cognitive games for education and healthcare. By analyzing the data collected from a number of cognitive games on digital cubes containing accelerometers, the company diagnoses children with psychological disorders such as attention deficit disorder and autism. In both cases, data mining enhances the current processes and decision-making abilities. Clinical decision-making can thus be enhanced by data mining and predictive modelling techniques (Tan et al.).

Social Network Analysis: Social network analysis has found applications in the healthcare domain for the study of problems such as disease contagion, smoking cessation, and physician collaboration (Valente 2010). On social networking sites such as PatientsLikeMe, there is a wealth of information from a public health perspective that can enable identification of communities of health, players with high influence, etc. We believe that as online technologies take shape and data collection becomes easier, two areas of social network analysis will come to the fore as key research areas. The first is the temporal evolution of these social graphs and the influence of several factors in social interaction; beliefs and knowledge sharing will be of great interest. The longitudinal nature of the analysis also provides ample scope for building interesting prediction models for different kinds of social contagion. The second area is the clinician-to-clinician interaction within a healthcare organization. While many networks exist for clinicians to interact outside their work for professional reasons, the mechanisms for social and informational interaction *within* an organization are scarce. We believe online technologies will reduce the necessity to be physically present for such interactions. The value of such interactions among clinicians for improving care is a growing area of interest and has shown promise. We believe that this should help provide the push for better online interactions among clinicians within a care organization to share knowledge and improve patient care.

Integrated Systems: As the healthcare industry moves from a fee-for-service model to a pay-for-performance model, healthcare organizations are being held responsible for the overall health of the population. This has resulted in an increasing trend of healthcare organizations merging into large systems that serve an entire local area and provide both care and insurance. As such systems merge and integrate, there is increased scope for data sharing, access, and analysis. We believe that, akin to the merging of brick-and-mortar healthcare systems, in the future we will see more integrated online healthcare systems that cater to different functions: care, insurance, social networking, etc. The digital world has seen large organizations such as Google trying to integrate a variety of services to capture as many users as possible while attempting to provide each user with a personalized, one-stop experience. The advancements in mobile applications, cloud technologies, and increasing data collection and analysis for healthcare will lead to such unified frameworks. Such technologies will be a win-win for the patient as well as the organization. While organizations can start reducing costs by better understanding the process and the patient's preferences, patients will benefit from an integrated view of their health and will be able to search and make better choices for their own care.

Conclusions

This entry examines how online technology can aid the healthcare industry in gaining information and intelligence and reducing costs associated with brick-and-mortar systems. We provide an overview of online healthcare systems. We briefly update the reader regarding the different types of system in this arena. We then provide example applications to illustrate the use of online technologies to aid in certain areas of healthcare management. These examples provide evidence of success and the potential of online healthcare systems. Finally, we point out the challenges in existing technologies and how data mining and social network analysis can aid future analytics that should be of interest to the healthcare community at large

Acknowledgments

Prof. Jaideep Srivastava, the research supervisor for this study, Consults with Allina Health, the research sponsor. Their relationship has been reviewed and managed by University of Minnesota in accordance with its conflict of interest policies.

Cross-References

▶ Community Evolution
▶ Corporate Online Social Networks and Company Identity
▶ Data Mining
▶ Ethical Issues Surrounding Data Collection in Online Social Networks
▶ Game Theory and Social Networks
▶ Inferring Social Ties
▶ Modeling Social Preferences Based on Social Interactions
▶ Online Social Network Privacy Management
▶ Patterns in Productive Online Networks: Roles, Interactions, and Communication
▶ Social Networks in Emergency Response
▶ Social Networking Sites
▶ Web Communities Versus Physical Communities

References

Anderson JG, Rainey MR, Eysenbach G (2003) The impact of CyberHealthcare on the physician–patient relationship. J Med Syst 27(1):67–84

Ball MJ, Smith C, Bakalar RS (2007) Personal health records: empowering consumers. J healthcare Inf Manag 21(1):77

Bertakis KD, Azari R (2011) Patient-centered care is associated with decreased healthcare utilization. J Am Board Fam Med 24(3):229–239

Doc2Doc (2013) BMJ Group. http://doc2doc.bmj.com/. Accessed 15 Mar 2013

Doximity (2013) Doximity, Inc. https://www.doximity.com/. Accessed 15 Mar 2013

E-Visits (2013) http://www.allinahealth.org/ahs/medical-services.nsf/page/evisits_MyChart. Accessed 10 Mar 2013

Eysenbach G (2001) What is e-health? J Med Internet Res 3:e20. [PMC free article] [PubMed]

Forkner-Dunn J (2003) Internet-based patient self-care: the next generation of healthcare delivery. J Med Internet Res 5(2) http://www.jmir.org/2003/2/e8/

Greene JA, Choudhry NK, Kilabuk E, Shrank WH (2011) Online social networking by patients with diabetes: a qualitative evaluation of communication with Facebook. J Gen Intern Med 26(3):287–292

Halamka, J (2006) The perfect storm for electronic health records. J Healthc Inf Manag 20(3):25

Hesse BW, Nelson DE, Kreps GL, Croyle RT, Arora NK, Rimer BK, Viswanath K (2005) Trust and sources of health information: the impact of the internet and its implications for healthcare providers: findings from the first health information national trends survey. Arch Intern Med 165(22):2618

Intuit Inc. (2013) Intuit Health patient portal. http://healthcare.intuit.com/portal/. Accessed 15 Mar 2013

Kummervold PE, Chronaki CE, Lausen B, Prokosch HU, Rasmussen J, Santana S, . . . , Wangberg SC (2008) eHealth trends in Europe 2005–2007: A population-based survey. J Med Internet Res 10(4) http://www.jmir.org/2003/2/e8/

Latitude (2013) Boston scientific. http://www.boston-scientific.com/lifebeat-online/live/latitude-remote-monitoring.html. Accessed 15 Mar 2013

LinkedIn (2013) About us. http://www.linkedin.com/about-us. Accessed 15 Mar 2013

Lorig K, Ritter PL, Laurent DD, Plant K, Green M, Jernigan VBB, Case S (2010) Online diabetes self-management program: a randomized study. Diabetes Care 33(6):1275–1281

Medicare (2013) Personal Health Records (PHRs). Medicare.gov. Accessed 1 Mar 2013

Medscape (2013) Medscape connect. http://www.medscape.com/connect. Accessed 15 Mar 2013

Nalari Health (2013) Benefits of online care. http://www.nalarihealth.com/online-healthcare/benefits-of-online-care/. Accessed 15 Mar 2013

OpenEMR Project (2013) http://www.open-emr.org/. Accessed 15 Mar 2013

PatientsLikeMe (2013) http://www.patientslikeme.com/. Accessed 15 Mar 2013

Roots K, Heller M, Srivastava J, Srivastava S, Schumann J (2013) Efficacy of a novel video game to help diagnose ADHD. In: 17th Annual MN-HSR conference, St. Paul, Minnesota

Tan P, Steinbach M, Kumar V (2006) Introduction to data mining. Addison-Wesley, Reading

Tang PC, Ash JS, Bates DW, Overhage J, Sands DZ (2006) Personal health records: Definitions, benefits, and strategies for overcoming barriers to adoption. J Am Med Inform Assoc 13(2):121–126

Valente TW (2010) Social networks and health: models, methods, and applications. Oxford University Press, Oxford

Virtuwell (2013) http://www.virtuwell.com/. Accessed 15 Mar 2013

Wallwiener M, Wallwiener CW, Kansy JK, Seeger H, Rajab TK (2009) Impact of electronic messaging on the patient-physician interaction. J Telemed Telecare 15(5):243–250

Zaroukian MH (2008) Patient Web portals, clinical messaging and e-visits. HIMSS. HIMSS Enterprise Integration Task Force, 12 June 2008. http://www.himss.org/files/HIMSSorg/content/files/Patient_portals Zaroukian.pdf. Accessed 1 Mar 2013

Zipnosis Inc (2013) http://www.zipnosis.com/. Accessed 15 Mar 2013

Online Knowledge Networks

► Networks of Practice

Online Privacy Heuristics

► Online Social Network Privacy Management

Online Privacy Paradox and Social Networks

Verena Kreilinger and Sebastian Sevignani
Unified Theory of Information Research Group (UTI), Vienna, Austria

Synonyms

Awareness; Control; Economic surveillance; Knowledge; Privacy; Social networking sites; Targeted advertising; Trade-offs

Glossary

Social Networking Sites (SNS) Public or semipublic web-based services that allow facilitating or maintaining social relations and community building among people who share something such as a common interest or background

Paradox An unexpected, ostensibly or effectively irreconcilable contradiction

Privacy Paradox (PP) Contradiction between individuals' intention to disclose private issues and individuals' actual disclosure behavior

Utility Maximization Theory Assumes that individuals aim at maximizing utility

Privacy Calculus Model Individual's trade-off between perceived risks and benefits of information disclosure

Personalized Advertising Advertisements that are customized to individuals or groups, based on personal information such as demographic data, interests, and purchasing behavior

Definition

Today, informational privacy is most often defined either as control over flows of information or as the access to information. For Westin, "privacy is the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others" (Westin 1967, 7). Westin focuses on the control of information, which makes him a prototypical proponent of "control theories" of privacy (Tavani 2008, 142–143). On the other hand, there are "access theories" of privacy (Tavani 2008, 141–142): For Gavinson, privacy "is related to our concern over our accessibility to others: the extent to which we are known to others, the extent to which others have physical access to us, and the extent to which we are the subject of others' attention" (Gavinson 1984, 347). If we combine these two major strands of privacy approaches, one can speak about privacy as the individual's control over the access to personal information (Tavani 2008).

In this entry, we deal with paradoxes that occur in online social networking actors' management of control over the access to their personal information. At large, a paradox is an unexpected, ostensibly or effectively irreconcilable contradiction. Such contradictions may occur in or between dimensions, such as online social networking actors' knowledge, experience, attitude, and behavior.

Within the literature, several PP have been identified. For example, Barnes (2006) speaks about PP but actually means that different social groups deal differently with privacy issues. Another frequently mentioned paradox is associated with companies' interest in product personalization and referred to as the personalization PP (Awad and Krishnan 2006; Radin 2001; Lee and Cranage 2011; Xu et al. 2011). Aiming at competitive advantages and profitability, companies personalize their products and services to better address customers. Personalization tends to raise privacy concerns among potential consumers, and this could affect companies' interests negatively. On the other hand, companies' efforts to give consumers more control over their information tend not to perform among people who have privacy concerns. Although this paradox has relevance for online social networking since targeted advertising is the main business model of online SNS, it is not a paradox regarding privacy. In fact, it is a paradox relating to companies' interests because they come into contradiction with consumers needs.

Actually, the most often mentioned PP is the gap between individuals' intention to disclose private issues and individual's actual disclosure behavior (Norberg et al. 2007; Spiekermann et al. 2001; Pötzsch 2009; Holland 2010; Brandimarte et al. 2010; Oetzel and Gonja 2011; Bosau et al. 2008; Dwyer et al. 2007; Tufekci 2008; Utz and Krämer 2009; Gross and Acquisti 2005). Norberg et al. (2007) conducted a study in order to shed light upon why people actually disclose more personal information than they intend to provide. They call the contradictory relationship between "individuals' intentions to disclose personal information and their actual personal information behaviors" the "privacy paradox" (Norberg et al. 2007, 100). They found that consumers not only provided more demographic data or information about personal preferences than they intended to but even disclosed sensitive data such as personally identifying or financial information. Even people with negative perceptions about disclosing personal information actually provide personal details on request. They conclude "that, in the realm of privacy, behavioral intentions may not be an accurate predictor of actual behavior" (Norberg et al. 2007, 118). Therefore, they stress the relevance of strictly differentiating between attitudes, concern, intentions to behave, and actual behavior. In the following, we will explore this genuine PP more precisely.

Historical Background

Impelled by the development and evolution of the Internet and contemporary information and communication technologies, which allow for a whole new way of information processing, privacy has become a major concern. Huge amounts of personal user information can easily be collected, stored, analyzed, distributed, and manipulated. SNS, personalized services, targeted advertising, individualized search engines, location-based applications, cloud storage, and the digitalization of formerly merely accessible data are just some of the phenomena that bring along new privacy issues. It might be no coincidence that, driven by consumer research, the PP was examined in the context of e-commerce early on (Norberg et al. 2007, 100) since a lot depends on consumers' willingness to data disclosure for these businesses. Especially in the light of SNS, further research has been conducted then in order to explain the PP.

The Privacy Paradox in Context of Social Media

The PP, observable in the digital realm, is even more striking in the context of SNS (Utz and Krämer 2009; Bosau et al. 2008; Dwyer et al. 2007; Tufekci 2008; Kreilinger 2013). In one of the first studies examining students' attitude and behavior within online social networks, Acquisti and Gross (2006) found that though privacy concerns restrained some students from joining SNS, there were no significant differences in the amount of information shared between highly concerned and unconcerned users. Many studies about the usage of SNS found that though users generally were concerned about their privacy, they disclosed and shared a lot of sensitive information (in their profiles).

Possible explanations to the PP often focus on users' misconception and lack of knowledge (Norberg et al. 2007; Barnes 2006; Utz and Krämer 2009; Pötzsch 2009; Brandimarte et al. 2010; Oetzel and Gonja 2011). Especially in context of SNS, users tend to overlook the fact that their profiles are accessible not only to their online friends but often to all Internet users. People seem to be less concerned about their privacy if they have established relationships with trusted members of such a network and forget about "all the rest of the digital world" (Sheehan and Hoy 2000; Jagatic et al. 2007; Pötzsch 2009). Others even conclude that the PP is slowly becoming resolved due to increased knowledge and experience with SNS and web applications. When examining the level of information disclosure among SNS users over time, a movement towards more restrictive privacy settings can be observed (Utz and Krämer 2009; Lewis et al. 2008). A possible explanation for this change is that when SNS were still a very new phenomenon, users enthusiastically put a lot of sensitive information on their profiles. When SNS became mainstream, people learned about the first negative consequences and cases of privacy infringement within the SNS. As a consequence, users have become more aware of privacy issues. At the same time, users became more experienced and learned how to deal with new technologies and applications (Utz and Krämer 2009).

In contrast, other authors (Brandimarte et al. 2010; Oetzel and Gonja 2011; Sandoval 2011) point to the deliberate disguise and misinformation on behalf of the companies benefitting from user's willingness to provide great amounts of personal data. As a result, confusion could also be an explanatory reason: "One answer is that much of the public is simply confused by the battles over responsibility for the environmental risks regarding information privacy" (Turow and Hennessy 2007, 315).

Further explanations provided by Norberg et al. (2007, 118) are contextual factors, such as the physical setting, social factors reflecting the relationship between the individual and the person(s) or institution collecting the information, and cognitive factors. Interestingly, they found that though perceived risks influence individuals' intention to disclose, they do not influence their actual behavior. However, no such clear relation could have been found for "trust." Gross and Acquisti (2005) argue that trust within social networks may be assigned differently and have a different meaning than in the offline environment. Social relations become more leveled: Users provide the same information to all online friends – regardless of the very different nature of these relations in the offline world.

Some authors consider a risk-benefit tradeoff to cause the PP (Chellappa and Sin 2005;

Dinev and Hart 2006). This trade-off can be described by a utility function (Awad and Krishnan 2006, 18):

$$U(X) = \text{Benefit} - \text{Cost}$$

According to utility maximization theory, users strive to maximize utility (Awad and Krishnan 2006, 17). Perceived benefits for disclosing personal data in context of SNS may include social exchange, relationships, collaborations, and reputation; perceived costs may be identity theft, marketing spam, stalking, and negative reputation in other contexts (Pötzsch 2009, 230). Chellappa and Sin (2005) found that the perceived benefits gained from a personalized service are two times more influential on actual data disclosure behavior than users' concern for privacy. Brandimarte et al. (2010) outline some other mechanisms that may contribute to explain inconsistent privacy-related decisions including prospect theory, i.e., wrong estimation of probabilities (Kahneman and Tversky 1979), trust (Culnan and Armstrong 1999), hyperbolic time discounting and immediate gratification (Acquisti and Grossklags 2003; Acquisti 2004), and rational models of decision making (Posner 1978; Stigler 1980).

Other authors consider methodological discrepancies to cause or at least contribute to the PP (Pötzsch 2009; Holland 2010; Brandimarte et al. 2010). They argue that awareness, attitude towards privacy, privacy concern, etc., are conceptually vague and empirically hard to measure in contrast to actual behavior. The first group of variables is usually measured based on users' self-perception, whereas the second can be measured by exploring factually disclosed personal information by methods of data mining, experimental settings, etc.

Pötzsch (2009, 230–231) highlights that if people are asked in general about privacy, many of them are to some extent privacy aware. She argues: "However in real situations the concrete value of privacy (costs) is hard to estimate and is no longer salient to people. The quantity of possible price premiums or the 'universe of new friends' (benefits) is primarily advertised; it is just a few clicks and disclosure of a few personal data items away" (Pötzsch 2009, 230–231). In such situations, privacy concerns take a back seat and costs are perceived low, since there is a lack of stimuli at the moment of attention (Pötzsch 2009, 231).

Brandimarte et al. (2010) emphasize another possible methodological as well as theoretical cause for the PP: the vagueness of the concept of privacy itself. They argue that it is important to distinguish between control to publish and control over access and usage of personal data: "Since we have control over the publication of our private information, we give less importance to control (or lack thereof) over the accessibility and use of that information by others" (Brandimarte et al. 2010). Thus, they infer that "control over the publication of their private information decreases individuals' privacy concerns and increases their willingness to publish sensitive information, even when the probability that strangers will access and use that information stays the same or, in fact, increases" (Brandimarte et al. 2010).

How Can the Privacy Paradox Be Handled?

The answer to the question how to deal with the PP consequently depends on how the paradox was defined and explained previously. One possible account to the privacy paradox is to accept it as it is by assuming that people usually behave in a contradictory manner. Norberg et al. (2007, 105) point to several variables that influence the relationship between intention and behavior, such as the degree of correspondence between the measure of intention and the measure of behavior, the temporal stability of intention, and the degree to which the behavior was planned. Independent from intention, behavior can be influenced by factors such as routinization of behavior, effects of heuristic processing, and information selectivity in the decision-making process. Others call for further investigations of the PP in order to find more definite explanations of the phenomenon and to be able to resolve it ultimately. Another way is to circumvent the paradox. That may be to focus solely on one part of the paradox and to neglect the privacy concern, for instance. A review of literature suggests different approaches to practically resolve the paradox and allow users for noncontradictory behavior: Norberg, Horne, and Horne appeal to consumers' responsibility and statethat "unless consumers make the effort to truly understand what they are granting permission to, and to whom they are giving their personal information, their sense of personal privacy will continue to deteriorate" (2007, 120). Most frequently, it is argued that an informed user who is free to choose will resolve the PP ultimately. Therefore, many authors stress awareness and knowledge about privacy issues on SNS as the precondition to free and conscious behavior. They suggest educational programs and initiatives to raise awareness and knowledge (Norberg et al. 2007; Barnes 2006; Utz and Krämer 2009; Pötzsch 2009). For example, Barnes (2006) recommends that increasing awareness of privacy issues, including knowledge about the blurring line between the private and the public in the age of the Internet, is key to solve the problem. At the same time, this would be a precondition of any social, technical, and legal initiatives. Pötzsch (2009) outlines that an asymmetry in initiatives to raise awareness exists. Companies strongly advertise the benefits that consumers obtain by providing personal data, but at the same time, possible risks and ways to protect privacy are rarely promoted.

Some authors use market theory and the calculus model of the user to resolve the PP (Holland 2010; Chellappa and Sin 2005; Dinev and Hart 2006). Within this framework, one may argue that initiatives either to clearly maximize the benefits of personal information disclosure for the user or to clearly minimize the costs of it will resolve paradoxical privacy-related behavior. Others, who have identified market distortions, such as incomplete or asymmetric information, bounded rationality, and psychological deviations from rationality, little bargaining power, and badly enforced property rights in personal information, propose steps to eliminate those distortions. Spiekermann et al. (2001) suggest a technical solution in order to establish perfect pseudonymity. Then information can be disclosed by the users, but this information cannot be tracked to the real user and would therefore no longer be any privacy issue. Others highlight the importance of rethinking the concept of privacy in order to resolve the paradox (Holland 2010; Brandimarte et al. 2010; Tufekci 2008). Holland argues that existing economical, educational, technical, and regulatory suggestions and practical efforts fail to address the challenge "to facilitate privacy regulation while preserving the benefits of new social spaces that require strategic data disclosure in order to work" (2010, 926). He suggests a balanced approach between overprotection and under-regulation that gives space to emerging social norms of data disclosure and protection in the realm of social media instead of fixing the now outmoded notice-and-choice system of privacy protection. In a similar context, Tufekci wishes for better "understanding of the process of privacy optimization sought by students and a dialogue about how we, as a society, wish to draw the boundaries between public and private, disclosure and withdrawal, and past choices and future possibilities" (2008, 35). Brandimarte et al. (2012) alert us about problematic implications of the dominant control theory of privacy in general and about measures that solely increase users' control over publication in particular. They show that more subjective control can paradoxically result in less privacy "in the sense of higher objective risks associated with the disclosure of personal information" (2012, 6). Consequently, these authors prefer a rethought concept of privacy that is not restricted to control over personal information and remind us that the further usage of once disclosed information is a crucial, yet underrated, aspect of privacy protection.

