

# Exploring the Relationships among ICTs: A Scalable Computational Approach Using KL Divergence and Hierarchical Clustering

Chia-jung Tsui, Ping Wang, Kenneth R. Fleischmann, Douglas W. Oard, and Asad B. Sayeed  
 University of Maryland, College Park  
 [{ctsui,pwang,kfleisch,oard,asayeed}@umd.edu](mailto:{ctsui,pwang,kfleisch,oard,asayeed}@umd.edu)

## Abstract

*Different information and communication technologies (ICTs) are related in complex ways and, accordingly, their diffusion trajectories are related, too. How can the relationships among multiple ICTs be described and analyzed in a scalable way? In this study, we offer a scalable methodology, based on computational analysis of discourse, to examine the relationships among ICTs. Specifically, we employed Kullback-Leibler (KL) divergence to compare the semantic similarity of forty-seven ICTs discussed in the trade magazine InformationWeek over a decade. Using hierarchical clustering, we have found that the similarity of the technologies can be mapped in a hierarchy and similar technologies demonstrated similar discourses. The results establish the validity of our approach and demonstrate its scalability and richness. This analytical approach not only enables diffusion researchers to undertake multi-innovation, multi-source, and multi-period studies, but also helps practitioners effectively adopt and efficiently use new ICTs in their organizations.*

## 1. Introduction

Practitioners who consider adopting and using new information and communications technologies (ICTs) and scholars who study the diffusion of ICTs face a constant challenge: At any one time, we confront numerous seemingly promising ICT innovations. Some of them become widely adopted and used, making significant contributions to economic prosperity and social welfare; whereas others fade away, leaving little trace behind. While it has been argued that various ICTs are related to varying degrees and so are their diffusion trajectories [19], it is difficult to make sense of the relationships among ICT innovations.

For example, here is a partial list of contemporary ICT innovations: Service-oriented architecture (SOA), web services, open source software (OSS), web 2.0, YouTube, iPhone, blogs, and utility computing. How

are they related? How are their diffusion trajectories related? These are difficult questions, not only because of the complex relationships among ICT innovations, but also due to the lack of reliable methods to describe and analyze the relationships.

Indeed, ICT innovations are related in complex ways. First, a broader concept may be comprised of narrower, more specific concepts. Second, different concepts may represent the same core idea. Third, concepts may compete with each other as alternative solutions to similar problems or for the attention from the same group of people or organizations. Finally, concepts may complement each other to accomplish common tasks. Over time, these relationships may change, making it even harder to interpret.

Researchers of ICT diffusion are not well equipped to describe and analyze the complex and evolving relationships among ICT innovations. On the one hand, many studies in the dominant paradigm of ICT diffusion research have demonstrated that various organizational, technical, and environmental factors influence ICT adoption and use [8]. As this dominant paradigm is reaching "the point of diminishing returns as a framework for supporting ground-breaking research" [8, p. 314], it should be noted that most studies in the dominant paradigm employ single-innovation research designs, leaving the relationships among ICTs underexplored. On the other hand, the few multi-innovation studies have had to explicitly or implicitly rely on domain experts to evaluate ICT relationships (e.g., [6], [19]). Such expert evaluations are difficult to replicate, to generalize to other ICTs, or to *scale up* to examine the relationships among a large number of ICTs. Therefore, considering the current status of the ICT diffusion literature, we raise this research question: *How can the relationships among a large number of ICTs be described and analyzed in a scalable way?*

We answer this question in this study by offering a scalable methodology, based on computational analysis of the discourse about multiple ICTs, to understand innovation relationships. In the following, we first

briefly lay the theoretical foundation. Then we illustrate our approach with an empirical study of 47 ICTs over a decade. And finally we conclude by discussing the validity and benefits of our approach for ICT diffusion research and practice.