Future Directions: The Privacy Paradox and Society

Although the PP can be observed in sociocultural as well as political contexts, it is the economic sphere, where the paradox is notably flagrant:

"Commercial interests seek to maximize and then leverage the value of consumer information, while, at the same time, consumers voice concerns that their rights and ability to control their personal information in the marketplace are being violated. However, despite the complaints, it appears that consumers freely provide personal data" (Norberg et al. 2007, 100). Norberg et al. (2007, 100–101) underline that consumers are constantly faced with the not-so obvious trade-off between personalized products and services, based on detailed customer profiles on the one hand and the privacy encroachment that such disclosure causes on the other.

In the literature, privacy is largely seen as an individualistic issue (see, e.g., the privacy critique in Lyon (1994), Etzioni (1999), Cohen (2012), and Sevignani (2012)). The situation is, as it is understandable from the review of existing literature, similar with the related paradox: Most of the proposed explanations and approaches to the problem address the individual user and its subjective control over personal data; user awareness and knowledge rising is an important example in this context. However, in order to better understand paradoxes that occur on an individual level, it might be meaningful to consider them in a wider societal context.

Critical dialectical theory reminds us that "the conception of the contradictory nature of societal reality does not, however, sabotage knowledge of it and expose it to the merely fortuitous. Such knowledge is guaranteed by the possibility of grasping the contradiction as necessary and thus extending rationality to it" (Adorno 1976, 109). To situate the PP in a society analysis and therefore extending rationality to it has, on the one hand, the advantage that the PP appears and becomes understandable in relation to further and superordinated paradoxes and contradictions. For instance, (Turow and Hennessy 2007, 300) highlight users' ambiguous perception of the roles of institutional actors: "A key result is that a substantial percentage of internet users believes that major corporate or government institutions will both help them to protect information privacy and take that privacy away by disclosing information to other parties without permission" (Turow and Hennessy 2007, 300). On the other hand, such extension of view might provide evidence that the PP is specific to capitalist societies and their decision to organize crucial spheres of life according to a logic that seeks to maximize profits. Under this regime, people have to take disadvantages into account in order to achieve the advantages of SNS. This problem is exacerbated particularly in a situation of SNS monopolies lacking alternative, privacy-aware, and nonprofit-oriented SNS (Chen and Michael 2012). For example, users enjoy free access to SNS in order to communicate and cooperate. In doing so, they have to accept a loss of privacy since economic surveillance is an inherent part of the business model of commercial, advertising-funded SNS (Fuchs 2011a, b; Sevignani 2013a).

Further research should therefore assess the PP in its economic dimensions and relate it to critical societal analysis in order to properly understand the preconditions and heteronomies within which the PP occurs. Potential solutions of the PP should consequently go beyond individualistic, user-centric approaches and aim at altering the PP's societal preconditions. For instance, potential solutions would include political regulation (legislation and enforcement) that limits commercial use of personal data and, even more important, the support of alternative privacy-aware and noncommercial SNS (Sevignani 2013b).

Acknowledgments

The research presented in this entry was conducted in the project "Social Networking Sites in the Surveillance Society" (http://www.sns3.uti.at), funded by the Austrian Science Fund (FWF) (project number P 22445-G17; project coordinator Prof. Christian Fuchs).

Cross-References

▶ Anonymization and De-anonymization of Social Network Data
▶ Data Mining

▶ Ethical Issues Surrounding Data Collection in Online Social Networks
▶ Ethics of Social Networks and Mining
▶ Legal Implications of Social Networks
▶ Online Social Network Privacy Management
▶ Privacy and Disclosure in a Social Networking Community
▶ Privacy in Social Networks, Current and Future Research Trends on

References

Acquisti A (2004) Privacy in electronic commerce and the economics of immediate gratification. In: Proceedings of the 5th ACM conference on electronic commerce, New York

Acquisti A, Gross R (2006) Imagined communities: awareness, information sharing, and privacy on the Facebook. In: Golle P, Danezis G (eds) Proceedings of 6th workshop on privacy enhancing technologies, Robinson College, Cambridge, pp 36–58

Acquisti A, Grossklags J (2003) Losses, gains, and hyperbolic discounting: an experimental approach to information security attitudes and behaviors. In: Proceedings of the 2nd annual workshop on economics and information security, College Park

Adorno TW (1976) On the logic of the social sciences. In: Adorno TW, Albert H, Dahrendorf R, Habermas J, Pilot H, Popper KR (eds) The positivist dispute in German sociology. Heinemann, London, pp 105–122

Awad NF, Krishnan MS (2006) The personalization privacy paradox: an empirical evaluation of information transparency and the willingness to be profiled online for personalization. MIS Q 30(1):13–28

Barnes SB (2006) A privacy paradox: social networking in the United States. First Monday 11(9):11–15

Bosau C, Fischer O, Koll M (2008) StudiVZ: determinants of social networking and dissemination of information among students. Paper presented at the international congress of psychology, Berlin

Brandimarte L, Acquisti A, Loewenstein G (2012) Misplaced confidences: privacy and the control paradox. Soc Psychol Personal Sci 4(3):340–347

Chellappa RK, Sin RG (2005) Personalization versus privacy: an empirical examination of the online consumer's dilemma. Inf Technol Manag 6(2):181–202

Chen X, Michael K (2012) Privacy issues and solutions in social networking sites. IEEE Technol Soc Mag 31(4):43–53

Cohen JE (2012) Configuring the networked self: law, code, and the play of everyday practice. Harvard University Press, New Haven

Culnan MJ, Armstrong PK (1999) Information privacy concerns, procedural fairness, and impersonal trust: an empirical investigation. Organ Sci 10(1):104–115. doi:10.1287/orsc.10.1.104

Dinev T, Hart P (2006) An extended privacy calculus model for ecommerce transactions. Inf Syst Res 17(1):61–80. doi:10.1287/isre.1060.0080

Dwyer C, Passerini K, Hiltz SR (2007) Trust and privacy concern within social networking sites: a comparison of Facebook and MySpace. In: Proceedings of the thirteenth Americas conference on information systems, Keystone

Etzioni A (1999) The limits of privacy. Basic Books, New York

Fuchs C (2011a) An alternative view of privacy on Facebook. Information 2:140–165. doi:10.3390/info2010140

Fuchs C (2011b) New media, Web 2.0 and surveillance. Sociol Compass 5:134–147

Gavinson R (1984) Privacy and the limits of law. In: Schoeman F (eds) Philosophical dimensions of privacy: an anthology. Cambridge University Press, Cambridge/New York, pp 346–402

Gross R, Acquisti A (2005) Proceedings of the Workshop on Privacy in the Electronic Society (WPES), ACM, 71–80

Holland HB (2010) Privacy paradox 2.0. Widener Law J 19(3):893

Jagatic TN, Johnson NA, Jakobsson M, Menczer F (2007) Social phishing. Commun ACM 50(10):94–100. doi:10.1145/1290958.1290968

Kahneman D, Tversky A (1979) Prospect theory: an analysis of decision under risk. Econometrica 47(2): 263–291

Kreilinger V (2013) Social networking sites in the surveillance society. Research design & data analysis, presentation, and interpretation: part two. The internet & surveillance – Research paper series (Forthcoming report). Available from http://www.sns3.uti.at.dd29412.kasserver.com/?page_id=24

Lee CH, Cranage DA (2011) Personalisation–privacy paradox: the effects of personalisation and privacy assurance on customer responses to travel Web sites. Tour Manag 32(5):987–994. doi:10.1016/j.tourman.2010.08.011

Lewis K, Kaufman J, Christakis N (2008) The taste for privacy: an analysis of college student privacy settings in an online social network. J Comput Mediat Commun 14(1):79–100. doi:10.1111/j.1083-6101.2008.01432.x

Lyon D (1994) The electronic eye: the rise of surveillance society. University of Minnesota Press, Minneapolis

Norberg PA, Horne DR, Horne DA (2007) The privacy paradox: personal information disclosure intentions versus behaviors. J Consum Aff 41(1):100–126. doi:10.1111/j.1745-6606.2006.00070.x

Oetzel MC, Gonja T (2011) The online privacy paradox: a social representations perspective. In: Proceedings of the 2011 annual conference extended abstracts on human factors in computing systems, CHI EA'11, Vancouver. ACM, New York, pp 2107–2112. doi:10.1145/1979742.1979887

Posner, RA (1978) An economic theory of privacy. Regulation, May/June, 19–26

Pötzsch S (2009) Privacy awareness: a means to solve the privacy paradox? In: Matyáš V, Fischer-Hübner S, Cvrèek D, Švenda P (eds) The future of identity in the information society. Springer, Berlin/Heidelberg, pp 226–236

Radin TJ (2001) The privacy paradox: E-commerce and personal information on the Internet. Bus Prof Ethics J 20(3/4):145–170

Sandoval M (2011) Consumer surveillance on web 2.0. In: Fuchs C, Boersma K, Albrechtslund A, Sandoval M (eds) The Internet & surveillance. Routledge, New York, pp 147–169

Sevignani S (2012) The problem of privacy in capitalism and the alternative social networking site Diaspora. TripleC J Sustain Inf Soc 10(2):600–617

Sevignani S (2013a) Facebook vs. Diaspora: a critical study. In: Lovink G, Rasch M (eds) Unlike us reader: social media monopolies and their alternatives. Institute of Network Cultures, Amsterdam, pp 323–337

Sevignani S (2013b) Privacy in the Internet: commodity vs. common good. Sci Public Policy (Forthcoming article)

Sheehan KB, Hoy MG (2000) Dimensions of privacy concern among online consumers. J Public Policy Mark 19(1):62–73

Spiekermann S, Grossklags J, Berendt B (2001) E-privacy in 2nd generation E-commerce: privacy preferences versus actual behavior. In: Proceedings of the 3rd ACM conference on electronic commerce, Tampa. ACM, New York, pp 38–47

Stigler GJ (1980) An introduction to privacy in economics and politics. J Leg Stud 9(4):623–644

Tavani HT (2008) Informational privacy: concepts, theories, and controversies. In: Himma KE, Tavani HT (eds) The handbook of information and computer ethics. Wiley, Hoboken, pp 131–164

Tufekci Z (2008) Can you see me now? Audience and disclosure regulation in online social network sites. Bull Sci Technol Soc 28(1):20–36. doi:10.1177/0270467607311484

Turow J, Hennessy M (2007) Internet Privacy and Institutional Trust: Insights from a National Survey. New Media Society 9(2):300–318.

Utz S, Krämer NC (2009) The privacy paradox on social network sites revisited: the role of individual characteristics and group norms. Cyberpsychol J Psychosoc Res Cyberspace 3(2):1–11

Westin A (1967) Privacy and freedom. Atheneum, New York

Xu H, Luo X (Robert), Carroll JM, Rosson MB (2011) The personalization privacy paradox: an exploratory study of decision making process for location-aware marketing. Decis Support Syst 51(1):42–52. doi:10.1016/j.dss.2010.11.017.egime

Online Privacy Risk Management

▶ Online Social Network Privacy Management

Online Retail

▶ E-Commerce and Internet Business

Online Social Media

▶ User Behavior in Online Social Networks, Influencing Factors

Online Social Network Phishing Attack

Theodoros Tzouramanis and Loukas Karampelas
Department of Information and Communication Systems Engineering, University of the Aegean, Karlovassi, Samos, Greece

Synonyms

Scamming

Glossary

Social Engineering Attacks Attacks that rely on the psychological manipulation of the chosen victim

Spam Huge bulk of copied messages to an extremely large number of recipients

DNS Hijacking/DNS Redirection The practice of redirecting the resolution of Domain Name Systems (DNS) names to other DNS servers

Definition

Phishing is a kind of attack whereby an attacker endeavors to steal sensitive information. The technique is to direct the victim to provide the information voluntarily under the pretence that this information is requested by a legitimate source. Phishing attacks are usually motivated by the intention of stealing personal credit card or personal bank account details via e-mails but are also used for many other purposes always depending on the situation and the needs of the phisher. Social network phishing is the technique of stealing users' log-in and other sensitive personal information in online social network sites.

Introduction

Millions of personal computer users maintain a profile in at least one online social network. It would appear that the majority of users tend to use online social networks more than they use e-mails to communicate with their friends. This trend has caught the attention of phishers, who know that as the number of online users increases, the number of potential victims for their attacks increases as well. While this popularity is the basic attraction of social networks for phishers, social networks are also a choice area to focus on: on the one hand, a scammed social network profile allows more effective phishing targeting of both the owner of the profile and her/his friends with social engineering attacks, and on the other hand, phishers can simply spam these victims for advertisement and such apparent purposes, in order to install malicious code or ensnare them into a website. The most worrying issue with regard to phishing in social networks is that the majority of social network members who have a profile are not aware of how real the threat is and of how likely they are to become victims. The degree to which phishing has become widespread is clear from the fact that it is possible to become a phisher without advanced knowledge in computing, not counting the plethora of free-access phishing tutorials available on the Internet, or even on the popular video-sharing website YouTube, to amateurs more usually than not without good background knowledge in computing. While such activities might be perceived as harmless fun by amateur phishers, professional phishers, on the other hand, are extremely dangerous. Phishing comes under identity theft crimes and the penalty is heavy (CriminalDefenceLawyer.com 2012).

Historical Background

Phishers have been active for many years. Before the technology boom, phishing techniques were based on traditional forms of postal correspondence, aiming at the fraudulent financial extortion of victims in the guise of investment opportunities. The telephone subsequently replaced traditional correspondence, and it has now given way to electronic mail (e-mail).

The first official use of the term "phishing" was recorded back in 1996. The term is a variant of "fishing" and probably influenced by the word "phreaking" (Liberman 2004).

The first phishing incident of size targeted the America Online (AOL) Network Systems in 1990. In order to open an account and have access to AOL resources, one would have to provide credit card details. Hackers created false AOL accounts by using algorithms that generated fake credit card numbers. AOL afterwards developed tools to prevent this kind of attack, by testing each given credit card number and by accepting it only if it corresponded to a real bank credit card account. As hackers were stopped from having access to AOL resources, they devised another way of getting access: by phishing legitimate accounts from the users under the pretence that they were employed by AOL.

Since then, phishing attacks have focused on online banking, online auctions, and online social network websites.

Scientific Fundamentals

Techniques

The simplest way of carrying out a phishing attack on social networks is by using the profile cloning method: the profile of a victim is cloned and used to send a friend request to her/his friends. It is based on the expectation that a number of friends will accept these requests as they will assume that the requester created a second profile. Those friends that accept the request are the target victims. A combination of friendship responses and social engineering attacks gives the phisher access to information which she/he should not be allowed to access. In Bilge et al. (2009) an automated system was devised that uses a cross-site-cloning attack where victims' contacts are stolen and used to build a realistic fake profile in a social network in which the victim has not yet registered. In Huber et al. (2009) an automated system (a bot) was devised that carried out classic social engineering attacks in social networks: the idea is that the bot collects data from profiles and, then using these social data, designs and sends the social network users personalized intelligent messages aimed at manipulating them.

Scamming an original profile is another way of carrying out a phishing attack. This method improves the success of a social engineering attack by virtue of the fact that phishers are able to use the original identity of the victim. Once a profile has been scammed, the phisher can decide on the strategy that she/he will follow next. In most of the cases, mass spam or phishing messages will follow this attack to the victims and the list of their friends.

Table 1 summarizes all the different phishing techniques that are surveyed in this entry.

Phishing Steps

The first step in phishing design is the lure. The lure is the most important phase of the phishing process because it determines to the largest extent whether a user will fall victim of the attack. In electronic phishing, the lure could be contained in an e-mail, in a message, or in a malware program. In every case, the form and the content of the lure must be convincing. Consequently, the lure consists of a convincing story and a URL hyperlink that directs the victim to a website that is under the control of the phisher. If the lure is of a sufficiently high standard, it is possible for even a user who is aware of the technique and in a position to suspect that she/he is dealing with a lure, to become a victim, just as with the average user, as shown by a relevant survey (Dhamija et al. 2006), in which 91 % of the participants were convinced that a spoofed Web banking site was legitimate. Lures are sent directly to an individual user, so if the phisher wants to target users in mass, the success of her/his attack will depend also on the number of user profiles scammed.

The hook is the last step in the design of a successful phishing attack. In the social networks context, a phishing hook is a website that is a copy of an authentic website. The hook is used for stealing the log-in or other sensitive personal information of the victim. When a victim enters her/his personal details in a hook and attempts, for example, to log in, there are two conceivable scenarios: either the user becomes aware that she/he is been scammed or she/he gets caught. In the latter case, the time factor is to the advantage of the phisher; in the former case, the time factor stands against the phisher, and the intended victim may change her/his log-in details before the scammer has time to get hold of them. A relevant survey on phishing steps and techniques, in Jakobsson and Myers (2006), is highly recommended. Finally, it should be mentioned that it is common for users to stick to the same password across a range of websites that they use making it thus easier for phishers to test the passwords already in their possession to extend the Web of their phishing and attack the users' profiles in sites of financial interest such as e-auction and e-banking portals.

Sophisticated Techniques

There are techniques to increase the success of a lure, for example, by creating a convincing e-mail with company logos, achieving the degree of formality that will deceive the victim. In addition to this, phishing attacks also use spoofing. This technique consists of sending an e-mail with a fake id meant to look authentic, which can be made possible since the Simple Mail Transfer Protocol (SMTP) does not provide authentication. The lack of authentication in

Online Social Network Phishing Attack, Table 1 The various phishing attack techniques that are surveyed in this entry

Phishing attack technique	Description
Social engineering attack	Attack that relies on the psychological manipulation of the chosen victim
Classic phishing attack	Sending an e-mail that seems legitimate and contains a hyperlink which directs the victim to the attacker's server
Profile cloning	Creating a fake identity of a victim and using it for cheating her/his friends in the same network
Cross-site cloning	Creating a fake identity of a victim and using it for cheating her/his friends in other social networks where the victim does not hold an account
Spam	Sending a huge bulk of copied messages to an extremely large number of recipients
Using an already scammed profile	Stealing the original identity of a victim and using it for cheating her/his friends
Spoofing	Sending an e-mail that appears to be coming from an original legitimate source
DNS hijacking/DNS redirection	Redirecting the resolution of Domain Name Systems (DNS) names to other DNS servers
Phraming	Hacking into an original DNS server and redirecting all the traffic to the attacker's server

e-mail technology makes it possible for anyone to pretend to another's identity, deceiving even the experts.

Also, as mentioned above, a hyperlink is contained in the lure; therefore, this link needs to be appropriately selected in order to look like a legitimate one. This link is the one with the hook behind it and refers directly to a server that is under the control of the phisher. If the link differs significantly from the legitimate one, the targeted user may become suspicious. Therefore the competent phisher adds to the hyperlink with the hook, the name of the company that provides the original URL and some keywords to draw the attention of their victim, such as "log-in" or "update now."

A more efficient approach still in the context of this technique is a DNS attack whereby phishers can use exactly the same hyperlink as the original. DNS attacks are also termed hijacking attacks (Internet Corporation for Assigned Names and Numbers ICANN; Michelakis et al. 2004). They consist of two types of attacks. The first type is when the attacker hacks a DNS server and gains access to DNS records and subsequently modifies these, so that requests for the genuine webpage will be redirected to the page selected by the hacker, meaning that all the traffic of the network that is served from this server is under the control of the attacker. An online social networking phisher directs to such a fake website all the traffic that concerns the social network that she/he is targeting. The range of the network traffic that a hacker is able to control depends on the DNS server that is hacked, since the higher it is in the hierarchy, the bigger the network it is serving. This kind of attack has been termed "pharming" from the interpretation that if an attack against a DNS server succeeds, the phisher will literally harvest or "farm" the personal details of all the users across the network, with the hacked DNS server under her/his control. The danger which such an attack represents, when successful, is to be measured by the size of the network: the larger the network, the more disastrous the potential consequences.

The second type of hijacking attack is when the perpetrator simply registers a domain name similar enough to a legitimate one that users are likely to type in, either by mistaking the actual name or through a typo. One big issue that phishers are called to confront is their trace. Their trace depends on two factors: one regards the lure and the other the hook. If the e-mail that contains the lure is sent from a phisher's personal computer, it is possible that she/he will be traced, let alone if this e-mail has been sent to thousands of individuals. To cover up her/his traces, the phisher usually uses peer-to-peer networks and, more specifically, botnets that are collections of compromised computers by Trojan horses which the phisher controls remotely. Botnets allow the easy forwarding of e-mails by the phisher, since a large number of computer systems are under her/his control and the spread/promotion of "lure e-mails" is extremely fast.

As mentioned above the hook URL forwards the victim directly to a server that is under the control of the phisher. Since a phisher cannot use her/his computer address for fear of being traced, she/he uses hacked computers as servers. This can be proved a disadvantage: if the hacked computer is located by the authorities, the phisher will inevitably be traced during the next access session for the purpose of collecting the results of phishing.

Common Techniques

The majority of phishers are not in possession of either the skills or the knowledge to use the techniques referred above and therefore use easier tactics albeit not so effective. The most common of these tactics is an attack related to DNS and is a "shortened URL method" (Chhabra et al. 2011). It will be remembered that the hook, meaning the phisher's webpage, is loaded in a server. Some servers allow anyone to load a webpage free of charge and of control. While this is a welcome convenience, two problems, however, confront phishers: the first regards the danger of being traced back, because these servers are known targets for these kinds of attacks. The other problem is that these servers provide phishers with a huge URL for their hook, which is likely to raise the potential victims' suspicion. With the shortened URL method, there is a server that acts as a proxy to the real server and provides the phisher with a URL that is much shorter than the original one. There are servers that provide this facility for free, and they are easy to use for the average user. Such servers make it possible for phishers to hide their tracks more efficiently because the server containing the webpage is not used directly but indirectly. In addition, shortened URLs can more easily beat filter mechanisms such as anti-spam and anti-phishing. Another positive for the phishers is that if their shortened hyperlink is marked by anti-phishing filters as a phishing webpage, the phishers can create new shortened URLs that will be clear.

Consequences of Social Network Phishing

Social network phishing stands as a major threat to individuals and to society. The consequences of social network phishing activity are considerable, with more than a billion users today using online social networking sites. These are spread geographically across countries and include also users that extend their influence to entire communities, for example, politicians, journalists, and generally people who are prominent in the media.

In Jagatic et al. (2007), sensitive information was phished from a social network when an e-mail was sent to the participants under the pretence that it originated from a legitimate source. In this study 72 % of the participants were phished because the lure e-mail was supposed to have been sent from a friend of the victim, thus making it more attractive and convincing than a lure from an unknown source. It is a statistical fact (Gross and Acquisti 2005) that the majority of the users' profiles in online social networks stores the real personal information of the corresponding users, such as names, e-mails, phone numbers, family status, dates of birth or birthdays, work and professional information, education, religion, current addresses, travelling, habits, preferences, and much more information that can be of use in a social engineering attack. By default settings this sensitive information is visible to friends, possibly across all social networks. Therefore a scammed profile with many contacts is a significant social information database. This information can be downloaded by either a manual or an automated process (Nazir et al. 2008), for example, when the victim has installed an application that is under the control of the attacker. Again, a manual or an automated process with a specialized software can then be used to devise specific lures targeting all these new contacts. E-mail addresses collected from a social network's database, for example, can be used for sending the corresponding users a mass e-mail, with the lure content. As for the subject of the content of the lure e-mail, the phisher may prefer an issue that a large number of users connected by friendship share in common (of course, the success rate of this attack will depend also on the use of the abovementioned phishing tactics). The phisher needs then to repeat the same or a similar strategy to the contacts of these contacts, and this process can go on infinitely. When a phisher collects an important number of scammed profiles of unsuspicious users, she/he can implement phishing tactics to phish the information she/he wants, for example, credit cards or bank accounts, systematically enhanced with social information. A victim that has already been phished in a social network domain will be easily also phished in other domains as well.

Anti-phishing Measures

In phishing attacks phishers usually do not use malicious code for their purposes and so the existing security mechanisms cannot confront this threat with a good success rate. Many anti-phishing applications have recently been developed such as (Cranor et al. 2006; Zhang et al. 2007) and some anti-spam applications such as TechTarget (2013). Most of these usually operate in collaboration with existing databases of blacklisted URLs, such as the Phishtank (2013), which is based on online URL scan. Some other of these databases can filter even shortened URLs (Evans 2011). In addition to these systems which are able to limit to a certain extent the phishing phenomenon, social network sites use also their own protocol to ensure security and privacy. For example, Facebook sets a limit on the number of friend requests that a user may issue in a specific time period, meaning that every user can add only a specific number of new friends to her/his friend list in this specific time period, setting thus limitations to the automatic creation of malicious user profiles. Users can also report other users as malicious for the administrators to check the incident.

Beyond the scope of the technical measures that the security officers of an online social network can take to protect its users against phishing (McGeehan 2009), individuals carry themselves the major share of responsibility. Companies nowadays go to considerable expense to educate their employees in the matter of security (Gordon et al. 2006). Some of these companies, in order to keep their employees in a state of awareness and preparedness, train employees by sending them fake e-mails with a phishing content, an example of which is the Phishme application (PhishMe 2013) that creates fake e-mails and sends them to a given list of recipients, keeping records on the recipients who fall in the phisher's trap. Phishing awareness has also been conducted in a classroom setting, producing good results (Robila and Ragucci 2006). Much anti-phishing education material (Anandpara et al. 2007; Kumaraguru et al. 2007) and online phishing tests (SonicWALL 2013) are available to learn distinguishing between phishing and legitimate e-mails and websites. The common advice that prevails is to be aware of suspicious e-mails and hyperlinks and to always access a website by typing into the Web browser its real address. Measures will both reduce the threat to individual users and the effect of the phenomenon as a whole. Concluding, social networks should seriously consider limiting the personal data of a user profile that are visible to the public and to the friends list, especially in the default privacy settings of the profile.

Cross-References

▶ Social Engineering/Phishing

References

Anandpara V, Dingman A, Jakobsson M, Liu D, Roinestad H (2007) Phishing IQ tests measure fear, not ability. In: Proceedings of the 11th international conference on financial cryptography and 1st international conference on usable security, Scarborough, Trinidad, pp 1–6

Bilge L, Strufe T, Balzarotti D, Kirda E (2009) All your contacts are belong to us: automated identity theft attacks on social networks. In: 18th international conference on World Wide Web (WWW), Madrid, pp 551–560

Chhabra S, Aggarwal A, Benevenuto F, Kumaraguru P (2011) Phi.sh/$social: the phishing landscape through short urls. In: Annual collaboration, electronic messaging, anti-abuse and spam conference (CEAS), Perth, pp 92–101

Cranor L, Egelman S, Hong J, Zhang Y (2006) Phinding phish: an evaluation of anti-phishing toolbars. Technical report, Carnegie Mellon University

CriminalDefenceLawyer.com: Phishing: sentencing and penalties. Address to download: http://www.criminaldefenselawyer.com/crime-penalties/federal/phishing.htm. Retrieved on 1 Oct 2012

Dhamija R, Tygar JD, Hearst M (2006) Why phishing works. In: Proceedings of the SIGCHI conference on human factors in computing systems, Montreal, pp 581–550

Evans J (2013) Manage your Facebook privacy and reputation profile, 2011. Address to download: http://www.julianevansblog.com/2011/04/manage-your-facebook-privacy-and-reputation-profile.html. Retrieved on 18 Apr 2013

Gordon LA, Loeb MP, Lucyshyn W, Richardson R (2006) CSI/FBI computer crime and security survey. Technical report, Computer Security Institute

Gross R, Acquisti A (2005) Information revelation and privacy in online social networks. In: Proceedings of the workshop on privacy in the electronic society, Alexandria, pp 71–80

Huber M, Kowalski S, Nohlberg M, Tjoa S (2009) Towards automating social engineering using social networking sites. In: IEEE international conference on computational science and engineering, Vancouver, vol 3. pp 117–124

Internet Corporation for Assigned Names and Numbers (ICANN) Security and Stability Advisory Committee (SSAC) (2005) Domain name hijacking: incidents, threats, risks, and remedial action. Address to download: http://archive.icann.org/en/announcements/hijacking-report-12jul05.pdf. Retrieved on 3 Nov 2013

Jagatic TN, Johnson NA, Jakobsson M, Menczer F (2007) Social phishing. Commun ACM 50(10):94–100

Jakobsson M, Myers S (eds) (2006) Phishing and countermeasures: understanding the increasing problem of electronic identity theft. Wiley, Hoboken

Kumaraguru P, Sheng S, Acquisti A, Cranor LF, Hong J (2007) Teaching Johnny not to fall for phish. tech. Cranegie Mellon University, Pittsburgh

Kumaraguru P, Sheng S, Acquisti A, Cranor LF, Hong J (2010) Teaching Johnny not to fall for phish. ACM Trans Internet Technol 10(2):31

Liberman M (2013) Phishing, in the language log, 2004. Address to download: http://itre.cis.upenn.edu/~myl/languagelog/archives/001477.html. Retrieved on 18 Apr 2013

McGeehan R (2009) Protect yourself against phishing. Technical report, Facebook. Address to download: https://www.facebook.com/blog.php?post=81474932130. Retrieved on 18 Apr 2013

Michelakis E, Androutsopoulos I, Paliouras G, Sakkis G, Stamatopoulos P (2004) Filtron: a learning-based anti-spam filter. In: Proceedings of the 1st conference on email and anti-spam, Mountain View

Nazir A, Raza S, Chuah C-N (2008) Unveiling facebook: a measurement study of social network based applications. In: Proceedings of the 8th ACM SIGCOMM conference on internet measurement. ACM, Vouliagmeni, Greece, pp 43–56

PhishMe Inc (2013) Phishme. Address to download: http://www.phishme.com/. Retrieved on 18 Apr 2013

Phishtank, http://www.phishtank.com. Retrieved on 18 Apr 2013

Robila SA, Ragucci JW (2006) Don't be a phish: steps in user education. In: Proceedings of the 11th annual SIGCSE conference on innovation and technology in computer science education, Bologna, Italy. ACM, New York, pp 237–241

SonicWALL, Inc (2013) SonicWALL phishing IQ test. Address to download: http://www.sonicwall.com/furl/phishing. Retrieved on 18 Apr 2013

TechTarget: Hijacking – Definition. Address to download: http://searchsecurity.techtarget.com/definition/hijacking. Retrieved on 18 Apr 2013

Zhang Y, Egelman S, Cranor LF, Hong J (2007) Phinding phish: evaluating anti-phishing tools. In: Proceedings of the 14th annual network & distributed system security symposium (NDSS), San Diego

Online Social Network Privacy Management

Catherine Dwyer
Seidenberg School of Computer Science and
Information Systems, Pace University,
New York, NY, USA

Synonyms

Managing online information flows; Online privacy heuristics; Online privacy risk management

Glossary

Self-Censoring To consciously limit the content of comments and posts to avoid controversial topics

Obfuscation To preserve privacy by producing misleading, false, or ambiguous data with the intention of confusing an adversary or simply adding to the time or cost of separating bad data from good

Impression Management Maintaining a desired social presence, which includes adjustments and behavioral changes in response to feedback within a social setting

Socio-technical A concept that examines activity with the understanding that it is influenced by both social and technical factors

Definition

Online social network privacy management describes the strategies participants in online social networks follow, in order to benefit from involvement in an online community while controlling the risks of unintended information leakage. Online social network privacy management is a socio-technical process that combines the use of technology-based privacy tools along with behavioral mechanisms to control information flow.

Boyd and Ellison define online social networks as "web-based services that allow individuals to (1) construct a public or semi-public profile within a bounded system; (2) articulate a list of other users with whom they share a connection; and (3) view and traverse their list of connections and those made by others within the system" (2007). Online privacy incidents often result from an individual revealing information to unintended audiences. The ability to traverse social networks also enables information to leak across contexts. When others in your network disclose information about you, this is an extremely difficult problem to manage, because the privacy problem results from the actions of others, beyond the control of individual privacy settings.

Given the leaky nature of online social networks, individuals must undertake sophisticated strategies to control information flow across contexts (Chen and Michael 2012). One example is the construction of multiple profiles (identities) in an online social network, in order to present separate profiles to different audiences (Stutzman and Hartzog 2009). This strategy can be constrained by the specified Terms of Service (TOS) of online social networks that require participants to use their "real name" (e.g., Facebook) and prohibit multiple accounts. This strategy and its challenges demonstrate the socio-technical nature of online privacy management.

The use by teenagers of online social networks has been the subject of public concern and extensive study. Boyd and Marwick (2011) report that teenagers follow varied practices to manage their online privacy. Rather than specifically depending on technology-based tools, teenagers use behavioral mechanisms, such as self-censoring or obfuscation that keep communication within the intended scope of the interaction. Teenagers may resumé-ify their profiles, pitching their self-presentation to those who have power over their future. This is an adult-approved approach, but one that is disconnected from teens' social dynamics, that prioritizes socialization over adult acceptance. Another technique is the use of false names and information in their profile. This is also encouraged by adults, without thought as to what it means to suggest lying to solve social woes.

In considering online privacy management, it is important to recognize the cognitive complexity of privacy management in everyday life (Dwyer 2008). Privacy management involves the selection of what information is shared with other parties, in order to cultivate a specific public reputation. It is a complex task that evolves during years of socialization. The evolution of privacy management within children has been noted. For example, when a child first comprehends they are naked in a public space is marked as a cognitive milestone.

Privacy management involves real-time information boundary management decisions (Petronio 2002). Studies in the area of cognitive psychology have found that privacy management decisions are constrained by cognitive limitations due to the need to make judgments quickly (Carey and Burkell 2009). We can think of people applying the same "bounded rationality" (Simon 1955) to privacy management decisions that are applied to other types of decisions.

Online privacy management is increased in complexity by several properties of online information flow. Boyd describes four characteristics of information shared in online social networks relevant to privacy management (2007): (1) information shared online is searchable; (2) information persists in these spaces beyond the life of an encounter; (3) information can be replicated or copied, making it difficult to determine validity; and (4) information shared is public to unknown and unknowable audiences.

In managing their online interactions, individuals grapple with tensions between presenting their identity and managing their privacy. Research has shown a gap between individual's perceptions of their level of privacy, compared to the actual availability of their profile data (Madejski et al. 2011). People do share information about themselves online even if they perceive privacy risks. It is more than a simple cost-benefit analysis. There is no way to calculate the "value" of a social relationship in a tangible way, so that it can be weighed against the risk of privacy problems (Dwyer and Hiltz 2008). Risk cannot be simply measured, and behavioral economists show that people are very poor judges of risk (Carey and Burkell 2009).

Participants in online social networks consider privacy management to be their own responsibility: "*I don't post anything that is private*." They do not just rely on the privacy tools provided by online social networks (Madden 2012). Instead individuals have developed blended approaches that combine the use of technology with behavioral strategies (Collins et al. 2012). Future development of privacy management tools will be more successful if they are designed from a socio-technical perspective.

Cross-References

▶ Ethical Issues Surrounding Data Collection in Online Social Networks

▶ Graphical User Interfaces for Privacy Settings

▶ Online Privacy Paradox and Social Networks

References

Boyd D (2007) Social network sites: public, private, or what? Knowledge tree. Retrieved 25 July 2007, from http://kt.flexiblelearning.net.au/tkt2007/?page_id=28

Boyd D, Ellison NB (2007) Social network sites: definition, history, and scholarship. J Comput-Mediat Commun 13(1), article 11

Boyd D, Marwick AE (2011) Social privacy in networked publics: teens attitudes, practices, and strategies. In: A decade in internet time: symposium on the dynamics of the internet and society, Oxford

Carey R, Burkell J (2009) A heuristics approach to understanding privacy-protecting behaviors in digital social environments. In: Kerr I, Steeves V, Lucock C (eds) Lessons from the identity trail. Oxford University Press, New York

Chen X, Michael K (2012) Privacy issues and solutions in social network sites. Technol Soc Mag IEEE 31(4): 43–53

Collins R, Dwyer C, Hiltz S, Shrivastav H (2012) Do I know what you can see? social networking sites and privacy management. Paper presented at the AMCIS, Seattle

Dwyer C (2008) Appropriation of privacy management within social networking sites. PhD, New Jersey Institute of Technology, Newark. Retrieved from http://library.njit.edu/

Dwyer C, Hiltz SR (2008) Designing privacy into online communities. Paper presented at the proceedings of internet research 9.0, Copenhagen

Madden M (2012) Privacy management on social media sites. Pew internet report. Retrieved from http://pewinternet.org/Reports/2012/Privacy-management-on-social-media.aspx

Madejski M, Johnson M, Bellovin S (2011) The failure of online social network privacy settings, Technical report CUCS-010-11, Columbia University

Petronio S (2002) Boundaries of privacy: dialectics of disclosure. State University of New York Press, Albany

Simon H (1955) A behavioral model of rational choice. Q J Econ 69(1):99–118

Stutzman FD, Hartzog W (2009) Boundary regulation in social media. SSRN eLibrary. doi:10.2139/ssrn.1566904

Online Social Networks

▸ Topology of Online Social Networks

Online Social Networks' Concepts

▸ Socio-Graph Representations, Concepts, Data, and Analysis

Online Social Networks: Online Social Networking Platforms, Online Social Media

▸ Privacy Preservation and Location-Based Online Social Networking

Ontologies

▸ Semantic Social Networks

Ontology

▸ Automatic Document Topic Identification Using Social Knowledge Network

Ontology Alignment

▸ Ontology Matching

Ontology Matching

Pavel Shvaiko
Informatica Trentina SpA, Trento, Italy

Synonyms

Ontology alignment; Schema matching

Glossary

OM Ontology matching
OAEI Ontology alignment evaluation initiative

Definition

The *matching* operation determines an *alignment* A' for a pair of ontologies $O1$ and $O2$. Thus, given a pair of ontologies, which can be very simple and contain one entity each, the *matching task* is that of finding an alignment between these ontologies. There are some other parameters that can extend the definition of matching, namely, (i) the use of an input alignment A, which is to be extended or completed; (ii) the matching parameters, for instance, weights or thresholds; and (iii) external resources, such as common knowledge and domain-specific thesauri, see Fig. 1.

An *alignment* is a set of correspondences between entities belonging to the matched ontologies. Alignments can be of various cardinalities: 1:1 (one-to-one), 1:m (one-to-many), n:1 (many-to-one), or n:m (many-to-many).

Given two ontologies, a *correspondence* is a 3-uple:

$$\langle e_1, e_2, r \rangle,$$

such that:

- e_1 and e_2 are entities, e.g., classes and properties, of the first and the second ontology, respectively.
- r is a relation, e.g., equivalence ($=$), more general ($\sqsupseteq$), disjointness ($\perp$), and holding between e_1 and e_2.

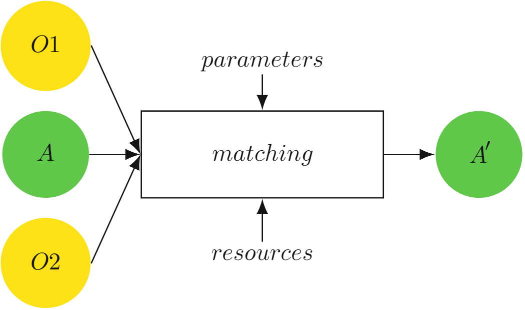

Ontology Matching, Fig. 1 The ontology matching operation

The correspondence $\langle e_1, e_2, r\rangle$ asserts that the relation r holds between the ontology entities e_1 and e_2. For instance, $\langle \mathsf{Book}, \mathsf{Monograph}, \sqsupseteq\rangle$ asserts that Book in $O1$ is more general ($\sqsupseteq$) than Monograph in $O2$. Correspondences have some associated *metadata*, such as the correspondence author name. Another frequently used metadata element is a confidence in the correspondence, which is typically in the $[0, 1]$ range. The higher the confidence, the higher the likelihood that the relation holds (Euzenat and Shvaiko 2007; Shvaiko and Euzenat 2013).

Illustrative Example

In order to illustrate the matching problem, let us use the two simple ontologies $O1$ and $O2$, which are presented in Fig. 2. Classes are shown in rectangles with rounded corners, e.g., in $O1$, Book being a subclass of Product, while relations are shown without the latter, such as price being an attribute defined on the integer domain and creator being a property. Albert Camus: La chute is a shared instance. Correspondences are shown as thick arrows that connect an entity from $O1$ with an entity from $O2$. They are annotated with the relation that is expressed by the correspondence, for example, Person in $O1$ is less general ($\sqsubseteq$) than Human in $O2$.

Let us suppose that an e-commerce company acquires another one. Technically, this acquisition requires the integration of their information sources and, hence, of the ontologies of these companies. The documents or instance data of both companies is stored according to ontologies $O1$ and $O2$, respectively. In our example these ontologies contain subsumption statements, property specifications, and instance descriptions. The first step in integrating ontologies is matching, which identifies correspondences, namely, the candidate entities to be merged or to have subsumption relationships under an integrated ontology. Once the correspondences between two ontologies have been determined, they may be used, for instance, for generating query expressions that automatically translate instances of these ontologies under an integrated ontology. For example, the attributes with labels title in $O1$ and in $O2$ are the candidates to be merged, while the class with label Monograph in $O2$ should be subsumed by the class Product in $O1$ (Shvaiko and Euzenat 2013).

The www.OntologyMatching.org contains links to a number of ontology matching projects which provide code for their implementations of the matching operation. Some examples include Alignment API (David et al. 2011) or S-Match (Shvaiko et al. 2009). In turn, a large collection of data sets commonly used for experiments can be found at http://oaei.ontologymatching.org, namely, Ontology Alignment Evaluation Initiative (OAEI), which is a coordinated international initiative that organizes annual evaluations of the increasing number of matching systems (Euzenat et al. 2011).

Key Applications

Ontology matching is an important operation in traditional applications, e.g., ontology evolution, ontology integration, data integration, data warehouses, or the recently raised linked data (Noy and Klein 2004; Isaac et al. 2009; Melnik 2004; Doan and Halevy 2005; Bellahsene et al. 2011; Gal 2011; Heath and Bizer 2011). These applications are characterized by heterogeneous models, e.g., database schemas or ontologies, which are analyzed and matched manually or semiautomatically at *design time*. In such applications, matching is a prerequisite to running the actual system.

There are some emerging applications that can be characterized by their dynamics, such as peer-to-peer information sharing, Web service composition, or query answering (Montanelli et al. 2011; Vaccari et al. 2012; Chang 2009). Such applications, contrary to traditional ones, require (ultimately) a *run time* matching operation and often take advantage of more explicit conceptual models (Shvaiko and Euzenat 2005).