## 2. A relational view of ICT innovations

An ICT innovation is an information and communication technology that is perceived as new by an individual or other unit of adoption [15]. No innovation emerges and evolves in isolation. To varying degrees, the diffusion trajectories of related ICT innovations (as may be indicated by popularity or cumulative adoption) are interrelated. However, findings from the few sporadic studies of the diffusion of related innovations are inconsistent. On the one hand, for example, Wang [19] found that the popularity waves (as measured by the number of articles about each innovation) of four ICT innovations (manufacturing resource planning – MRP II, business process reengineering – BPR, customer relationship management – CRM, and supply chain management – SCM) that were closely related to enterprise resource planning (ERP) were negatively related to the popularity of ERP, indicating competing or substituting relationships between ERP and the other innovations. On the other hand, Berger and Heath [3] found positive correlation among related cultural innovations, confirming a main thesis in social cognition theory: When a particular concept is activated, related concepts may be activated as well [9].

To resolve this inconsistency, we must understand how ICT innovations are related. One useful approach is to classify ICT innovations by their attributes and functions. For example, Ein-Dor and Segev [6] surveyed the definitions of 17 ICT innovations in the Information Systems literature, identified from the definitions 31 attributes and 27 functions, and then described the innovations by two bit-vectors: a vector of attributes and a vector of functions. Further they performed quantitative methods such as multidimensional scaling (MDS) to visualize the relationships among the innovations in terms of their relative similarity/dissimilarity. While this approach can be implemented rigorously within the scope of each study, the choice of factors (such as attributes and functions of ICTs) for classification varies, depending on the specific opinions or background knowledge of the domain experts (most often the authors themselves) who participate in the study. In addition, as the number of ICT innovations increases, the effort by human experts to describe each innovation according to its

attributes and functions increases and the reliability of that classification work may decrease.

In search of an alternative, complementary approach to understanding ICT innovation relationships, we propose to conceptualize a network of innovations as part of an ecological system, where innovations can be likened to species in a competitive and symbiotic resource space [19, 21]. Innovations rely on the attention from communities of organizations and people with interests in producing and/or using the innovations. Each community emerges to make sense of an innovation and orchestrate material activities. The membership of the community evolves dynamically, as the collective attention to the innovation evolves. The flows of attention among innovations are both reflected and enabled by discourse – what have been said and written about the innovations. The discourse about an innovation candidly records human actions undertaken to describe, promote, or denounce the innovation [14]. Therefore, analysis of the discourse about multiple innovations can help us assess the complex relationships among ICT innovations. Although discourse analysis does not necessarily require the help of domain experts, where there is a large amount of discourse data, researchers often have to trade off between depth and breadth in discourse analysis.

## 3. Computational approach to the analysis of ICT relationships

Recent advances in computational linguistics have made it possible to achieve both depth and breadth in the analysis of very large sets of discourse data. Computational or automated analysis of discourse is a large, active interdisciplinary field with a variety of theories and techniques (see [13] for a non-technical primer). In this paper, we offer a computational methodology that can describe and analyze innovation relationships in a scalable way. We have conducted an empirical study to illustrate this methodology. Before we present the details of our approach, we need to describe the discourse data we have collected for this illustration.

### 3.1. Data collection

There are numerous discourse outlets, including books, magazines, conferences, blogs, wikis, and many others, where discourse data may be collected. In order to illustrate how our scalable computational methodology works in ICT innovation research, we decided to focus on a particular ICT trade magazine, *InformationWeek*, as the data source for this

illustration. As described below, the scale of the data we have collected from *InformationWeek* is large enough for us to demonstrate the scalability of our approach. While we could have selected data of a similar or larger scale from another outlet or set of outlets, we chose to focus on a trade magazine for two main reasons. First, compared to other types of outlets, the trade press involves most, if not all, types of actors in an innovation community (e.g., vendors, users, consultants, academics, regulators, investors, and journalists) in the production (e.g., writing, editing, and advertising) and consumption (e.g., subscribing, reading, and discussing) of trade publications. Second, it is the business of ICT trade magazines to actively represent what others have to say about various ICT innovations [20]. While sometimes a particular trade magazine may be over influenced by a particular actor (such as an influential vendor or important user groups), for any trade to take place, the trade press by definition has to strike a balance between the pursuits of vendors and user needs [20]. Such comprehensive and diverse representation of the actors and their ideas and actions makes trade press a desirable source of data in this study. Further, as one of the major ICT trade magazines, *InformationWeek* reaches approximately 440,000 business technology professionals at more than a quarter million unique locations. Over nearly three decades, the magazine has been helping IT leaders define and frame their business technology strategies. For an illustrative study like ours, *InformationWeek* provides an appropriately focused source of data.