Cross-References

- ▶ Graph Matching
- ▶ Linked Open Data

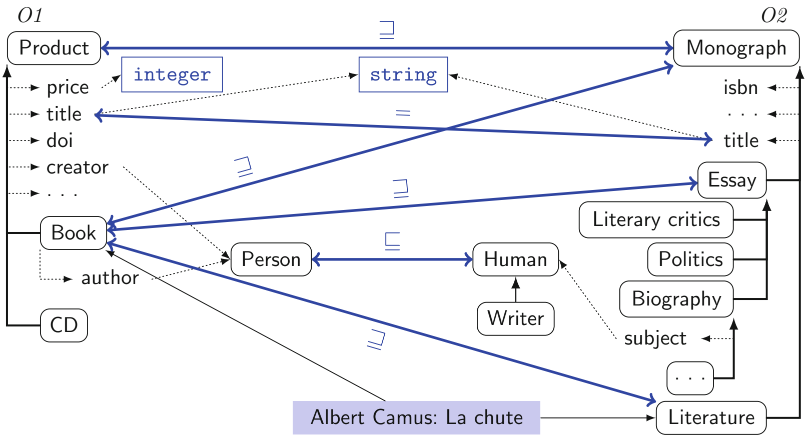

Ontology Matching, Fig. 2 Two simple ontologies and an alignment

▶ Service Discovery

▶ Web Ontology Language (OWL)

References

Bellahsene Z, Bonifati A, Rahm E (eds) (2011) Schema matching and mapping. Springer, Berlin/Heidelberg/New York

Chang K (2009) Deep-web search. In: Liu L, Özsu MT (eds) Encyclopedia of database systems. Springer, New York/London, pp 784–788

David J, Euzenat J, Scharffe F, Trojahn C (2011) The Alignment API 4.0. Semant Web J 2(1):3–10

Doan A, Halevy A (2005) Semantic integration research in the database community: a brief survey. AI Mag 26(1):83–94

Euzenat J, Shvaiko P (2007) Ontology matching. Springer, Berlin/New York

Euzenat J, Meilicke C, Stuckenschmidt H, Shvaiko P, Trojahn C (2011) Ontology alignment evaluation initiative: six years of experience. J Data Semant XV:158–192

Gal A (2011) Uncertain schema matching. Morgan & Claypool, San Rafael

Heath T, Bizer C (2011) Linked data: evolving the web into a global data space. Morgan & Claypool, San Rafael

Isaac A, Wang S, Zinn C, Matthezing H, van der Meij L, Schlobach S (2009) Evaluating thesaurus alignments for semantic interoperability in the library domain. Intell Syst 24(2):76–86

Melnik S (2004) Generic model management: concepts and algorithms. Springer, Berlin/New York

Montanelli S, Bianchini D, Aiello C, Baldoni R, Bolchini C, Bonomi S, Castano S, Catarci T, Antonellis VD, Ferrara A, Melchiori M, Quintarelli E, Scannapieco M, Schreiber FA, Tanca L (2011) The ESTEEM platform: enabling p2p semantic collaboration through emerging collective knowledge. J Intell Inf Syst 36(2):167–195

Noy N, Klein M (2004) Ontology evolution: not the same as schema evolution. Knowl Inf Syst 6(4):428–440

Shvaiko P, Euzenat J (2005) A survey of schema-based matching approaches. J Data Semant IV:146–171

Shvaiko P, Euzenat J (2013) Ontology matching: state of the art and future challenges. Trans Knowl Data Eng 25(1):158–176

Shvaiko P, Giunchiglia F, Yatskevich M (2009) Semantic matching with S-Match. In: Virgilio RD, Giunchiglia F, Tanca L (eds) Semantic web information management. Springer, Berlin/London, pp 183–202

Vaccari L, Shvaiko P, Pane J, Besana P, Marchese M (2012) An evaluation of ontology matching in geo-service applications. GeoInformatica 16(1):31–66

Open Encyclopedia

▶ Wikipedia Collaborative Networks

Open Innovation Social Networks

▶ Social Networking for Open Innovation

Open Source

▶ Gephi

Opinion Diffusion and Analysis on Social Networks

Cheng-Te Li, Hsun-Ping Hsieh,
Tsung-Ting Kuo, and Shou-De Lin
Graduate Institute of Networking and
Multimedia, National Taiwan University,
Taipei, Taiwan

Synonyms

Opinion mining; Preference propagation; Sentiment detection and analysis; Topic information diffusion

Glossary

Microblogging A broadcast medium in the form of blogging
Diffusion The process by which a new idea or new product is accepted by people
Sentiment Feelings and emotions
Preference An individual's attitude toward a set of objects

Introduction

With the bloom of the social networking and microblogging services, such as Facebook, Twitter, and LinkedIn, people can easily express their feelings and share ideas with friends. Through these services, messages posted by some persons can be seen, responded, or even broadcasted by others. It can be viewed as that through a social network service, opinions and the useful information are propagated from one to the other. With the time proceeds, opinions can be spread and evolved in a social network.

In this entry, we aim to review a number of studies discussing opinion detection, spread, and change on social networks. The existing studies about opinion mining on social networks can be categorized into three categories: (1) topic-based information diffusion, (2) preference propagation, and (3) sentiment detection. Given a *topic* (e.g., a news topic, an event, a public figure, a hyperlink, a picture, or video clip), the first task aims at modeling the diffusion of different topics in a social network (e.g., Petrovic et al. 2011; Zhu et al. 2011) and determines which nodes and links are important for the topic-based information diffusion (e.g., Budak et al. 2011; Leskovec et al. 2007). Second, by considering the personal *preference* (e.g., voting candidates in elections, choosing brands of certain commercial object, and coordinating ideas when making decisions) as an individual's opinion, researchers have been trying to model the spread of preference over individuals in a social network (e.g., Lou et al. 2013; Liu 2009). Finally, considering the sentiments (e.g., positive, negative, neutral; joy, surprise, sadness, anger, fear) of microblogging short messages, the third type of tasks aims to automatically detect the sentiment for a message, a person, or even a group of individuals (e.g., Go et al. 2009; Li et al. 2011). Next, we will cover some background on information diffusion and sentiment detection before moving into the discussion of the above three types of tasks.

Historical Background

Opinion diffusion and analysis usually require the integration of two major components: opinion analysis and information diffusion. Below, we will first introduce each of them independently.

Information Diffusion on Networks The diffusion of information on social networks has been studied for decades. Generally, the proposed strategies are either model driven or data driven. The model-driven strategies tend to explain the diffusion behavior using certain manually crafted, usually intuitive, models. Note that the past diffusion data are generally not fully exploited inside those models. For instances, models such as independent cascade (IC) model and linear threshold (LT) model (Kempe et al. 2003) are the foundation for a number of more sophisticated diffusion models. In the IC model, binary signal is propagated within a network with certain probability to activate other nodes. In the LT model, a real-value weight is propagated through the network while nodes that receive information surpassing a given threshold are considered as activated. On the other hand, the data-driven strategies usually utilize learning-based approaches to predict the future propagation given historical records of prediction (Fei et al. 2011; Galuba et al. 2010; Petrovic et al. 2011). Data-driven strategies usually fit the observation better comparing with the model-driven approaches due to its consideration of the past diffusion behavior (Galuba et al. 2010).

Sentiment Analysis The development of more sophisticated machine learning techniques together with the advance of the computational capability as well as the emergence of Internet provides the means and the massive data for researchers to design effective sentiment analysis methods (Pang and Lee 2004), which aims to identify certain feelings for the given text. The most commonly used sentiment is the *polarity*, positive or negative. There are more fine-grained sentiments such as "joy," "sad," "angry," and "fear." The traditional study of sentiment analysis in natural language processing can be viewed from three levels: document level (Pang and Lee 2008), sentence level (Kim and Hovy 2004), and world/phrase level (Esuli and Sebastiani 2006). To detect the sentiments of any given texts, most of existing studies focus on investigating diverse kinds of features (especially on lexical features) or design novel supervised learning methods for sentiment classification (Esuli and Sebastiani 2006; Kim and Hovy 2004; Pang and Lee 2004, 2008).

Topic-Based Information Diffusion

Topic-based information diffusion brings some new challenges to opinion mining on social networks. It aims to investigate how social relationships and content information interact with each other to affect the spread of topic information on social networks (Fei et al. 2011; Petrovic et al. 2011; Zhu et al. 2011), especially on the prediction of future diffusion of a certain topic. Specifically, many researches focus on predicting which messages are diffused to which neighbors in the network, given a certain topic. Most of existing studies assume that in order to train a model and predict the future diffusion of a topic, it is required to obtain historical records about how such topic had propagated in a social network. Petrovic et al. (2011) devise a time-sensitive passive-aggressive algorithm, which maintains a linear decision boundary to make the binary retweeting decision of the old topic (i.e., be retweeted or not be retweeted) in Twitter. Given a tweet, Zhu et al. (2011) design a logistic regression model to predict the retweeting decision of each user in the Twitter social network and eventually conclude that time decay can significantly affect the users' retweet behavior.

Some researchers, however, argue that a more realistic application in this direction is to forecast the propagation of novel or unseen topics (Kuo et al. 2012). For example, a company would like to know which users are more likely to be the source of "viva voce" of a newly released product for advertising purpose. A political party might want to estimate the potential degree of responses of a half-baked policy before deciding to bring it up to public. To achieve such goal, it is required to predict the future propagation behaviors of a topic even before any actual diffusion happens on this topic (i.e., no historical propagation data of this topic are available). Lin et al. (2011) propose a model-driven idea aiming at predicting the inference of implicit diffusions for novel topics. Kuo et al. (2012) propose a data-driven approach focusing on the prediction of explicit diffusion behaviors. Despite that no diffusion data of novel topics is available, their model extracts the latent information from data of known topics to for prediction.

Opinion Diffusion Analysis

Besides diffusion prediction, some researchers have worked on other kinds of topic-driven or topic-free opinion diffusion analysis on social networks. One common purpose is to design algorithms that can find important nodes that play critical roles of diffusion. Leskovec et al. (2007) propose to detect a set of social sensors such that their placements can efficiently detect the propagation of certain user-given topic in a social network. Lappas et al. (2010) propose to find a set of effectors that can faithfully reproduce the given active nodes observed by a certain topic campaign in a social network. Budak et al. (2011) propose to find a set of protectors, when a certain topic campaign is broadcasted over social network, to minimize the number of individuals that adopt the misinformation. Chen et al. (2011) model how the positive and negative topical influences compete to active and inactive nodes and then how to find a seed set to maximize the positive topical influence spread. Singer (2012) considers the fact that people have different interests and costs to become early adopters and integrates such idea into the selection of seed nodes in the influence maximization problem. Kimura et al. (2009) alternatively propose a method to find a limited number of links to block in a network for minimizing the propagation of undesired topical things, such as computer virus and malicious rumors. Scripps et al. (2007) propose to find a set of bridging nodes which can maximize the influence spread of diverse topical communities. Erdos et al. (2012) propose a way to find a set of filter nodes responsible for removing the redundant topics relayed through them, such that the information multiplicity phenomenon is minimized.

Preference Diffusion

With the success of viral marketing, people benefit from the power of crowd opinions and believe that individual options or preferences could be highly affected by acquaintances even though individuals generally possess intrinsic preferences. For example, in an election, people would argue and even attempt to convince others for their favorite candidates. With the rise of social networking service, people create and reply posts to promote their favorite candidates to the public. In this case, it is the preference toward some candidates that is propagated from some people to others and diffused over people in a social network. Up to date, there are only a few computational approaches with systematic and quantifiable studies on preference diffusion. Nevertheless, being able to model the human preference does possess its own value in the real-world applications. Social scientists might wonder how the opinion or preference exchange between friends can affect each other's evaluation on an object. In commercial areas, campaign companies can benefit from effective promotion of an item given a limited budget through a social network.

Bartholdi et al. (1992) first investigate the process to determine needed actions by the organizer to change candidates for manipulating election result. It is recognized as the classical social choice theory. However, they do not propose any computational models for the propagation and negotiation of preference between voters. Gibbard (1973) and Satterthwaite (1975) show that every election scheme with at least three possible outcomes is subject to individual's manipulation. The minority can manipulate the group decision to have a preferred outcome. However, their models assume that the individuals are independent of each other. Nevertheless, these works open the door for future research about preference diffusion.

Researchers quickly find that traditional information diffusion models such as LT and IC models cannot be adopted directly for preference diffusion. Liu (2009) attempted to check whether the preference distribution of an individual changes if the number of political experts in a social network increases. They use an agent-based model for simulation. In their model, the agent maintains a binary value toward a candidate (instead of a real value or ranking), and the model simply propagates such values to the other agents in the matrix of candidates. Yoo et al. (2009) propose a semi-supervised importance propagation model. Their idea is to add the original preference score of the candidate into the accumulated score obtained from the neighbors. The above solutions only partially handle the preference diffusion problem as the media for propagation in their frameworks are either binary or real values; while in the preference diffusion problem, each social entity possesses its own preference toward a set of candidates; therefore, it is a preference list that needs to be propagated in the network.

Lou et al. (2013) later design a framework to model the preference diffusion on social networks. They argue that a proper preference diffusion model should satisfy some properties: hyper-dimensional media, input dependency, deterministic convergence, and consensus. The properties are intuitively inspired by the natural real-world phenomena and are summarized as follows. First, the media (which represents preference toward different candidates) should be propagated throughout the process being a real-valued vector that sums to one, because in the real world, each node (or individual) usually has equal right in casting votes. Second, the preference distribution should be affected significantly by the initial intrinsic preference as well as the social network topology. Finally, the diffusion should converge eventually, and the common trends in the real world would finally appear after a great number of interactions. They proposed a preference diffusion model that can satisfy all properties among the existing models at present.

Sentiment Detection on Microblogs

With the popularity of online microblog services such as Facebook and Twitter, there are increasing amount of data available facilitating the study on sentiment analysis on this platform (Barbosa and Feng 2010; Bollen et al. 2010). Microblogging services generally possess some signature properties that are different from conventional weblogs and forum (Li et al. 2011). First, microblogs deal with almost *real-time* messaging, including instant information, expression of feelings, and quick ideas. It also provides a source of crowd intelligence that can be used to investigate common feelings or potential trends about certain news or concepts. However, such real-time property can lead to the production of an enormous number of messages that recipients must digest. Second, microblogging is *time traceable*. The temporal information is crucial because contextual posts that appear close together are, to some extent, correlated. Third, the style of microblogging posts tends to be *conversation based* with a sequence of responses. This phenomenon indicates that the posts and their responses are highly correlated in many respects. Fourth, microblogging is *social influenced*. Posts from a particular user can also be viewed by his/her friends and impact them (e.g., the empathy effect) implicitly or explicitly. Therefore, posts from friends in the same period may be correlated sentiment-wise as well as content-wise.

The sentiment detection on social networks can be defined as to determine the sentiment associated a length-limited microblog post. The main strategies of microblog sentiment detection (Go et al. 2009; Li et al. 2009; Bartholdi et al. 1992; Bermingham and Smeaton 2010; Bifet and Frank 2010; Davidov et al. 2010; Pak and Paroubek 2010; Sun et al. 2010) can be divided into two categories. The first focuses on feature engineering to investigate various textual and linguistics features (e.g., n-gram, POS tagging, prior subjectivity and polarity of words, punctuation, word pattern) or other generic feature (e.g., URL, hashtag, emoticon). The second investigates different kinds of learning algorithms. The commonly used methods for detecting sentiments are either unsupervised (e.g., keyword matching and k-nearest neighbor, KNN) or supervised (e.g., probabilistic graphical models, discriminant models). Here, we summarize the related works from the perspective of features and learning methods using Table 1.

In general, supervised methods have better performance on sentiment detection. Such approach has heavy demand on the quality of training data. Therefore, it becomes crucial to be able to generate accurate sentiment labels for microblog posts. The most common approach for labelling the microblog posts is using emoticons, which is a kind of facial expression composed of characters, such as :-), :-o, :-D, :-(, and ^^ Read (2005) and Go et al. (2009) use emoticons to label Twitter data as positive or negative sentiment automatically. As a result, a great number of labelled training data can be obtained. Recent advances also follow such labelling approach (Bifet and Frank 2010; Davidov et al. 2010; Pak and Paroubek 2010; Sun et al. 2010).

Most of the above works consider only the post content as features for sentiment detection. More recently, Li et al. (2011) improve the performance through considering three advanced features: *response* factor, *contextual* factor, and

Opinion Diffusion and Analysis on Social Networks, Table 1 Summary of related works that detect sentiments in microblogs

	Feature	Method
Pak and Paroubek (2010)	Statistic counting of adjectives	Naive bayes
Riley (2009)	n-grams, smileys, hashtags, replies, URLs, usernames, emoticons	Naive bayes
Go et al. (2009)	Usernames, sequential patterns of keywords, POS tags, n-grams	Naive bayes, maximum entropy, SVM
Li et al. (2009)	Several dictionaries about different kinds of keywords	Keyword matching
Bartholdi et al. (1992)	Retweets, hashtag, replies, URLs, emoticons, uppercases	SVM
Sun et al. (2010)	Keyword counting and Chinese dictionaries	Naive bayes, SVM
Davidov et al. (2010)	n-grams, word patterns, punctuation information	k-nearest neighbor
Bermingham and Smeaton (2010)	n-grams and POS tags	Binary classification

social factor. First, they believe the sentiment of a post is highly correlated with (but not necessary similar to) that of responses to the post. For example, an angry post usually triggers angry responses, but a sad post usually solicits supportive responses. Second, they assume that the sentiment of a microblog post is correlated with the author's previous posts (i.e., the "context" of the post). Third, they also assume that the friends' sentiments are correlated with each other. This is because friends affect each other, and they are more likely to be in the same circumstances, and thus enjoy/suffer similarly. Eventually, Li et al. devise a Markov transition model to integrate such information (i.e., response, contextual, and social) and obtain a better performance on detecting sentiments in microblogging social networks.

In addition, there are some variations on microblog sentiment detection. Jiang et al. (2011) aim to detect the sentiment of posts given a certain given topic, such as "Obama" and "iPad." They propose several rules to generate *topic-dependent* features. Taking the message "I love iPad" as an example, they create and use the word pattern "love_ipad" as a composite feature. Calais et al. (2011) employ the transfer learning technique to detect the sentiment for a query topic. They first learn the bias of users toward the specific topic and transfer the information learned to the problem of sentiment classification of positive and negative sentiment in the microblogging social network.

Future Directions

- **Topic-Based Information Diffusion**
  – Considering topics evolve over time, one possible future direction is to model and track the change and diffusion of the topic of interest on a social network.
  – Exploiting the topic-based information diffusion for other kinds of social network analysis tasks such as community detection and link prediction.
  – Most of the existing topic-based information diffusion models assume homogeneous social network, while how to expand them to heterogeneous social networks can be a promising future direction.
- **Preference Diffusion**
  – Integrating the personal attributes and social content into the modeling of preference diffusion rather than assuming each person's preference or interests are provided in advance.
  – As the existing methods for preference diffusion are mostly model driven, it is preferable to see the development of certain data-driven approach for this task.

- **Sentiment Detection**
  - Analyze the relationship between topic diffusion and sentiment diffusion on social networks.
  - Brainstorm useful applications for microblog sentiment detection (e.g., opinion leader identification, social-based advertisement).
  - Incorporating natural language processing methods to handle slangs, jargons, or any new form or writing derived due to the space limitation of microblog posts.

Acknowledgments

This work was supported by National Science Council, National Taiwan University, and Intel Corporation under Grants NSC101-2911-I-002-001, NSC101-2628-E-002-028-MY2, and NTU102R7501.

Cross-References

► Actionable Information in Social Networks, Diffusion of
► Modeling Social Preferences Based on Social Interactions
► Sentiment Analysis in Social Media
► Social Influence Analysis
► Topic Modeling in Online Social Media, User Features, and Social Networks for
► Transforming and Integrating Social Networks and Social Media Data
► Twitter Microblog Sentiment Analysis
► User Sentiment and Opinion Analysis

References

Barbosa L, Feng Robust J (2010) Sentiment detection on twitter from biased and noisy data. In proceedings of international conference on computational linguistics COLING, Mumbai, pp 36–44

Bartholdi JJ, Tovey CA, Trick MA (1992) How hard is it to control an election. Math Comput Model 16:8–9, 27–40

Bermingham A, Smeaton AF (2010) Classifying sentiment in microblogs: is brevity an advantage? In: Proceedings of acm international conference on information and knowledge management CIKM, Toronto, pp 1183–1186

Bifet A, Frank E (2010) Sentiment knowledge discovery in twitter streaming data. In: Proceedings of international conference on discovery science (DS'10), Canberra, pp 1–15

Bollen J, Pepe A, Mao H (2010) Modeling public mood and emotion: twitter sentiment and socio-economic phenomena. In: Proceedings of acm international world wide web conference WWW, Raleigh

Budak C, Agrawal D, Abbadi AE (2011) Limiting the spread of misinformation in social networks. In: Proceedings of ACM international world wide web conference WWW, Hyderabad, pp 665–674

Calais P, Veloso A, Meira W Jr, Almeida V (2011) From bias to opinion: a transfer-learning approach to real-time sentiment analysis. In: Proceedings of ACM SIGKDD international conference on knowledge discovery and data mining KDD, San Diego

Chen W, Collins A, Cummings R, Ke T, Liu Z, Rincon D, Sun X, Wang Y, Wei W, Yuan Y (2011) Influence maximization in social networks when negative opinions may emerge and propagate. In: Proceedings of SIAM international conference on data mining SDM, Mesa, pp 379–390

Davidov D, Tsur O, Rappoport A (2010) Enhanced sentiment learning using twitter hashtags and smileys. In: Proceedings of international conference on computational linguistics COLING, Mumbai, pp 241–249

Erdos D, Ishakian V, Lapets A, Terzi E, Bestavros A (2012) The filter-placement problem and its application to minimizing information multiplicity. In: Proceedings of the VLDB endowment PVLDB, Istanbul, vol 5, no 5, pp 418–429

Esuli A, Sebastiani F (2006) SentiWordNet: a publicly available lexical resource for opinion mining. In: Proceedings of LREC, Genoa

Fei H, Jiang R, Yang Y, Luo B, Huan J (2011) Content based social behavior prediction: a multi-task learning approach. In: Proceedings of ACM international conference on information and knowledge management CIKM, Glasgow

Galuba W, Aberer K, Chakraborty D, Despotovic Z, Kellerer W (2010) Outtweeting the twitterers – predicting information cascades in microblogs. In: Proceedings of international conference on online social networks, Boston, MA, USA

Gibbard A (1973) Manipulation of voting schemes: a general result. Econometrica 41(4):587601

Go A, Bhayani R, Huang L (2009) Twitter sentiment classification using distant supervision. Technical report, Stanford University

Jiang L, Yu M, Zhou M, Liu X, Zhao T (2011) Target-dependent twitter sentiment classification. In: Proceedings of annual meeting of the association for computational linguistics ACL 2011, Portland

Kempe D, Kleinberg J, Tardos E (2003) Maximizing the spread of influence through a social network. In: Proceedings of ACM SIGKDD international conference on knowledge discovery and data mining KDD, Washington, DC

Kim S-M, Hovy E (2004) Determining the sentiment of opinions. In: Proceedings of coling, Geneva

Kimura M, Saito K, Motoda H (2009) Blocking links to minimize contamination spread in a social network. ACM Trans Knowl Discov Data TKDD 3(2):9

Kuo T-T, Hong S-C, Lin W-S, Peng N, Lin S-D, Lin W-F (2012) Exploiting latent information to predict diffusions of novel topics on social networks. In: Proceedings of annual meeting of the association for computational linguistics ACL, Jeju Island

Lappas T, Terzi E, Gunopulos D, Mannila H (2010) Finding effectors in social networks. In: Proceedings of ACM SIGKDD international conference on knowledge discovery and data mining KDD, Washington, DC, pp 1059–1068

Leskovec J, Krause A, Guestrin C, Faloutsos C, Van-Briesen J, Glance N (2007) Cost-effective outbreak detection in networks. In: Proceedings of ACM SIGKDD international conference on knowledge discovery and data mining KDD, San Jose, pp 420–429

Li S, Zheng L, Ren X, Cheng X (2009) Emotion mining research on microblog. In: Proceedings of IEEE symposium on web society, Shuo Chen, PP 71–75

Li C-T, Wang C-Y, Tseng C-L, Lin S-D (2011) Meme-Tube: a sentiment-based audiovisual system for analyzing and displaying microblog messages. In: Proceedings of annual meeting of the association for computational linguistics ACL, Portland

Lin CX, Mei QZ, Jiang YL, Han JW, Qi SX (2011) Inferring the diffusion and evolution of topics in social communities. In: Proceedings of the IEEE international conference on data mining, Vancouver

Liu FC (2009) Modeling political individuals using the agent-based approach: a preliminary case study on political experts and their limited influence within communication networks. J Comput 19(4):819

Lou J-K, Wang F-M, Tsai C-H, Hong S-C, Kung P-H, Lin S-D (2013) Modeling the diffusion of preferences on social networks. In: Proceedings of SIAM international conference on data mining SDM, Austin

Pak A, Paroubek P (2010) Twitter as a corpus for sentiment analysis and opinion mining. In: Proceedings of international conference on language resources and evaluation LREC, Valletta, pp 1320–1326

Pang B, Lee L (2004) A sentimental education: sentiment analysis using subjectivity summarization based on minimum cuts. In: Proceedings of annual meeting of the association for computational linguistics ACL, Barcelona, pp 271–278

Pang B, Lee L (2008) Opinion mining and sentiment analysis. Found Trends Inf Retr 2(1–2):1–135

Petrovic S, Osborne M, Lavrenko V (2011) RT to win! Predicting message propagation in twitter. In: Proceedings of international AAAI conference on weblogs and social media ICWSM, Barcelona

Read J (2005) Using emoticons to reduce dependency in machine learning techniques for sentiment classification. In: Proceedings of annual meeting of the association for computational linguistics ACL 2005, University of Michigan

Riley C (2009) Emotional classification of twitter messages. Technical report, UC Berkeley

Satterthwaite MA (1975) Strategy-proofness and arrow's conditions: existence and correspondence theorems for voting procedures and social welfare functions. J Econ Theory 10(2):187–217

Scripps J, Tan P-N, Esfahanian A-H (2007) Node roles and community structure in networks. In: Proceedings of ACM SIGKDD international workshop on social network analysis SNA-KDD, Paris, pp 26–35

Singer Y (2012) How to win friends and influence people, truthfully: influence maximization mechanisms for social networks. In: Proceedings of ACM international conference on web search and data mining WSDM, Seattle, pp 733–742

Sun YT, Chen CL, Liu CC, Liu CL, Soo VW (2010) Sentiment classification of short chinese sentences. In: Proceedings of conference on computational linguistics and speech processing rocling, Nantou, pp 184–198

Yoo S, Yang Y, Lin F, Moon I-C (2009) Mining social networks for personalized email prioritization. In: Proceedings of ACM SIGKDD international conference on knowledge discovery and data mining KDD, Paris, pp 967–976

Zhu J, Xiong F, Piao D, Liu Y, Zhang Y (2011) Statistically modeling the effectiveness of disaster information in social media. In: Proceedings of IEEE global humanitarian technology conference, Seattle

Opinion Mining

- ▸ Multi-classifier System for Sentiment Analysis and Opinion Mining
- ▸ Opinion Diffusion and Analysis on Social Networks
- ▸ Sentiment Analysis in Social Media
- ▸ User Sentiment and Opinion Analysis

Optimal Scaling

- ▸ Correspondence Analysis

ORA

▶ ORA: A Toolkit for Dynamic Network Analysis and Visualization

*ORA

▶ ORA: A Toolkit for Dynamic Network Analysis and Visualization

ORA: A Toolkit for Dynamic Network Analysis and Visualization

Kathleen M. Carley
Carnegie Mellon University, Pittsburgh,
PA, USA
Netanomics, Sewickley, PA, USA

Synonyms

ORA; *ORA; ORA-NetScenes; ORA toolkit

Glossary

Social Network Analysis Graphical, statistical, and visualization metrics, algorithms, and techniques for analyzing structural data that can be represented as nodes and relations. Social network analysis is also referred to as network analysis, dynamic network analysis, network science, SNA, and DNA

Social Media Data generated by online social networking tools such as Twitter, Facebook, or Foursquare. Social media networks are networks derived from social media data such as the Twitter retweet network. Social media are also referred to as online sources, open source, and e-media

Dynamic Networks Networks that vary through time. An example is the network of who talks to whom within a company by day. Dynamic networks are also referred to as temporal networks, time variant networks, and dynamical networks

Spatial Networks Networks embedded in space such that each node has one or more locations at which it occurs. An example is the network of who interacts with whom among political activists such that each activist is also linked to the locations at which they have taken part in demonstrations. Spatial networks are also referred to as geo-spatial networks, geo-intelligence networks, geographical networks, networks through space, and spatially embedded networks

Meta-Networks A network of networks in which there are generally multiple classes of nodes and multiple classes of links. Meta-networks are also referred to as high-dimensional networks and geo-temporally embedded meta-networks

Definition

ORA: A network analysis toolkit for graphical, statistical, and visual analytics on both social networks and high-dimensional networks that can vary by time and/or space. ORA is a full function network analytics package that supports the user in creating, importing, exporting, manipulating, editing, analyzing, comparing, contrasting, and forecasting changes in one or more networks. ORA is a multi-platform network toolkit that can operate in stand-alone mode or as a service within a web architecture.

ORA ID card:

* Tool name, title – ORA
* Creation year – 1995
* Author – Kathleen M. Carley
* Scope – general
* Copyright – academic student version, commercial full professional version
* Type – program
* Size limits – over one million nodes in batch
* Platforms – PC, LINUX, and MAC
* Programming language – Java GUI and C++ backend
* Orientation – multi-disciplinary

Introduction

ORA is a network analytic tool developed by CMU and Netanomics that allows the user to fuse, analyze, visualize, and forecast behavior given network data. Using ORA the user can reason about networks at the node, group, or network level. ORA includes a wide range of capabilities that supports multiple types of analyses; e.g., the user can analyze a social network using standard social network metrics; examine geo-temporal networks; identify key actors, key topics, and hot spots of activity; identify communities, subgroups, and patterns of interest; examine changes in which nodes are key; and examine changes in group membership. ORA is designed to support the analysis networks that vary in size (e.g., from small to large networks) and type (e.g., the networks can be social, communication, semantic, task, or other), and the networks can be high dimensional, aka meta-networks, that are dynamic and may be embedded in geographic regions.

Key Points

There are several ways in which ORA differs from other social network toolkits. First, it actively supports high-dimensional network data, often referred to as meta-network data, including changed in such data through time and embedded in space. Thus, ORA has not just one-mode metrics but two-mode and multimode networks. Many of these metrics and algorithms are only available in ORA such as measures of cognitive demand, redundancy, and the fuzzy grouping algorithm – FOG. Second, ORA supports analysis on very large network data sets. There is both a GUI version of ORA and a batch mode version – the latter of which has been used with networks with 10^6 nodes. Third, ORA is interoperable with a large number of other tools. For example, data in CSV, TSV, UCINET, or Twitter JSON formats and many other formats can be directly imported. Further, ORA can export to Google Earth or to KML files, thus enabling interoperation with GIS tools. Fourth, ORA has been designed to reduce training time and effort on the part of the user. For example, output is organized by topic which enables the user to select a topic and then all the major metrics used for that topic are automatically computed and printed to a web page. This means that the user does not have to remember what metrics to run. As another example, common sequences that are done by users are automated and occur with a single keystroke. Fifth and finally, ORA has an integrated online help system, tool tips, QuickStart Guide (Carley, 2013, ORA: quick start guide, unpublished manuscript), user's guide (Carley et al. 2013a), and associated Google groups to provide help. Moreover, two textbooks provide guidance in network analysis and show those analyses in ORA (Everton 2012; McCulloh et al. 2013).

Historical Background

ORA was originally developed as a network toolkit to support reasoning about changes in network within organizations from a meta-network perspective. This led to the organizational management report which not only identifies key actors but also assesses the level of shared situation awareness, individuals who might be over or underworked, and other factors that can lead to organizational problems such as the identification of insider threats. ORA has evolved enormously from this beginning.

ORA was extended to support link analysis in addition to social network analysis, for complex sociocultural systems. This led to improved visualization, data entry, and multimode metrics. This also led to support for reasoning about the network in context, placing networks on maps, and visualizing the trails that led to the formation of network ties.

ORA was extended to support semantic network analysis and to enable the user to build networks from text. That led to the inclusion of new rhetorical-based network metrics and interoperation with text-mining software.

ORA was extended to support network pattern of life analysis. This led to the inclusion of

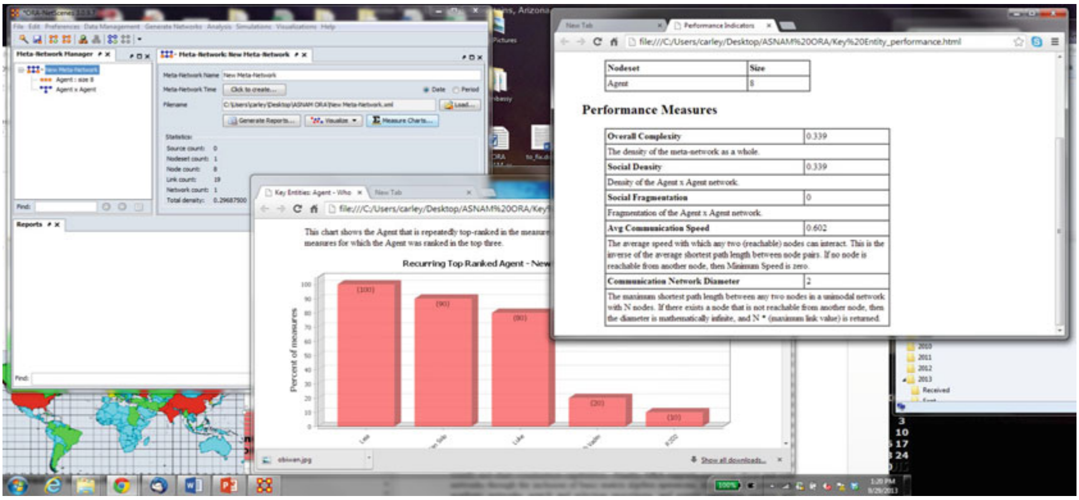

Meta-Network Name: New Meta-Network
Meta-Network Time: Click to create... ● Date ○ Period
Filename: C:\Users\carley\Desktop\ASNAM ORA\New Meta-Network.xml

Statistics:
Source count: 0
Nodeset count: 1
Node count: 8
Link count: 19
Network count: 1
Total density: 0.29687500

file:///C:/Users/carley/Desktop/ASNAM%20ORA/Key

This chart shows the Agent that is repeatedly top-ranked in the measure measures for which the Agent was ranked in the top three.

file:///C:/Users/carley/Desktop/ASNAM%20ORA/Key%20Entity_performance.html

Nodeset	Size
Agent	8

Performance Measures

Overall Complexity	0.339
The density of the meta-network as a whole.	
Social Density	0.339
Density of the Agent x Agent network.	
Social Fragmentation	0
Fragmentation of the Agent x Agent network.	
Avg Communication Speed	0.602
The average speed with which any two (reachable) nodes can interact. This is the inverse of the average shortest path length between node pairs. If no node is reachable from another node, then Minimum Speed is zero.	
Communication Network Diameter	2
The maximum shortest path length between any two nodes in a unimodal network with N nodes. If there exists a node that is not reachable from another node, then the diameter is mathematically infinite, and N * (maximum link value) is returned.	

ORA: A Toolkit for Dynamic Network Analysis and Visualization, Fig. 1 Illustration of ORA in operation

spectral analytics for networks and network change detection.

Today, ORA supports full spectrum network analytics and visualization for small to very large networks. It has evolved into a widely used tool that is interoperable with many other technologies. It supports network analytics across a large number of domains and has been used in both theoretical and very applied contexts.

ORA

ORA is intended to be a comprehensive analytic and visualization engine for network analytics and forecasting that transcends traditional disciplinary boundaries as it supports standard social network analysis, link analysis, geo-network analysis, and dynamic network analysis. The algorithms in ORA are thus from a large number of fields (Carley and Pfeffer 2012). In a similar vein, ORA supports data entry from and exports to many other tools common in fields where network analysis is used. ORA supports analysis of many types of networks – e.g., social networks, social media networks, communication networks, activity networks, and semantic networks. As such, there are many metrics in ORA that are specialized to certain types of network data – e.g., geo-temporal networks. ORA is organized by the questions the user asks, not the metrics. As such, common workflows are automated and all metrics relevant to a particular question are automatically included in a report providing the network answer to the question of interest. Visual analytics are a key feature in ORA which employs both graphical, statistical, and visual analytic algorithms to help the user assess, visualize, and forecast behavior for social networks or high-dimensional networks that can be dynamic or spatially situated. ORA supports 2D and 3D network visualization, geo-spatial network visualization, heat maps, loom or trail visualization, nodel (a node variant of wordl), and chart visualization capabilities. Finally, ORA supports reasoning about and with networks through the inclusion of basic matrix algebra operations, algorithms for generating synthetic networks, search and selection procedures, and simple simulation engines and comparative statics for examining the potential impact of change. The basic interface to ORA and part of several reports are shown in Fig. 1.

What Can Be Done in ORA?

Using ORA the analyst can identify key actors, key topics, and key locations; characterize and visualize networks; assess changes in the networks and the key actors, key topics, and key locations; visualize trends; identify patterns in the networks; examine the networks and key actors, key topics, and key locations in terms of where they are by using the geo-spatial mapping functions; and multiple other tasks. Within ORA there are wizards for many multistep functions such as data importing, cleaning a visualization, and running common metrics. Many of these wizards lead the user through a set of decision when analyzing the data and generating reports. A list of the reports, functions, and generators in ORA is in Table 1. The system is organized by questions that the user might ask of the data, and many common workflows are built in to reduce time to do standard analyses. For example, a typical network question is what are the key nodes? The Key Entity report automatically runs various centrality metrics and 2-mode networks for identifying key nodes. Which metrics are run depends on the type of entity – i.e., whether the node is a who, what, where, why, or how. Figure 1 shows parts of the Key Entity report.

ORA can also go beyond node analysis. For example, it contains a wide range of grouping and clustering and pattern and community detection algorithms. These are available in the Locate Group report. Moreover, within the visualizer the user can color nodes by the groups they are in. ORA also supports change detection and spectral analysis for network metrics at the node and graph level. Using these techniques, dynamic data can be assessed, and changes and regularities in the way network metrics and nods behave over time can be characterized.

User Control and Awareness

There are an increasing number of people who want to analyze networks; however, many of them have had little training in network analytics. This can lead them to make errors in analysis without realizing it such as identifying a nodes centrality when the "network" contains multiple types of links and nodes. ORA tries to minimize this by providing more user guidance, not allowing functions that are known to not be applicable, and making the user be explicit about what node class a node is in and the nature of the relationship. Moreover, ORA makes explicit whether the results are based on the status of the underlying network such as whether it is binary or symmetrized. Studies show that ORA results map onto other common widely used network analysis tools when the user manipulates the data to match the way the other tool alters the data by default (Wei et al. 2011).

High-End Visualization

ORA enables the user to interact with the data visually in a number of ways. Many of these are shown in Fig. 2. One of the keys is that the visualizer and the data editor are linked so that a change in one effects a change in the other. Thus, the user can directly enter or edit data in the visualizer, or the user can directly see the impact on the network image as nodes are merged or deleted in the editor. In addition, standard graphing tools are used in reports for things such as pie charts, spider graphs, scatter plots, and histograms.

Temporal Analysis and Geo-Temporal Analysis

A unique feature of ORA is that it enables the user to examine networks as they change through time and the movement of networks through space. Many of the reports, if the user has two time periods, automatically do a comparison. In the visualizer, the network over time allows the user to step through changes in the network in movie like fashion and view changes over time for measures or nodes of the user's choice.

For multi-time period, data spectral analysis and change detection are available (McCulloh et al. 2012). Spectral analysis supports the user in assessing the regularities and anomalies in temporal network data. Graph or node level metrics can be examined over time and the "patterns of life" identified, such as the drop in Twitter activity at 4 am or an increase in the centrality of a team lead in e-mail traffic immediately after a group meeting. Change detection is a forensic technique that supports the user in identifying when a change in a dynamic network occurred that led to the current signal.

There are two different techniques for a what-if analysis that support reasoning about change.

ORA: A Toolkit for Dynamic Network Analysis and Visualization, Table 1 ORA analytic capabilities

Reports	Reports	Network generators
All measures	Locate groups	Stylized
Belief propagation	Management	Erdos-Renyi
Capabilities	Missing links	Core periphery
Communications network	Network distribution	Scale free
assessment	Optimizer	Cellular
Communicative power	Part of speech	Lattice
Communicators	Potential errors	Small world
Context	QAP/MRQAP analysis	Fixed degree distribution
Core network	Role view	Calculated from existing
Critical sets	Semantic network	data
Detect spatial patterns	Shortest path	Expected interaction
Drill down	Simmelian ties analysis	Expertise
Geo-spatial assessment	Sphere of influence	Similarity
Group talk	Standard network analysis	Distinctiveness
Hot topics (content analysis)	Statistical change detection	Resemblance
Immediate impact	Statistical distribution	Command and control
Influence net	Topic analysis	Infer beliefs
K-centrality	Trails	Influence network
Key entity	Trails analysis	Matrix algebra
Large scale	Twitter	User entered network
Local patterns	Unique trails report	Through editor
		Through visualizer
Grouping algorithms and pattern identifiers	Visualizers	Functions
Grouping:	Measure charts	Correspondence analysis
Girvan-Newman	View measures over time	Geary C and Moran I
Louvain	View networks over time	Analysis
Components	Network drill down	Simulators
Concor	Node cloud	Micro-sims
Johnson Hierarchical	Color grid	Information diffusion
K-means	Network block	Disease propagation
Attribute-based	Geo-spatial networks	Goods dispersion
K-Fog	Region viewer	OrgAhead
Alpha-Fog	Trails	Near term analysis
Local Patterns:	Trails in GIS	
Minimum spanning tree	2D network	
Cliques	3D network	
Hidden links		
Stars		
Checkerboards		
Balls and chains		
Cycles		

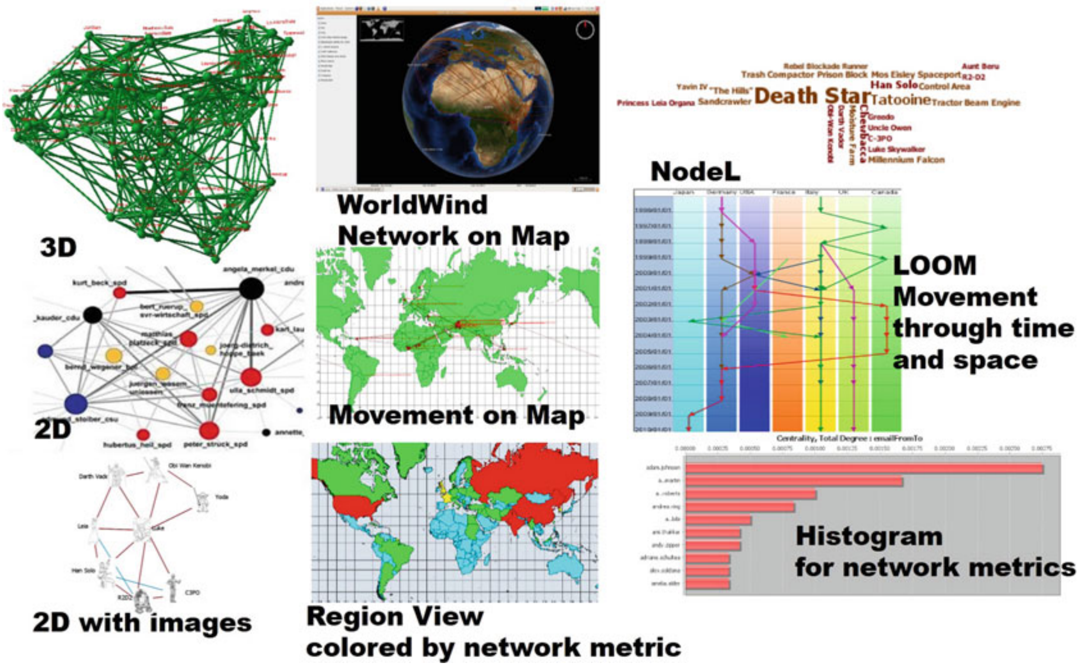

ORA: A Toolkit for Dynamic Network Analysis and Visualization, Fig. 2 Illustrative visualization capabilities

The first of these is the immediate impact report where the user can select a set of nodes or links to remove and then a static comparison is run between the original network and the hypothetical new network. The second technique is the near-term impact report where ORA calls the network simulator Construct already instantiated with data from the network being analyzed. Then a simple simulation-based virtual experiment is run.