We downloaded all of the articles published during a ten-year period (1998-2007) in *InformationWeek* using the Lexis/Nexis online database. Meanwhile, we compiled a list of 47 ICT innovation concepts (Table 1), ranging from enterprise software (e.g., CRM) to personal gadgets (e.g., iPod), from abstract concepts (e.g., artificial intelligence) to concrete products/services (e.g., YouTube), and from highly popular (e.g., e-business) to less well-known concepts (e.g., digital subscriber line – DSL). Admittedly, this list is *ad hoc*, but it serves the illustration purpose well because the list covers a broad range of ICT innovations in the examination period. We then extracted from the *InformationWeek* articles all paragraphs that contain any of the ICT innovations on the list. In doing so, we considered possible labels for each innovation, plural forms, and acronyms unique to the innovation. For example, in extracting paragraphs containing “digital subscriber line,” we also included paragraphs mentioning “digital subscriber lines” and “DSL.” Some ICT innovations had many paragraphs in the 10-year period while others had only a few. For example, there were more than 5,000 paragraphs

mentioning ERP. In total, 71,113 paragraphs were extracted, with about 1,500 paragraphs on average for each innovation.

### 3.2. Data analysis

To make sense of the relationships among the ICT innovations, we focused for this paper on the initial step of exploring the similarity of the innovations. One approach is to infer innovation similarity from the semantic similarity of the discourses about the innovations. Specifically, we employed Kullback-Leibler (KL) divergence, a probabilistic measure for differences in the pattern of word choices by authors, as a proxy for comparison of the semantic similarity of any two paragraphs extracted from *InformationWeek*. We also used hierarchical clustering analysis to aggregate the innovations in a hierarchical structure.

**3.2.1. KL divergence.** Originally introduced in 1951 [11] and considered a foundation of information theory [5], KL divergence is a statistic that quantifies in bits how close a probability distribution  $P$  is to another distribution  $Q$ . For probability distributions of discrete random variables, the KL divergence of  $Q$  from  $P$  is defined as:  $D_{KL}(P \parallel Q) = \sum_i P(i) \log(P(i) / Q(i))$ . KL

divergence is always non-negative. It equals zero if and only if the two distributions match exactly. KL divergence is commonly used for comparing the relative frequency of term use in pairs of discourses [12].

In this dataset, each ICT innovation is represented by concatenating all of the paragraphs that were automatically detected as mentioning the innovation. The use of language in the paragraphs constitutes a probability distribution of normalized unigram word counts and we calculated the KL divergence for each pair of innovations. The calculation generated an asymmetric 47x47 matrix with each column and row representing one of the 47 innovations. After symmetrization (by averaging the KL divergence in each direction, i.e.,  $\frac{D_{KL}(P \parallel Q) + D_{KL}(Q \parallel P)}{2}$ ), the value in

each cell of the matrix defines a distance (in the formal sense, satisfying the triangle inequality) between a pair of innovations. In order to classify the innovations and visualize their relationships, we performed clustering analysis on the symmetrized KL divergence matrix.

**3.2.2. Hierarchical clustering.** Cluster analysis is the process of grouping objects into unknown clusters such that the within-group variation is minimized and the between-group variation maximized [7]. The

agglomerative hierarchical clustering method groups objects on a series of levels, from the finest partition, in which each individual object forms its own cluster, and successively combines smaller clusters into larger ones until all objects are in one cluster. Agglomerative hierarchical clustering employs an aggregation criterion, or "linkage rule," to determine how the distance between two clusters should be calculated based on the distance scores of pairs of objects. The most well known aggregation criteria are single link, complete link, and average link [10]. The distance between two clusters is represented by the minimum, maximum, or average distance between any pair of objects, one object from each cluster. In single link clustering, two clusters with the smallest minimum pairwise distance are merged in each step. In complete link clustering, two clusters with the smallest maximum pairwise distance are merged in each step. And average link clustering is a compromise between the other two methods. We used the complete link in this study because it forms tight clusters with good global cluster quality [12].