When the network is geo-temporal, it is often useful to think in terms of trails. The trail format captures data of the form who/what was where when. An example would be information on people moving through space. Another example, in which the where is virtual, would be authors publishing papers in different journal through time. In this case the "where" is the journal. ORA supports changing data from trails to networks and networks to trails. Moreover, it supports direct analysis of trails in the visualization subtool – loom. There are also specialized reports, one that identifies the geo-temporal clusters using a trial clustering algorithm (Gullapalli and Carley 2013).

Big Data Analytics

ORA is designed to support the analysis of networks varying in size through a set of wizards and optimizations that change how the system operates depending on the size of the data. Recently, approximation algorithms have been added such as the k-centrality algorithms (Pfeffer and Carley 2012). Special features for very big data include allowing the user to choose to not run slow metrics such as betweenness and automated grouping prior to visualization for very large networks. The professional version of ORA can be run either through the interface or from the command line. All metrics have been optimized for sparse matrices. Tests have shown ORA capably of running, through command line or on very large multi-processor machines analytics on networks with 10^6 nodes and 10^7 links.

Interoperability

ORA is designed to work synergistically with other applications. To begin with, there is the data to model workflow (Carley et al. 2007) designed so that the user could step through extracting networks from texts (AutoMap, Carley et al. 2013c), then analyze those networks (ORA, Carley et al. 2013a), and then simulate changes in those networks (Construct, Carley et al. 2012a; Carley et al. 2009b). This interoperability is made possible through the use of DyNetML, a variant of GraphML that supports large-scale dynamic network of high-dimensional data. The format is completely open. This workflow has been recently improved, and the entire process increased in speed (Carley et al. 2012c). Secondly, ORA can import from and export to a number of formats. For example, data can be imported in a number of ways, including from Excel, CSV, TSV, UCINET, Pajek, Analyst Notebook, PenLink, Personal Brain, GraphML, shape files, and various Twitter JSON formats. Images can be saved in pdf, tiff, jpeg, svg, and a proprietary editable format. Shape files in KML can be exported from the geo-spatial visualizer which can also port to NASA World Wind and Google Earth. Reports are in HTML and work with most browsers, CSV and often PDF formats are also available.

Key Applications

ORA has been used in a wide number of contexts. Indeed, new applications appear frequently. Selected applications are listed in Table 2. The items in Table 2 are meant to be illustrative, not exhaustive. In essence, the use of ORA is only limited by the user's imagination. The meta-network flexibility admits analysis of many kinds of data, from social networks of who interacts with whom, to Twitter networks, to studies of who is critical and what they do, where, and when. Illustrative examples of how ORA has been used in various contexts include assessment of covert activity (Carley et al. 2009a), citation analysis (Kas et al. 2012), social media analysis (Carley et al. 2013b), public health (Merrill et al. 2012), hospital safety outcomes (Effken et al. 2013), structure of ethnic violence (Van Holt et al. 2012), and terror groups (Kenney et al. 2012).

Future Directions

ORA is evolving in several ways. These include supported platforms, data entry and manipulation, algorithms, big data analytics, statistical reasoning, and interoperation.

Platforms: Currently, there is both a free student version and a more comprehensive professional COTS system able to handle larger data faster and with better 3D and geo-spatial capabilities. Multi-platform versions for the Mac and Linux are in beta version. An enterprise version that supports scripting and a web version for operation in Ozone exist and are being extended.

Data Entry and Manipulation: ORA is currently being extended to directly read in additional JSON formats for Twitter. New optimized workflows for merging and deleting nodes for multiple networks at once are in development.

Algorithms: Currently there are over 150 metrics and 20 grouping algorithms in ORA. New incremental and approximation metrics and new group level measures are being added.

Big Data Analytics: ORA has been optimized for large-scale sparse matrices. The scripting version can handle 10^6 nodes. Metrics are ordered by their speed so the user can choose through the interface to run only fast measures, to run all but the slowest measures, or to run all measures. Currently, new metrics to support even faster processing are being implemented. These include the incremental metrics for closeness (Kas et al. 2013a) and betweenness (Kas et al. 2013b). In addition, since ORA supports dynamic analysis, the current system is being refactored to support improved speed for handling multiple time periods at once. Again, the incremental metrics support temporal analysis as well.

Statistical Reasoning: ORA currently has a version of QAP and MRQAP. The MRQAP routine is being extended to allow the user to specify and run multiple models with a scripting approach. Further, an ERGM module for ORA is currently

ORA: A Toolkit for Dynamic Network Analysis and Visualization, Table 2 Illustrative uses of ORA

Organizations	Law and military	Health	General planning	Text analysis
Design	Counter-terrorism	Fall and medication errors	Housing design	Semantic networks
Team management	Counter-narcotics	Public health coordination	Barrier reef protection	Topic groups
Silo identification	Human trafficking	Disaster Response	Global cyber	Hashtag networks
Productivity Assessment	Piracy	Disease spread	Security Curriculum	Communicative reach analysis
Command level cyber security	Counter-IEDs	Health-related Twitter	Evaluation	
Software design teams	Atmospherics	Disaster related SMS	P2P networks	
Enron e-mail	Sociocultural geography	Journal networks		
Citation networks	Open-source analysis			
Interlocking directorates	Global hostility and alliance Networks			

under development. Finally, most reports provide some guidance as to whether the metrics observed are within or beyond the bounds expected, given a normal distribution; comparison against a second baseline – a scale-free distribution – is underway. Further, key reports are being revised using "color" to quickly identify those nodes whose values are significantly greater than or less than other nodes in the same network on the metric of interest.

Interoperation: ORA is intended to work synergistically with many other tools. Currently, ORA can import from many other social network analytic tools, CSV, TSV, and some forms of JSON, and can output in CSV, HTML, and KML and is interoperable with many tools such as the Analyst Notebook and ArcGIS. Wizards for increased interoperation with technologies that extract social media or do forecasting are currently being developed. APIS to support integration into workflow systems and to support user addition of metrics are planned.

Acknowledgments

This research is part of the Dynamics Networks project in the Center for Computational Analysis of Social and Organizational Systems (CASOS; http://www.casos.cs.cmu.edu) at Carnegie Mellon University. The development of ORA was supported in part by CMU, Carley Technologies Inc, KAIST, ONR, AFOSR, ARL, ARO, NSF, and private donations.

Cross-References

- ▶ Analysis and Visualization of Dynamic Networks
- ▶ Anonymization and De-anonymization of Social Network Data
- ▶ Centrality Measures
- ▶ Classical Algorithms for Social Network Analysis: Future and Current Trends
- ▶ Combining Link and Content for Community Detection
- ▶ Community Detection in Social Network: An Experience with Directed Graphs
- ▶ Community Detection, Current and Future Research Trends
- ▶ Community Identification in Dynamic and Complex Networks
- ▶ Inferring Social Ties
- ▶ Motif Analysis
- ▶ Network Anomaly Detection Using Co-clustering

- ▶ Pajek
- ▶ Spatial Networks
- ▶ Spectral Analysis
- ▶ Spectral Evolution of Social Networks
- ▶ UCINET
- ▶ Visual Methods and Tools for Social Network Analysis
- ▶ Visualization of Large Networks

References

Carley KM, Diesner J, Reminga J, Tsvetovat M (2007) Toward an interoperable dynamic network analysis toolkit. Adv Inf Shar Data Min Collab Syst (DSS special issue on cyberinfrastructure for homeland security) 43(4):1324–1347

Carley KM, Martin MK, Hirshman B (2009a) The etiology of social change. Top Cogn Sci 1.4:621–650

Carley KM, Martin MK, Hancock JP (2009b) Dynamic network analysis applied to experiments from the decision architectures research environment. In: Advanced decision architectures for the warfighter: foundation and technology, chapter 4

Carley KM, Pfeffer J (2012) Dynamic network analysis (DNA) and ORA. In: Schmorrow DD, Nicholson DM (eds) Advances in design for cross-cultural activities part I. CRC, Boca Raton, pp 265–274

Carley KM, Filonuk DT, Joseph K, Kowalchuck M, Lanham MJ, Morgan GP (2012b) Construct user guide, Carnegie Mellon University, School of Computer Science, Institute for Software Research. Technical report, CMU-ISR-12-112

Carley KM, Bigrigg MW, Diallo B (2012c) Data-to-model: a mixed initiative approach for rapid ethnographic assessment. Comput Math Organ Theory 18(3):300–327

Carley KM, Reminga J, Storrick J, Pfeffer J, Columbus D (2013a) ORA user's guide 2013, Carnegie Mellon University, School of Computer Science, Institute for Software Research. Technical report, CMU-ISR-13-108

Carley KM, Pfeffer J, Liu H, Morstatter F, Goolsby R (2013b) Near real time assessment of social media using geo-temporal network analytics. In: Proceedings of 2013 IEEE/ACM international conference on advances in social networks analysis and mining (ASONAM), Niagara Falls, 25–28 Aug 2013

Carley KM, Columbus D, Landwehr P (2013c) AutoMap user's guide 2013, Carnegie Mellon University, School of Computer Science, Institute for Software Research. Technical report, CMU-ISR-13-105

Effken JA, Gephart S, Carley KM (2013) Using ORA to assess the relationship of handoffs to quality and safety outcomes. CIN: Comput Inform Nurs 31(1): 36–44

Everton S (2012) Disrupting dark networks. Cambridge University Press, New York

Gullapalli A, Carley KM (2013) Extracting ordinal temporal trail clusters in networks using symbolic time series analysis. Soc Netw Anal Min, 1–16. Springer, Vienna

Kas M, Carley KM, Richard Carley L (2012) Who was where, when? Spatiotemporal analysis of researcher mobility in nuclear science. In: Proceedings of the international workshop on spatio temporal data integration and retrieval (STIR 2012), held in conjunction with ICDE 2012, Washington, DC, 1 Apr 2012

Kas M, Carley KM, Richard Carley L, (2013a) Incremental closeness centrality for dynamically changing social networks. In: Workshop on the semantic and dynamic analysis of information networks (SDAIN'13). Proceedings of the 2013 IEEE/ACM international conference on advances in social networks analysis and mining (ASONAM), Niagara Falls, 25–28 Aug 2013

Kas M, Wachs M, Carley KM, Richard Carley L (2013b) Incremental computation of betweenness centrality for dynamically growing networks. In: Proceedings of the 2013 IEEE/ACM international conference on advances in social networks analysis and mining (ASONAM'13), Niagara Falls, 25–28 Aug 2013

Kenney MJ, Horgan J, Horne C, Vining P, Carley KM, Bigrigg M, Bloom M, Braddock K (2012) Organizational adaptation in an activist network: social networks, leadership, and change in al-Muhajiroun. Appl Ergon 44(5):739–747

Merrill J, Orr MG, Jeon CY, Wilson RV, Storrick J, Carley KM (2012) Topology of local health officials' advice networks: mind the gaps. J Public Health Manag Pract 18(6):602–608

McCulloh IA, Johnson AN, Carley KM (2012) Spectral analysis of social networks to identify periodicity. J Math Sociol 36(2):80–96

McCulloh I, Armstrong H, Johnson A (2013) Social network analysis with applications. Wiley, Hoboken

Pfeffer J, Carley KM (2012) k-centralities: local approximations of global measures based on shortest paths. In: Proceedings of the WWW conference 2012, 1st international workshop on large scale network analysis (LSNA 2012), Lyon, pp 1043–1050

Van Holt T, Johnson JC, Brinkley J, Carley KM, Caspersen J (2012) Structure of ethnic violence in Sudan: an automated content, meta-network and geospatial analytical approach. Comput Math Organ Theory 18:340–355

Wei W, Pfeffer J, Reminga J, Carley KM (2011) Handling weighted, asymmetric, self-looped and disconnected networks in ORA, Carnegie Mellon University, School of Computer Science, Institute for Software Research. Technical report, CMU-ISR-11-113

Recommended Reading

Carley KM (2002) Smart agents and organizations of the future. In: Lievrouw L, Livingstone S (eds) The handbook of new media, chapter 12. Sage, Thousand Oaks, pp 206–220

Carley KM (2005) Organizational design and assessment in cyberspace. In: Rouse WB, Boff KR (eds) Organizational simulation. Wiley, Hoboken

Carley KM (2013) ORA: quick start guide, unpublished manuscript

Carley KM, Pfeffer J (2012) Dynamic network analysis (DNA) and ORA. In: Schmorrow DD, Nicholson DM (eds) Advances in design for cross-cultural activities part I. CRC, pp 265–274

Carley KM, Reminga J, Storrick J, Pfeffer J, Columbus D (2013a) ORA user's guide 2013, Carnegie Mellon University, School of Computer Science, Institute for Software Research. Technical report, CMU-ISR-13-108

Carley KM, Pfeffer J, Liu H, Morstatter F, Goolsby R (2013b) Near real time assessment of social media using geo-temporal network analytics. In: Proceedings of 2013 IEEE/ACM international conference on advances in social networks analysis and mining (ASONAM), Niagara Falls, 25–28 Aug 2013

ORA-NetScenes

▸ ORA: A Toolkit for Dynamic Network Analysis and Visualization

ORA Toolkit

▸ ORA: A Toolkit for Dynamic Network Analysis and Visualization

Organismic Computing

▸ Creating a Space for Collective Problem-Solving

Organization Sets

▸ Inter-organizational Networks

Organizational Blueprint

▸ Intra-Organizational Networks

Organizational Design

▸ Social Network Analysis in Organizational Structures Evaluation

Organizational Grapevine

▸ Intra-Organizational Networks

Organizational Marketing

▸ Business-to-Business Marketing

Organizational Network Analysis

▸ Social Network Analysis in Organizational Structures Evaluation

Organizational Networks

▸ Inter-organizational Networks

Organizational Online Grids

▸ Corporate Online Social Networks and Company Identity

Organizational Social Capital

▸ Intra-Organizational Networks

Origins of Social Network Analysis

David Knoke
Department of Sociology, University of
Minnesota, 909 Social Sciences, Minneapolis,
MN, USA

Synonyms

Interpersonal relations; Link analysis; Relational analysis; Social interaction; Social structural analysis

Glossary

Dyad Two actors and one or more relations among them

Field Observation The collection of information in natural social settings

Social Network A set of actors and the set of relations among them

Social Network Analysis Methods for collecting and analyzing data on actors and their relations

Sociometry Methods for measuring social relations in a small group

Sociogram A diagram of labeled points and lines representing a set of relations among individuals in a small group

Tertius Gaudens ("the third who benefits") A third party who affects a relation between two other parties

Triad three actors and one or more relations among them

Origins of Social Network Analysis

Social network analysis originated during the first half of the twentieth century in the disciplines of psychology, sociology, social psychology, and anthropology. Core concepts and principles of social structural relations were developed by a handful of scholars. The social psychological and anthropological lineages are discussed in separate entries in this volume ▸ Networks in Social Psychology, Beginning with Kurt Lewin by Patrick Doreian; ▸ Network Analysis in French Sociology and Anthropology by Karl van Meter. This entry discusses the psychological and sociological strands, drawing heavily from historical accounts by John Scott (1991) and Linton Freeman (1996, 2004). The following subsections describe the analytic ideas of Georg Simmel, early psychological contributions and sociometric diagrams of Jacob Moreno, and social structural research at Harvard University in the 1930s and 1940s.

Simmelian Triads

Although the German sociologist Georg Simmel (1858–1918) never used the term social network, many network analysts consider him the godfather of the field. Simmel's ideas about "sociation," the basic forms of interactions among persons, inspired subsequent theories and analyses of network microstructures. He argued that the fundamental social unit is the triad, not the dyad. The relationship between two people is intensified by "a third element, or by a social framework that transcends both members of the dyad" (Simmel 1950 [1908], p. 136), regardless of the content of those ties (e.g., friendship, business, kinship). For example, "a marriage with one child has a character which is completely different from that of a childless marriage, but is not significantly different from a marriage with two or more children" (p. 138). Whereas an isolated dyad favors greater individuality of both persons, because no majority can outvote an individual, adding a third member makes possible such a majority. As a result, strong individuality is devalued in a triad, and conformity to group norms is greater.

A third party may manipulate a dyad for personal gain or may mediate and ameliorate dyadic conflicts. In *The Web of Group Affiliations* (1955 [1922]), Simmel identified the *tertius gaudens* ("the third who benefits") as a third party who affects the relations between two other parties. The *gaudens* takes two basic forms: "either two parties are hostile toward one another and therefore compete for the favor of the third element; or they compete for the favor of the third element and therefore are hostile toward one another" (p. 155). An example of the first type is a consumer who plays two merchants against one another by bargaining with each for the lowest price and highest quality of merchandise in the market. This strategy requires the third party to keep the other two triad members unconnected, by controlling the information available to each. An example of a second, nonpartisan *gaudens* is a conciliator in labor management negotiations. "The non-partisan either produces the concord of two colliding parties, whereby he withdraws after making the effort of creating direct contact between the unconnected or quarreling elements; or he functions as an arbiter who balances, as it were, their contradictory claims against one another and eliminates what is incompatible in them" (p. 146–147). Simmel's ideas received empirical attention in research on small-group dynamics that emerged in the second half of the twentieth century.

Early Psychological Contributions

A few Canadian and American scholars of educational and developmental psychology in the 1920s conducted some of the earliest investigations of interpersonal relations among children, using concepts and research methods that later became prevalent in social network analysis (Freeman 1996). John Almack (1922) interviewed children, asking them to name others in their class whom they would invite to a party. Beth Wellman (1926) systematically recorded her observations of which sets of preschool children played together during free play periods. At the University of Toronto, Helen Bott (1928) also used ethnographic methods of data collection by observing preschool children at play. She identified five basic types of interaction: taking, interfering, watching, imitating, and cooperating. Her approach anticipated collecting information on multiple types of relational contents. Bott organized her data in matrix format, showing the tie between pairs of children in row-and-column entries, again anticipating a widespread format for data storage and computer analysis by later network researchers. (Helen Bott was the mother of Elizabeth Bott, a prominent member of the Manchester school of anthropologists who developed network analysis in that discipline.) Finally, Elizabeth Hagman (1933) observed children at free play, then interviewed them at the end of the school year about who were their playmates. In finding discrepancies between observational and interview data, she "defined a research problem that remains a key issue to a great many recent investigations in the social network field" (Freeman 2004, p. 21).

Sociometric Diagrams

The Austrian-American psychiatrist Jacob Moreno (1889–1974) moved to New York in 1925 where he developed group psychotherapy theory and psychodrama methodology. His seminal contribution to social network analysis appeared in *Who Shall Survive?* (Moreno 1934), which was "a signal event in the history of social network analysis … a turning point for the development of the field" (Freeman 2004, p. 7). Moreno used the term "network" in the modern sense, meeting three of four key criteria – structural ties linking social actors, systemic collection of empirical data, and graphic imagery – but he did not use mathematical models (Freeman 2004, p. 3). Moreno introduced the sociogram, a diagram (graph) of labeled points and lines representing a set of relations among individuals in a small social group. Relational data might be collected through multiple methods, including observations, experiments, interviews, or questionnaires. "For Moreno, social configurations had definite and discernible structures, and the mapping of these structures into a sociogram allowed a researcher to visualize the channels through which, for example, information could flow from one person to another and through which one individual could influence another" (Scott 1991, p. 10).

Constructing a sociogram enabled a therapist to identify group leaders and isolated persons and to reveal indirect connections, reciprocities, and asymmetric choices. For example, a sociometric "star" receives many friendship choices from others, reflecting her or his popularity and leadership in the group. Moreno argued that structural features revealed by sociograms of social configurations – formed through choices, attractions, repulsions, friendships, and other interpersonal interactions – could help to explain psychological well-being. Group relations served simultaneously as constraints and opportunities for social action and psychological development. Moreno applied sociometric methods to investigations of Sing Sing prison (Moreno 1932) and to a runaway problem at the Hudson School for Girls in New York (Moreno 1934). In the latter study, he demonstrated that, by mapping the girls' preferences for whom to sit with at meals and implementing changes based on the diagrams, runaways were decreased dramatically. Freeman (2004, pp. 35–36) contended that Moreno's psychology graduate student collaborator, Helen Jennings, made "immense" contributions to those projects, particularly in research design, data collection, and analysis.

In 1936, Moreno founded a journal, *Sociometry*, to which he recruited many prominent psychologists and sociologists as contributors and editorial board members. A Columbia University sociologist, Paul Lazarsfeld, calculated the baseline probabilities for a random sociometric choice model, which Moreno and Jennings (1938) published in the inaugural volume. This mathematical treatment satisfied the fourth defining criterion of social network analysis (Freeman 2004, p. 39). Despite the initial visibility and professional recognition accorded to Mareno's innovative contributions to network analysis, the sociometric model failed to take off after the 1930s. Freeman (2004, pp. 39–42) attributed this marginalization to Mareno's flawed character, a "dark side" of self-centered, self-serving megalomania which repulsed many of his early supporters. Consequently, social network analysis "continued to lack a unified structural paradigm" (p. 42).

Social Structural Research

In the 1930s, Harvard University became a center of social network research studies conducted by anthropologists and sociologists. A common focus of their projects was on empirically indentifying network subgroups, such as cliques and clusters, in natural social groups.

In 1924, engineers and personnel managers at the Western Electric Corporation began investigating the effects of physical conditions, such as lighting and heating, on worker productivity at its Hawthorne Works factory in Cicero, Illinois. Paradoxically, they discovered that any type of change would increased productivity. In 1927, the project head invited Elton Mayo (1880–1949), an Australian anthropologist and Harvard Business School professor, to create a new research team to refocus the Hawthorne studies on social psychological dimensions of the workplace. Mayo concluded that the earlier paradoxical findings were due to workers' increased motivation to perform because they were pleased that management took an interest in them. The Harvard team conducted intense observations of informal relations among 14 workers who assembled telephone switching equipment in a bank-wiring room. The men decreased their output, fearing the company would cut base pay rates and fire employees if the group became too productive. Cliques of bank wirers controlled rate busters who deviated from group norms with verbal ridicule, insult games, and punching on the upper arm ("binging"). The main report on the Western Electric project included many sociograms of the bank-wiring employees' informal relations, including friendship, antagonism, helping, and job-trading networks (Roethlisberger and Dickson 1939). The book was "the first major investigation to use sociograms to describe the actual relations observed in real situations" (Scott 1991, p. 18). However, Mareno's pioneering work on graphs was not discussed.

In 1929, W. Lloyd Warner (1898–1970), an American anthropologist who had conducted field work on an Australian tribe, returned to Harvard as an instructor in anthropology. Mayo recruited him as a half-time Business School faculty member and advisor to the Western Electric project. Between 1930 and 1935, Warner also led a group of graduate students in collecting ethnographic, historical, and interview data on interpersonal networks in "Yankee City" (Newburyport, Massachusetts), a small industrial city. The Yankee City study uncovered numerous cliques, defined as informal subgroups whose members develop group feeling, intimacy, and norms (Warner and Lunt 1941, p. 32). "Having discovered the existence of these cliques from the comments made by those they studied, Warner and his associates claimed that they were second in importance only to the family in placing people in society. People are integrated into communities through 'informal' and 'personal' relations of family and clique membership, not simply through the 'formal' relations of the economy and political system" (Scott 1991, p. 21). Because many people belong simultaneously to several cliques, "such overlapping in clique membership spreads out into a network of interrelations which integrate almost the entire population of a community in a single vast system of clique relations" (Warner and Lunt 1941, p. 111). Thus, network subgroups structure the larger social systems within which they are embedded. When fieldwork in Yankee City ended, Warner left Harvard for the University of Chicago. There he organized a team of ethnographers to collect data on racial differences in the social stratification of "Old City" (Natchez, Mississippi), a segregated Southern city about the size of Newburyport. The project team observed interactions among community members, identified 60 cliques, examined clique internal structures, and analyzed inter-clique connections (Davis et al. 1941).

George Homans (1910–1989) was exposed during the 1930s to structural perspectives as a Junior Fellow in the Harvard Society of Fellows, where he worked with Mayo and formed friendships with some of Warner's students. Subsequently, as a faculty member in the Harvard sociology department, Homans synthesized a theory of small group behavior based on the research findings of experimental social psychologists and observational sociologists and anthropologists. He developed a threefold classification that interconnects frequency of interaction, intensity of sentiment, and joint activity. In *The Human Group* Homans (1950) reanalyzed data from several classic network studies to demonstrate how his threefold classification scheme explained structural relations. Using data from the Old City project, he display a two-mode matrix of attendance 18 Southern white women at 14 social events ordered by date (see also Freeman 2003). By rearranging the matrix row and columns, Homans showed that the women were divided into two cliques whose members attended different subsets of parties. Although his method of matrix permutation appeared to be "simply a trial-and-error process which continued until he was able to spot an apparently significant pattern" (Scott 1991, p. 24), it anticipated more rigorous blockmodel methods developed two decades later at Harvard by Harrison White and his colleagues (*Networks at Harvard University Sociology* by Mark Pachucki). Homans turned his later attention to developing a social exchange theory deduced from behavioral psychology and economics. With retirements and departures, the early Harvard contributions to social network analysis came to an end in the 1940s, and no major network centers emerged in the USA during the next quarter century (Freeman 2004, pp. 61–64).

Acknowledgments

I thank Linton C. Freeman for his valuable counsel in developing the set of topics for the history of social network analysis appearing in the Encyclopedia. He was unable to write this entry, but his researches on the era remain an indispensable guide to the origins of the field.

Cross-References

- ▸ Networks at Harvard University Sociology
- ▸ Network Analysis in French Sociology and Anthropology
- ▸ Networks in Social Psychology, Beginning with Kurt Lewin

References

Almack JC (1922) The influence of intelligence on the selection of associates. School Soc 16:529–530
Bott HM (1928) Observations of play activities in a nursery school. Genet Psychol Monogr 4: 44–88
Davis A, Gardner BB, Gardner MR (1941) Deep south: a social anthropology of caste and class. University of Chicago Press, Chicago
Freeman LC (1996) Some antecedents of social network analysis. Connections 19:1–42
Freeman LC (2003) Finding social groups: a meta-analysis of the southern women data. In: Breiger R, Carely K, Pattison P (eds) Dynamic social network modeling and analysis. The National Academies Press, Washington, DC, pp 39–77
Freeman LC (2004) The development of social network analysis: a study in the sociology of science. Empirical, Vancouver
Hagman EP (1933) The companionships of preschool children. Univ Iowa Stud Child Welf 7: 10–69
Homans GC (1950) The human group. Harcourt/Brace, New York
Moreno JL (1932) Application of the group method to classification. National Committee on Prisons and Prison Labor, New York
Moreno JL (1934) Who shall survive? A new approach to the problem of human interrelations. Nervous and Mental Disease, Washington, DC
Moreno JL, Jennings HH (1938) Statistics of social configurations. Sociometry 1:342–374
Roethlisberger FJ, Dickson WJ (1939) Management and the worker. Harvard University Press, Cambridge
Scott J (1991) Social network analysis: a handbook. Sage, London
Simmel G (1950) [1908] The sociology of georg simmel. compiled and translated by Kurt Wolff. Free Press, Glencoe
Simmel G (1955) [1922] Conflict and the web of group affiliations. Kurt W, Reinhard B (Trans, eds) Free Press, Glencoe
Warner WL, Lunt PS (1941) The social life of a modern community. Yale University Press, New Haven
Wellman B (1926) The school child's choice of companions. J Educ Res 14:126–132

Outlier Detection

► Outlier Detection with Uncertain Data Using Graphics Processors

Outlier Detection with Uncertain Data Using Graphics Processors

Takazumi Matsumoto and Edward Hung
The Department of Computing, The Hong Kong Polytechnic University, Hung Hom, Hong Kong

Synonyms

Data mining; Graphics processors; Outlier detection; Parallel processing; Uncertain data

Glossary

Outlier Detection A data mining task in which data points that are outside expected patterns in a given dataset are identified

Central Processing Unit (CPU) A hardware component of a computer that is responsible for managing various operations as directed by a program

Graphics Processing Unit (GPU) A specialized hardware component of a computer that is designed to rapidly calculate large numbers of mathematical operations, primarily for displaying graphics. Modern GPUs can be programmed to perform a variety of other tasks

General Purpose Computing using GPUs (GPGPU) Programming GPUs for computational tasks other than graphics

Parallel Processing Computing multiple calculations simultaneously in order to more quickly solve a problem

Definition

Outlier detection is a widely employed data mining technique in which unusual events are extracted from a large body of data. This is used in areas such as network intrusion detection (Sequeria and Zaki 2002), fraud detection (Bolton and Hand 2002), and system fault detection (Lan et al. 2010) and is also often used in preprocessing large quantities of data before further analysis (e.g., by human experts). Recently there has been increased interest in mining uncertain data. A considerable portion of data collected in the real world contains some degree of uncertainty, as well as possibly erroneous and/or missing values.

One application in which outlier detection on uncertain data could be successfully applied is in the area of detecting fraudulent or erroneous activity on social networks. Since social network data contains large amounts of sensitive personal information, uncertainty can be added to preserve user privacy, while detecting anomalous or fraudulent activity. With increasing concern over the mining of personally identifiable information, even when obvious identifiers such as names and IP addresses have been stripped, uncertainty may be deliberately added to preserve privacy. A typical outlier detection method would discard this uncertainty information and treat the data as if it were certain; however, it has been noted (Aggarwal 2009) that incorporating uncertainty can be beneficial for outlier detection.

Modeling uncertainty adds a significant amount of complexity to outlier detection as uncertain objects are no longer represented by a single point, but rather a probabilistic object. However, fast multi-core CPUs and programmable GPUs have become popular in recent years, which provide substantial parallel computing resources at a very low cost. In particular, modern GPUs are massively parallel floating point processors attached to dedicated high-speed memory, providing large parallel computation capability for a fraction of the cost of traditional highly parallel processing computers.

Programming frameworks such as CUDA and OpenCL now allow programs to take advantage of this previously underutilized parallel processing potential in the form of GPGPU, accelerating computationally intensive tasks beyond typical applications in 3D graphics in scientific, professional, and personal applications. This performance can be leveraged when attempting to detect outliers with uncertain data in time-sensitive situations such as responding to sensor failure or network intrusion.

Key Points

In this entry, the performance and quality of the distance-based clustering/density-based outlier detection algorithm, originally described in Aggarwal and Yu (2008), is tested with serial and parallel implementations, using both the CPU only and the GPU.

This work also covers the following:

- Formal formulation of the problem and the algorithm in detail
- Notes on parallel implementation and the peculiarities of the OpenCL framework, as well as consideration of the representation and transformation of data from certain to uncertain forms
- Experimental methodology and the results of performance and quality tests
- Conclusion and consideration on future applications

Historical Background

Outlier detection is a well established and commonly used technique for detecting data points that lie outside of expected patterns. The two key factors under consideration are detection quality, which can be measured with metrics such as precision and recall (Aggarwal and Yu 2008; Tang et al. 2006) or sensitivity and specificity (Angiulli et al. 2006; Lan et al. 2010), and performance, typically observed with running time comparisons between implementations (as in Alshawabkeh et al. 2010). Outlier detection can be performance sensitive as it is often used as a preprocessing step before further analysis or as a service providing real-time alerts (e.g., intrusion detection system). There is generally a trade-off between performance and quality, depending on the various algorithms and implementations of outlier detection. This section will give an overview of some of the related work in the area of outlier detection and the use of GPGPU to accelerate outlier detection.

The prototypical approach to outlier detection is as a by-product of clustering algorithms (Alshawabkeh et al. 2010). In a clustering context, an algorithm such as DBSCAN (Ester et al. 1996) will exclude data points that are not sufficiently similar to other objects. This is usually done to improve cluster purity and prune otherwise spurious data points from consideration.

Alternatively, outlier detection can be described as a machine learning problem, so classification techniques such as k-nearest neighbors (k-NN), Support Vector Machines, and neural networks can be used to learn the distinction between an outlier and a non-outlier (also referred to as an in-lier). An outlier detection problem in this context can be considered as an unbalanced binary class problem (Nguyen and Gopalkrishnan 2010), since it is expected that the number of outliers is relatively small in proportion to non-outliers. In this entry, we focus on unsupervised techniques, since there may be no prior knowledge about the data or its classes (Heymann et al. 2012). The unbalanced nature of the classes create a challenge to classification systems (Nguyen and Gopalkrishnan 2010), while it is quite possible for unsupervised systems to observe the expected patterns from the data (Lan et al. 2010).

Outlier detection techniques can be broadly categorized into two main methodologies: *distance-based* approaches (Knorr and Ng 1998, 1999; Ramaswamy et al. 2000) and *density-based* approaches such as *Local Outlier Factor* (LOF) (Breunig et al. 2000).

Two distance-based algorithms were proposed in Knorr and Ng (1999): the Nested loop algorithm and the Cell-based algorithm. However, the former scales poorly with respect to dataset size and the latter scales poorly with respect to dataset dimensionality (Angiulli et al. 2006). A variant of this approach was proposed in Ramaswamy et al. (2000), which uses the concept of k-distance (the distance to p's k-nearest neighbors) to assert the objects with the highest k-distances as outliers. Ramaswamy et al. (2000) also uses partitioning clustering of the dataset and performs pruning to ignore partitions that do not contain outliers. On the other hand, LOF uses a similar concept to DBSCAN by using reachability in a local neighborhood around each point. Ultimately, the performance of either distance or density-based outlier detection schemes will depend on the characteristics of the input data (Tang et al. 2006) as well as the definition of outlier used.

In this entry, we focus on the outlier detection scheme proposed in Aggarwal and Yu (2008), which describes outlier detection on uncertain data database records, with each record having a number of attributes (dimensions). Each dimension has a *pdf*, and the objective is to find data points in an area with data density η (expressed as the η-probability of a data point) less than a threshold value δ, i.e., a (δ, η)-outlier. This algorithm is described in more detail later in this entry.

Since it is impractical to determine η-probabilities directly, a sampling approach of the *pdf*s is used. However, as sampling global density is a computationally expensive operation (quadratic scaling with respect to dataset size), a *microclustering* technique (Aggarwal and Yu 2008) is used to reduce the number of data objects by creating clusters of data objects and using the cluster centroids rather than the raw data to reduce the number of objects processed, similar to the approach taken in Ramaswamy et al. (2000). The initial cluster centroids are generated randomly, each with each data point being assigned to the nearest cluster. This is opposed to the common approach of dimensionality reduction by feature selection in classification approaches, and results in running time depends on the time taken for clustering and testing each data object (both constant time).

Data mining applications such as outlier detection are good candidates for parallelization (Hung and Cheung 2002) as in typical cases there is a large amount of data that is processed by a small number of routines. These tasks are "data parallel" and are well suited for execution on a GPU (Alshawabkeh et al. 2010). Unlike conventional parallel processing computers that have many complex CPU cores, a modern GPU consists of a large number of simple "stream processors" that are individually capable of only a few operations. However, the ability to pack many stream processors into the same space as a single CPU core gives GPUs a large advantage in parallelism. A traditional parallel computer such as the Intel Itanium 2 system used in Lozano and Acuna (2005) is significantly more costly.

The two most popular programming frameworks for GPGPU are C for CUDA, a proprietary solution developed by NVIDIA Corporation (2011), and OpenCL, an open standard managed by The Khronos Group (2011) and backed by multiple companies including Intel, AMD, NVIDIA, and Apple. With both CUDA and OpenCL, work is split from the host (i.e., the CPU) to kernels that execute on a computing device (usually GPUs). Kernels contain the computationally intensive tasks, while the host is tasked with managing the other computing devices. A single kernel can be executed in parallel by many worker threads on a GPU.

Several outlier detection algorithms including LOF (Alshawabkeh et al. 2010) have been adapted for acceleration with GPUs using CUDA (Tarabalka et al. 2009; Bastke et al. 2009; Huhle et al. 2008) and have seen significant reductions in running times (e.g., a 100 fold improvement in Alshawabkeh et al. 2010). In this work, we opted to use OpenCL as it provides a high degree of portability between different manufacturers of GPUs, as well as the ability to execute the same parallel code on a CPU for comparison.

While there is a significant body of work in the area of accelerating outlier detection on regular (certain) data, it is often the case that data has some degree of uncertainty or error (Aggarwal 2009). Moreover, some statistical techniques such as forecasting and privacy-preserving data mining will by nature be uncertain. The naïve approach would be to disregard the probability distributions in the uncertain dataset and take the mean value only as the data point, thus avoiding the problem of creating a model for the uncertainty. However, by applying common probability density functions (*pdf*s) such as the Gaussian distribution, a convenient closed form representation for storage and sampling points is available.

Algorithm for Outlier Detection with Uncertain Data

Problem Formulation

There are numerous definitions of outliers that have been proposed, the most well known of which are the distance-based definition proposed in Knorr and Ng (1998), which itself is based on an earlier definition in Hawkins (1980), and the local density-based definition used by LOF (Breunig et al. 2000). The formal definition of uncertain outliers as defined in Aggarwal and Yu (2008) follows in section "Density-Based Outlier Detection."

The intuitive definition of outliers is described in Hawkins (1980) as "*an observation that deviates so much from other observations as to arouse suspicion that it was generated by a different mechanism*." In addition, the distribution between outliers and non-outliers should be unbalanced, that is, there are significantly more outliers compared to non-ouliters (Fig. 1).

Let dataset $\mathcal{D}$ contain n uncertain data objects, with each object d_i having r dimensions. For each object, each dimension has a *pdf*, for a total of nr *pdf*s. $\mathcal{D}$ is normalized to the range [0,1].

The *pdf* for object d_i along dimension j is denoted by $h_i^j(\cdot)$ and the standard deviation of $h_i^j(\cdot)$ is denoted by $\psi_j(d_i)$. Since data is typically stored in certain form, uncertainty is estimated from the calculated standard deviation of each dimension.

Density-Based Outlier Detection

It is defined that the η-probability of object d_i is the probability that d_i lies in a subspace with overall data density of at least η. A subspace is defined as a subset of the full r dimensions, while overall data density is defined by the summation of the *pdf*s of each object. The probability p_i of d_i in a subspace of dimensionality a with overall data density of at least η can be found by solving the following integral:

$$p_i = \int_{G(x_1,\ldots,x_a)\geq\eta} \prod_{j=1}^{a} h_i^j(x_j)\,dx_j \qquad (1)$$

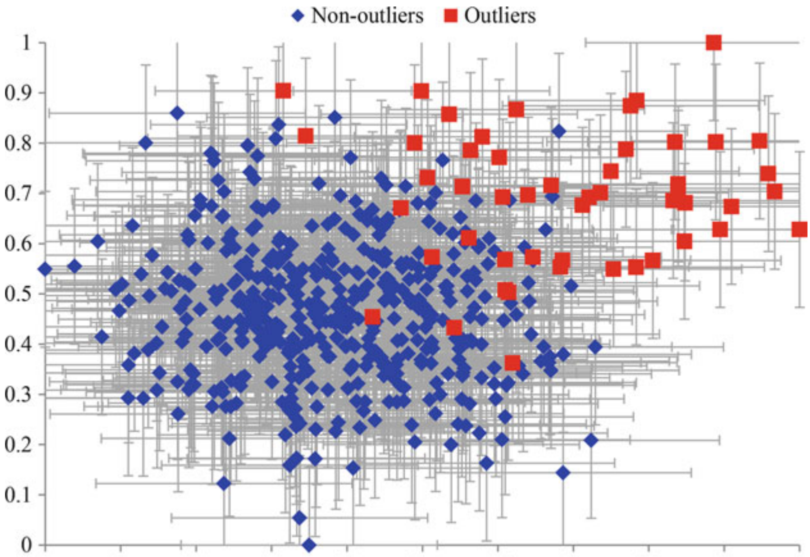

Outlier Detection with Uncertain Data Using Graphics Processors, Fig. 1 A sample synthetic dataset with 500 objects (450 normal objects, 50 outliers) and 2 dimensions. Error bars represent standard deviation

Note that $h(x_1, \ldots, x_a)$ represents the overall probability density function on all coordinates in the given a-dimensional subspace. However, as p_i is difficult to calculate precisely as in (1), it can be estimated using a density sampling algorithm. As it is necessary to estimate the overall density of the data space at any arbitrary point p, the common Gaussian distribution function, given by $f(p) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(p-\mu)^2}{2\sigma^2}}$, is used. Thus the density function $\overline{f^Q}$ given a uniformly random value u in $[0, 1]$ and uncertainty values $\psi(\cdot)$ is defined as follows:

$$\overline{f^Q}(u, \psi(d_i)) = \frac{1}{n} \sum_{i=1}^{n} \frac{1}{\sqrt{2\pi}(w + \psi(d_i))} \times e^{-\frac{(u-\overline{X_i})^2}{2(w^2+\psi(d_i)^2)}} \tag{2}$$

The parameter w is a smoothing factor, which was determined using the Silverman approximation rule as $1.06N^{-0.2}\sigma$ (Aggarwal and Yu 2008). The sampling algorithm is given in Fig. 2 in pseudo-code.

s denotes the number of samples, set as 800 by default as in Aggarwal and Yu (2008). a and a_{start} represent the last and first dimension of the current subspace, respectively. The inverse cumulative distribution function (inverse cdf) is computed numerically (Acklam 2003).

By sampling an object at multiple random points, calculating the overall data density at each sampled point and counting how many of those sampled points exceed η, the probability that object lies in a subspace with data density of at least η (i.e., η-probability) can be estimated. Finally, object d_i is defined as a (δ, η)-outlier if the η-probability of d_i in any subspace is less than a user parameter δ, that is, the object has a low probability of lying in a region of high data density. The overall algorithm in pseudo-code is given in Fig. 3.

The outlier detection algorithm as described in Aggarwal and Yu (2008) uses a "roll-up" method to explore the data space. This is an iterative approach that starts with a one-dimensional subspace and adds another dimension for each iteration until all dimensions are added. Each subspace is tested for outliers and any objects determined to be outliers are removed from further consideration in other subspaces. It is assumed that outliers will be lower density in some

Outlier Detection with Uncertain Data Using Graphics Processors, Fig. 2 Implementation of *EstimateProbability*

```
EstimateProbability($\mathcal{D}$, $d_i$, $\eta$, $a$, $s$, $a_{start}$)
Let $F_i^j(\cdot)$ be the inverse cumulative distribution function of pdf $h_i^j(\cdot)$
$success = 0$, $runs = 0$
for $s$ times do
    $density_i = 0$
    for $j = a_{start}$ to $a$ do
        $y$ = a uniform random value $[0, 1]$
        for all $d_k$ in $\mathcal{D}$ do
            $density_i = density_i + h_k^j(F_i^j(y))$
        end for
        $density_i = density_i / |\mathcal{D}|$
    end for
    if $density_i > \eta$ then
        $success = success + 1$
    end if
    $runs = runs + 1$
end for
return $(success/runs)$
```

subspaces than others, and conversely a non-outlier in this definition has high density in all tested subspaces.