## 4. Results

Our clustering analysis generated a hierarchy of clusters in a *dendrogram* (Figure 1), where vertical lines show joined clusters and the position of the lines on the scale from 1 to 25 indicates the distance at which clusters are merged. By inspecting the dendrogram, we have identified five natural clusters, all of which merged between 15 and 20 in the horizontal scale. These clusters are indicated by the five intersection points between the dendrogram and the vertical dotted line in Figure 1.

Cluster 1 includes 26 ICT innovations. Most of them are enterprise ICT applications. The hierarchical structure of this large cluster is shown in the dendrogram. For example, at the next granular level (around 15 in the horizontal scale), there exist two sub-clusters: one consisting of service-oriented ICT innovations such as OSS and web services and the other representing more traditional ICTs, which may be further differentiated at lower levels. Within the latter sub-cluster, we can see that, for example, the discourse on e-business is very similar to that on e-commerce. Similar relationships seem to exist in innovation pairs such as CRM and ERP, and knowledge management (KM) and groupware.

We counted the number of paragraphs each year mentioning each ICT innovation. The number of paragraphs about an innovation indicates the prevalence or popularity of the innovation in the discourse. For example, Figure 2 shows that the

popularity curves of the pair of very similar innovations (e-business and e-commerce) followed very similar patterns: both innovations enjoyed peak popularity around 2000 and then have lost much momentum since the dot-com crash. In contrast, the popularity curves of other very similar innovations followed very different patterns. Consider Figure 3, which shows the evolutionary trajectories of web services and SOA. The negatively correlated curves in the figure seem to suggest that the newer SOA replaced the older web services.

Cluster 2 includes five ICT innovations: DSL and virtual private network (VPN) are telecommunication technologies which can be applied to the other three innovations in the cluster. As Figure 4 shows, DSL and VPN had very similar popularity trajectories.

Only three innovations form Cluster 3 and four innovations form Cluster 5. These two clusters correspond to the so-called web 2.0 technologies that have become highly popular in recent years. Innovations in Cluster 3 share social network as a common feature. Cluster 5 represents text-based web 2.0 applications with user generated contents. The popularity curves in Figures 5 and 6 show that the innovations in these two clusters have generally experienced dramatic upswings circa 2004. Despite the similarity, the patterns of term use in the two clusters (as measured by symmetrized KL divergence) do not converge at the next level of aggregation. This interesting finding seems to suggest the substantial diversity of web 2.0 technologies.

Lastly, Cluster 4 has nine innovations all related to mobile or wireless technologies. Some, such as bluetooth and Wi-Fi, are the underlying mobile technologies. Others, such as TabletPC and iPod, are the devices enabled by the wireless/mobile technologies. Figure 7 shows that the rising popularity of iPhone coincided with the dwindling popularity of iPod, suggesting, once again, that the new replaces the old.

## 5. Discussion

### 5.1. Validity of the approach

The results from the KL-divergence and clustering analysis are consistent with our *a priori* knowledge about the relationship among these 47 ICT innovations. Such consistency provides reasonable confidence in the validity of the study's computational approach to understanding innovation relationships. Therefore, this study not only illustrates how our methodology works, but also validates our methodology, which may then be confidently applied to the more common scenarios

where *a priori* knowledge is unavailable, such as the cases of new or unknown innovations.