As a baseline, a simple global averaging method which uses the average data density over all dimensions is also demonstrated. This method only requires a single pass over the entire data space, allowing a fast linear scaling with input dataset size. Practically, this is achieved by setting $a = r$ instead of equal to the first dimension, which is equivalent to the final iteration of the roll-up method. This method is significantly faster compared to the roll-up method, but since only outliers that are in an area of low density in the entire data space can be found, the outlier detection ability is also significantly reduced.

Distance-Based Cluster Compression

It is evident that the requirement to calculate overall data density at each sampled point presents a large computational bottleneck. The overall complexity of the outlier detection step is high, with running time scaling up significantly as the number of data objects increases. To alleviate this, the number of objects n can be reduced (as in Ramaswamy et al. (2000)) using "microclustering" (Aggarwal and Yu 2008), effectively compressing many uncertain objects into a single representative uncertain object. This process (also referred to simply as clustering) is as follows:

Each cluster $\mathcal{C}$ is defined by a $3r + 1$ tuple $(\overline{CF2^x}(\mathcal{C}), \overline{EF2^x}(\mathcal{C}), \overline{CF1^x}(\mathcal{C}), n(\mathcal{C}))$.

- $\overline{CF2^x}(\mathcal{C})$ is a r-dimensional vector, with each entry p containing the sum of squares of mean values of the member objects' p-th dimension.
- $\overline{EF2^x}(\mathcal{C})$ is a r-dimensional vector, with each entry containing the sum of squares of errors of the member objects' p-th dimension.
- $\overline{CF1^x}(\mathcal{C})$ is a r-dimensional vector, with each entry p containing the sum of mean values of the member objects' p-th dimension.
- $n(\mathcal{C})$ contains the number of data objects in the cluster.

Outlier Detection with Uncertain Data Using Graphics Processors, Fig. 3 Implementation of *DetectOutlier*

```
DetectOutlier($\mathcal{D}$, $\eta$, $\delta$, $a$, $r$, $s$, $b$)
$\mathcal{O} = null$
while $|\mathcal{O}| < |\mathcal{D}|$ and $a \leq r$ do
    if $b > 0$ then
        $a_{start} = a - b$ or 0, whichever is larger
    end if
    $\mathcal{D}_a = \{$Subspace of $\mathcal{D}$ starting at dimension $a_{start}$ up to $a\} - \mathcal{O}$
    for Each object $d_i$ in $\mathcal{D}_a$ do
        p = EstimateProbability($\mathcal{D}$, $d_i$, $\eta$, $a$, $s$, $a_{start}$)
        if $p < \delta$ then
            Add $d_i$ to $\mathcal{O}$
        end if
    end for
    $a = a + 1$
end while
return $\mathcal{O}$
```

Initially, q centroids are generated by using a uniform random number generator in the range [0, 1], so each empty cluster is given an initial value. No new clusters are created, centroids remain fixed, and all data objects are assigned to the nearest centroid by Euclidean distance between the centroid's value and the mean value of the incoming data object. Since some clusters may contain zero members after clustering, these clusters are removed before outlier detection (Fig. 4).

As each cluster $\mathcal{C}$ is treated as a data object, a true error Δ is calculated by averaging the members, with bias equal to the distance from the cluster centroid and variance equal to the original error. This is defined for a given dimension j as

$$\Delta_j(\mathcal{C}) = \sum_{i=1}^{r} \frac{\text{bias}_j(\overline{X}, \mathcal{C})^2 + \psi_j(\overline{X})^2}{r} \tag{3}$$

Equation 3 can be rearranged to use the stored cluster tuple as follows:

$$\Delta_j(\mathcal{C}) = \frac{CF2_j^x(\mathcal{C})}{r} - \frac{CF1_j^x(\mathcal{C})^2}{r^2} + \frac{EF2_j^x(\mathcal{C})}{r} \tag{4}$$

The advantage of this approach is that it does not require further modification of the algorithm, as $\Delta(\mathcal{C})$ and $n(\mathcal{C})$ can be substituted into (2) as follows:

$$\overline{f^Q}(p, \Delta(\mathcal{C})) = \frac{1}{n} \sum_{i=1}^{q} n(\mathcal{C}_i) \frac{1}{\sqrt{2\pi}(h + \Delta(\mathcal{C}_i))} \times e^{-\frac{(p - \overline{X_i})^2}{2(h^2 + \Delta(\mathcal{C}_i))^2}} \tag{5}$$

Implementation Notes

Since the vast majority of data stored at present is in certain form, it is often required to convert it to uncertain form. This process of adding uncertainty can also used for privacy preservation. We only consider numerical values, with categorical values such as name and address removed. Further, numerical values (e.g., number of friends in

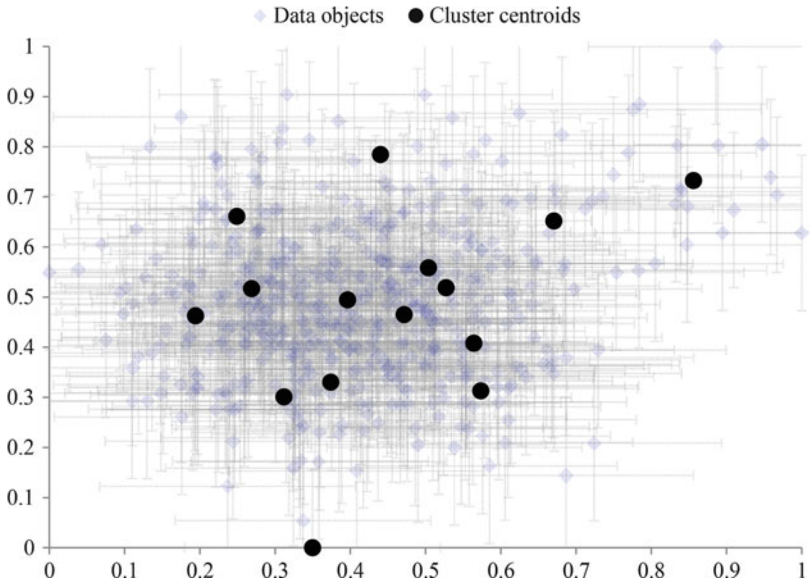

Outlier Detection with Uncertain Data Using Graphics Processors, Fig. 4 Fifteen populated cluster centroids ($q = 20$) overlaid on the 500 objects from Fig. 1

a social circle) are made uncertain by adding a random value and attaching a standard deviation.

The key parameters to consider are η and δ, which controls the algorithm's sensitivity to outliers, and the amount of uncertainty present. Of slightly less importance is the number of samples s and the number of clusters q if clustering is used. The selection of appropriate parameters is a challenge in any algorithm; in this entry, we consider the effect of the level of uncertainty modeled and the average difference between certain and uncertain data objects.

In order to measure the performance change from parallelization, we implemented the algorithm described previously in two ways: a baseline serial implementation in C++ (for the CPU) and a parallel implementation leveraging the OpenCL framework (for parallel execution on the CPU and GPU). For comparison purposes, the two implementations mirror each other as closely as possible.

Parallel Methods

When a kernel (subroutine) executes on a computing device such as a GPU, it should take into account the differences in architecture of a GPU compared to a standard CPU. To better leverage the GPU using OpenCL, this implementation uses several optimizations such as the use of single precision floating point values and special hardware-accelerated mathematical functions (*native* functions). Single precision floating point values are used extensively in graphics, and thus GPUs are optimized for many single precision functions. Typically, double precision floating point (FP64) calculations are slower by a significant margin compared to single precision (FP32), with FP64 usually noted as anywhere between 1/4 and 1/16 FP32 speed. Since earlier preliminary work has shown little quality impact from using FP64, FP32 has been used exclusively (both on the CPU and GPU).

In addition to an OpenCL implementation of *DetectOutlier*, we have implemented a number of other functions including the uniform random number generator (RNG) and the calculation of *pdf* and *cdf*. For portability we used the simple, predictable RNG *xorshift* (Marsaglia 2003). RNGs can be a challenge to implement correctly in a parallel environment, and for testing we required identical number generation between both host and GPU implementations that could be repeated consistently. A more robust RNG may be substituted if required.

The current OpenCL framework (version 1.2 at time of writing) and the GPU used impose certain additional restrictions, including a lack of recursion (i.e., a function calling itself) and lack of dynamic memory allocation on the GPU (i.e., all memory must be allocated when the kernel is first initialized). This is a problem as refinement methods used in math libraries often use recursion, requiring changes to remove recursion and to take advantage of additional OpenCL functionality (e.g., the complementary error function *erfc*). Since precision is limited to FP32, refinement of results to maximum precision is also simplified.

Other considerations include branching logic, which can cause a reduction in performance as GPUs must execute the same code path for each worker thread executing in parallel. Memory management is a more complex, multi-level affair on a GPU, with the fastest private memory available for each worker thread used as a scratch area and a slower global memory space that can be accessed by all workers used to store the large dataset $\mathcal{D}$. As copying data to and from the GPU is an expensive operation, data transfers should be minimized and processing done on the GPU maximized.

For any parallel computation task, the key concern is balancing the number of parallel tasks with the overhead of managing those workers. It is desirable to have as many threads as possible while having each thread do sufficient work and without overwhelming the processors. In this implementation, each data object is assigned one worker thread, and each worker is responsible for density sampling of that object's space. This allows for the unrolling of the most computationally expensive loop in *DetectOutlier*, as shown in Fig. 5.

```
ParallelDetectOutlier($\mathcal{D}$, $\eta$,$\delta$, $a$, $r$, $s$, $b$, $\mathcal{O}$)
i = worker thread number
while $d_i$ is not an outlier and $a \leq r$ do
    if $b > 0$ then
        $a_{start} = a - b$ or 0,whichever is larger
    end if
    p = EstimateProbability($\mathcal{D}$, $d_i$, $\eta$, $a$, $s$, $a_{start}$)
    if $p < \delta$ then
        Add $d_i$ to $\mathcal{O}$ as an outlier
    end if
    $a = a+1$
end while
```

Outlier Detection with Uncertain Data Using Graphics Processors, Fig. 5 Parallel implementation of *DetectOutlier*

Experimental Evaluation

Methodology

To test performance and quality, synthetic datasets of varying sizes were generated. The model used for generating the synthetic data is as follows – random numbers were generated following a Gaussian distribution with a mean value of 0 and a standard deviation of 1. The outliers are generated similarly, but an offset value is added to the mean (e.g., in Fig. 1, the offset value of the uncertain objects is 1.0), then the dataset is normalized. The number of outliers generated is 10 % of the total number of data objects.

Experimental testing was conducted on a PC running Microsoft Windows Vista SP2 with an Intel Core 2 Duo E8200 dual core CPU and an NVIDIA GeForce GT440 (96 stream processors) GPU. The serial and host code was compiled using Microsoft Visual Studio 2010. The OpenCL code was run using NVIDIA CUDA Toolkit 4.2 and driver 301.32 for the GPU and AMD APP 2.6 for the CPU.

Performance

In all tests, we used $\eta = \delta = 0.5$, $q = 150$, $s = 800$, $n = 100{,}000$ and $r = 5$ unless stated otherwise.

It is evident from Fig. 6 that all methods scale similarly with the number of objects, with the roll-up method on the GPU approximately 2× to 9× faster than on the CPU.

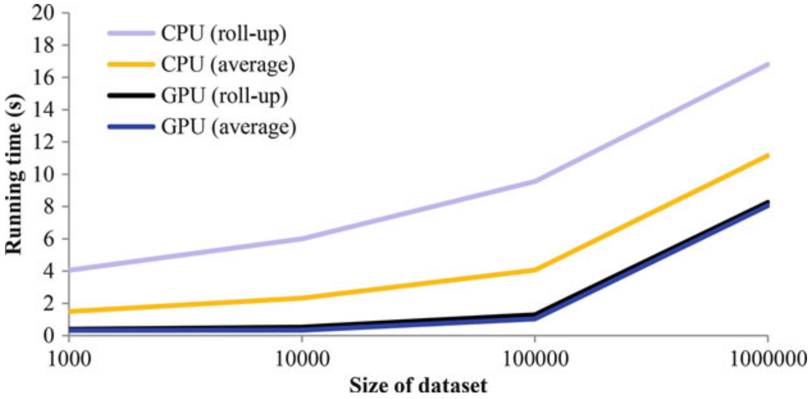

Outlier Detection with Uncertain Data Using Graphics Processors, Fig. 6 A comparison of running times with increasing numbers of data objects

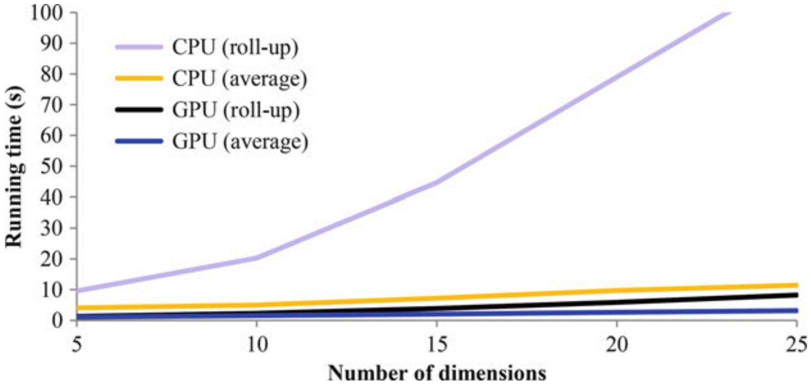

Outlier Detection with Uncertain Data Using Graphics Processors, Fig. 7 A comparison of running times with increasing dimensionality of data objects

The average method on the CPU is closer to the GPU speeds, but with significantly degraded quality. Note that as dataset size increases, the amount of time spent on clustering scales up rapidly while the time spent density sampling remains fairly constant.

Figure 7 demonstrates the processing power advantage of the GPU, with the time spent on clustering staying fairly constant while the time spent on density sampling (using the roll-up method) increases significantly. The CPU-based roll-up method rapidly scales up in running time while the GPU roll-up method is slightly faster with the much faster but much less accurate average method. Clearly for rapid detection of outliers with many dimensions, GPU acceleration offers high potential.

Quality

To test detection quality, precision and recall are measured using different uncertainties, from 1.0 to 4.0. This adjusts how much uncertainty (i.e., standard deviation) is assumed to be in the data. The offset value indicates on average how much difference there is between an outlier and a nonoutlier object. The lower the offset, the more mixed the outliers and non-outliers are (Fig. 8).

In this test, optimal balance between precision and recall was achieved with an uncertainty value of 3.0 in all cases. Clearly the best results are achieved when there is a large difference between outlier and non-outlier objects, but even when the objects are mixed together closely, the detection ability can be adjusted by changing the uncertainty value.

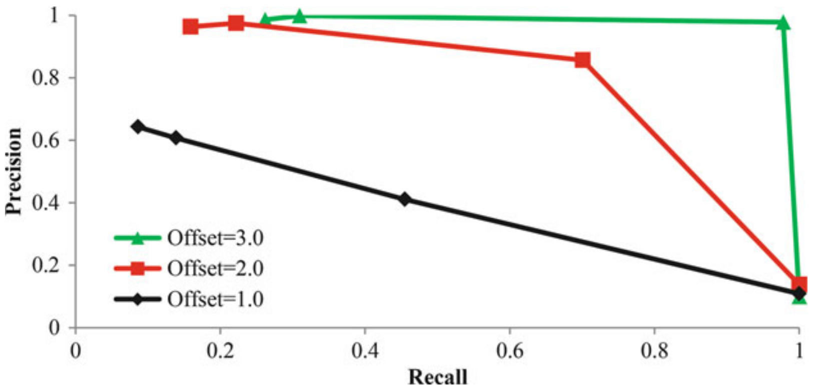

Outlier Detection with Uncertain Data Using Graphics Processors, Fig. 8 F1 score with different numbers of dimensions

Application

In this entry, we have demonstrated the use of a density sampling algorithm for outlier detection on uncertain data together with parallelization and optimization for the OpenCL framework. Experimental work on both a multi-core CPU and low cost GPU demonstrates significant reductions to running time, which could enable very fast detection of outliers in large uncertain datasets, such as social networking data that has been anonymized for privacy preservation.

Future Directions

In future work we would like to explore methods that can provide better detection quality under more challenging conditions where the outlier and non-outlier data objects are mixed closer together (i.e., more sensitivity). Avenues for more efficiency in the algorithm and implementation are also to be explored. Moreover, this algorithm does not currently take into account relationships that are present in datasets, which can be used to detect more subtle events. Finally, more extensive testing will be undertaken with a variety of datasets and sources.

Acknowledgments

The work described in this entry was partially supported by grants from the Research Grants Council of the Hong Kong Special Administrative Region, China (PolyU 5191/09E, PolyU 5182/08E, PolyU 5166/11E), and Hong Kong PhD Fellowship.

Cross-References

- ▶ Data Mining
- ▶ Distributed Processing of Networked Data
- ▶ Ethics of Social Networks and Mining
- ▶ Fraud Detection Using Social Network Analysis, a Case Study
- ▶ Mining Blackhole and Volcano Patterns for Fraud Detection
- ▶ Network Anomaly Detection Using Co-Clustering

References

Acklam PJ (2003) An algorithm for computing the inverse normal cumulative distribution function. Technical report

Aggarwal CC (ed) (2009) Managing and mining uncertain data. Springer, New York

Aggarwal CC, Yu PS (2008) Outlier detection with uncertain data. In: Proceedings of the SIAM international conference on data mining, Atlanta, 2008

Alshawabkeh M, Jang B, Kaeli D (2010) Accelerating the local outlier factor algorithm on a GPU for intrusion detection systems. In: Proceedings of the 3rd workshop on general-purpose computation on graphics processing units, Pittsburgh

Angiulli F, Basta S, Pizzuti C (2006) Distance-based detection of outliers. IEEE Trans Knowl Data Eng 18(2):145–160

Bastke S, Deml M, Schmidt S (2009) Combining statistical network data, probabilistic neural networks and the computational power of gpus for anomaly detection in computer networks. In: 1st workshop on intelligent security (security and artificial intelligence), Thessaloniki

Bolton RJ, Hand DJ (2002) Statistical fraud detection: a review. Stat Sci 17(3):235–255

Breunig MM, Kriegel HP, Ng RT, Sander J (2000) LOF: identifying density-based local outliers. In: Proceedings of SIGMOD 2000, Dallas

Ester M, Kriegel HP, Sander J, Xu X (1996) A density-based algorithm for discovering clusters in large spatial databases with noise. In: Proceedings of the 2nd international conference on knowledge discovery and data mining, Portland

Hawkins DM (1980) Identification of outliers. Chapman and Hall, London

Heymann S, Latapy M, Magnien C (2012) Outskewer: using skewness to spot outliers in samples and time series. In: 2012 IEEE/ACM international conference on advances in social networks analysis and mining, Istanbul

Huhle B, Schairer T, Jenke P, Strasser W (2008) Robust non-local denoising of colored depth data. In: IEEE computer society conference on computer vision and pattern recognition, workshop on time of flight camera based computer vision, Anchorage

Hung E, Cheung DW (2002) Parallel mining of outliers in large database. Distrib Parallel Databases 12(1):5–26

Khronos Group (2011) OpenCL. http://www.khronos.org/opencl/. Accessed 9 Oct 2012

Knorr EM, Ng RT (1998) Algorithms for mining distance-based outliers in large datasets. In: Proceedings of VLDB 1998, New York

Knorr EM, Ng RT (1999) Finding intensional knowledge of distance-based outliers. In: Proceedings of VLDB 1999, Edinburgh, pp 211–222

Lan Z, Zheng Z, Li Y (2010) Toward automated anomaly identification in large-scale systems. IEEE Trans Parallel Distrib Syst 21(2):174–187

Lozano E, Acuna E (2005) Parallel algorithms for distance-based and density-based outliers. In: Proceedings of the 5th IEEE international conference on data mining, Houston

Marsaglia G (2003) Xorshift RNGs. J Stat Softw 8(14): 1–6

Nguyen HV, Gopalkrishnan V (2010) Feature extraction for outlier detection in high-dimensional spaces. J Mach Learn Res Proc Track 10:66–75

NVIDIA Corporation (2011) CUDA. http://www.nvidia.com/object/cuda_home_new.html. Accessed 9 Oct 2012

Ramaswamy S, Rastogi R, Shim K (2000) Efficient algorithms for mining outliers from large data sets. In: Proceedings of the 2000 ACM SIGMOD, Dallas

Sequeria K, Zaki M (2002) ADMIT: anomaly-based data mining for intrusions. In: Proceedings of the 8th ACM SIGKDD, Madison

Tang J, Chen Z, Fu AW, Cheung DW (2006) Capabilities of outlier detection schemes in large datasets, framework and methodologies. Knowl Inf Syst 11(1): 45–84

Tarabalka Y, Haavardsholm TV, Kaasen I, Skauli T (2009) Real-time anomaly detection in hyperspectral images using multivariate normal mixture models and GPU processing. J Real-Time Image Process 4(3): 287–300

Over-Justification

▶ Incentives in Collaborative Applications

Ownership of Data

▶ Ethics of Social Networks and Mining