Further, we have triangulated the relationships we discovered in this study by employing multidimensional scaling (MDS), another visualization technique, to reduce the KL divergence matrix to two dimensions. Based upon a matrix of item-item similarities or dissimilarities, an MDS algorithm assigns a location to each item in a space such that the distances between the items correspond as closely as possible to the measured dissimilarities between the items. In other words, the proximity of items to each other in the space indicates how similar they are. We used the MDS procedure in SPSS based on the ALSCAL or alternating least squares scaling [17], the most popular MDS algorithm. In Figure 8, we present the 47 ICT innovations in the resulting two-dimensional scatter plot. This plot has an R-squared of 0.72, meaning that 72% of the variance of the scaled data can be accounted for by the MDS procedure. In this figure, we have used different colors to correspond to the five clusters identified in the previous clustering analysis. Generally speaking, most of the innovations in the same cluster (as shown in the dendrogram in Figure 1) are located close to each other in the MDS plot. This additional analysis has strengthened our confidence that this "KL divergence plus hierarchical clustering" analytical approach yields results that can be interpreted fairly easily, and that comport with our intuition.

## 5.2. Benefits of the approach

Our approach has several advantages. Foremost, computational analysis is scalable. The study has examined the discourse of 47 innovations in ten years, already surpassing the scale and scope of many ICT diffusion studies. While we have used just one trade magazine for this illustration, the capability of this approach is not limited in the number or type of discourse outlets. Further, although our own knowledge helped validate the methods in the illustration study, the methods themselves do not rely on expert knowledge. This feature differentiates our approach from other classification methods based on expert ratings or opinions (e.g., [6], [16]). Expert knowledge can be useful for specific research objectives, but methods relying on experts are not scalable. Moreover, unlike scalable analysis that relies on relatively thin observations, such as citations (e.g., [4]) or vocabulary (e.g., [1]), the KL divergence measure captures both the vocabulary and the rich context of the use of vocabulary. Overall, these advantages create a middle ground where both breadth and depth can be balanced in useful ways.

## 5.3. Implications for ICT diffusion research

The scalable computational approach we have demonstrated in this study can help to expand ICT diffusion research along three dimensions: innovations, data sources, and time. First, our approach can help ICT diffusion researchers overcome the limitation of single-innovation designs by concurrently exploring multiple innovations. By facilitating the study of multiple innovations and their relationships over an extended period of time, our approach enables researchers to develop theories of ICT diffusion on a more realistic footing. The reality of today's digital economy and information society calls for an understanding of how the diffusion processes of various ICTs interact with each other. In this study we have found that the popularity trajectories of semantically similar innovations sometimes demonstrate similar patterns (e.g., Figures 2 and 4) and sometimes suggest competition or substitution (e.g., Figures 3 and 7). Future research should build upon these findings to fully explore the contingencies.

Second, although we used only one data source (*InformationWeek*) to illustrate how the analysis is done, our approach is not limited to this or any other source. Any source may be biased by its own specifics and thus it would be useful to collect data from multiple sources and apply our methodology in the same way as we described above. On the one hand, a study may discover an objective representation of the innovation relationships by pooling the data from multiple sources in proportions that may represent key constituencies of innovation communities. On the other hand, researchers may analyze the data collected from each source and compare the results, revealing similarities and differences among various segments of the innovation communities that different sources represent. For example, a previous study using the KL divergence and MDS methods found that the ICT innovation relationships discovered in *InformationWeek* articles are very similar to those found in *ComputerWorld* articles [18]. Presently, we are in the process of collecting and analyzing data from other trade magazines and other sources such as the popular press, newspapers, academic journals, and informal media. In fact, we expect the clustering structures to differ across different sources due to their different styles or foci (such as the emphases on different actors in the innovation communities). Comparing the results across sources will enrich our understanding of the ICT innovations as well as the communities underlying these innovations.

Third, due to the dynamic nature of innovation relationships, it would be useful to conduct a series of

KL divergence and clustering analyses at multiple times. The evolving hierarchical structure of ICT innovations would reveal *what exactly is diffusing*, as innovations with changing meanings might cluster at different times with different innovations. This research capability could allow diffusion researchers to eventually understand the *speciation* of ICT innovations [2], that is, how various categories of ICT emerge, converge, and diverge over time. For example, the two separate clusters of web 2.0 technologies that we have found in this study may indicate a speciation moment, beyond which the two types of innovations could diffuse in different ways. To further illustrate the utility of our methodology in multi-period analysis, we sliced our dataset by year and performed the same analysis on each year's data. The page limit does not allow us to show all clustering results. Nevertheless, Figure 9 shows portions of the results for three years. Overall, the hierarchical structure of innovations did change over time. Specifically, CRM moved from a relatively separate innovation in 1998 to one closely associated with ERP in 2001 and 2004. However, in contrast, e-business and e-commerce are clustered together throughout the years.

In sum, this scalable computational methodology enables multi-innovation, multi-source, multi-period studies, potentially advancing knowledge about the dynamic relationships among ICT innovations.

#### 5.4. Implications for ICT adoption and use

This study provides the practitioners at the forefront of ICT adoption and use with a set of scalable tools for monitoring and understanding new and existing ICT innovations. In the absence of expert knowledge, as in the case where practitioners consider adopting a new and promising ICT, the same analysis can be applied to the discourse about the new technology and to the discourse about existing innovations as well. The new technology's cluster membership in the dendrogram may then serve to suggest its broad type. In using ICTs, practitioners might also find it useful to apply the KL divergence and clustering analysis to the discourse from internal sources such as emails, memos, manuals, and meeting minutes. The resulting dendrogram might then serve as taxonomy of ICTs. For providers of new IT products and services, such a taxonomy that emerges from data may complement the product categories designated in a top-down design process. The data-driven taxonomy of ICTs could also help identify redundancies and/or establish shared understanding of the ICTs in use across organizational units.

## 6. Conclusion

In this study, we have demonstrated a scalable computational approach, based on KL divergence and hierarchical clustering, for describing and analyzing the relationships among multiple ICT innovations. In essence, we utilize social artifacts (i.e., ICT discourse) to chart the technological terrains where ICTs are diffused and used. This *socio-technical* approach will realize its full potential in continued and sustained research on the socio-technical dynamics of ICT diffusion, adoption, and use.

## 7. Acknowledgements

This paper is based upon work supported by the National Science Foundation under Grants No. IIS-0729459 and SBE-0915645. We would like to thank Lidan Wang for her suggestion to use symmetrized KL divergence.

## 8. References

- [1] E. Abrahamson and M. Eisenman, "Employee-Management Techniques: Transient Fads or Trending Fashions?," *Administrative Science Quarterly*, vol. 53, no. 4, 2008, pp. 719-744.
- [2] R. Adner and D. A. Levinthal, "The Emergence of Emerging Technologies," *California Management Review*, vol. 45, no. 1, 2002, pp. 50-66.
- [3] J. A. Berger and C. Heath, "Idea Habitats: How the Prevalence of Environmental Cues Influences the Success of Ideas," *Cognitive Science*, vol. 29, no. 2, 2005, pp. 195-221.
- [4] L. M. A. Bettencourta, A. Cintron-Arias, D. I. Kaiser, and Castillo-Chavez, "The Power of a Good Idea: Quantitative Modeling of the Spread of Ideas from Epidemiological Models," *Physica A*, vol. 364, no. 2006, 2006, pp. 513-536.
- [5] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, Wiley, New York, 1991.
- [6] P. Ein-Dor and E. Segev, "A Classification of Information Systems: Analysis and Interpretation," *Information Systems Research*, vol. 4, no. 2, 1993, pp. 166-204.
- [7] B. Everitt, S. Landau, and M. Leese, *Cluster Analysis*, 4th ed., Oxford University Press, New York, 2001.
- [8] R. G. Fichman, "Going Beyond the Dominant Paradigm for Information Technology Innovation Research: Emerging Concepts and Methods," *Journal of the Association for Information Systems*, vol. 5, no. 8, 2004, pp. 314-355.

[9] S. T. Fiske and S. E. Taylor, *Social Cognition*, McGraw-Hill, New York, 1991.

[10] P. Hansen and B. Jaumard, "Cluster Analysis and Mathematical Programming," *Mathematical programming*, vol. 79, no. 1-3, 1997, pp. 191-215.

[11] S. Kullback and R. A. Leibler, "On Information and Sufficiency," *The Annals of Mathematical Statistics*, vol. 22, no. 1, 1951, pp. 79-86.

[12] C. D. Manning and H. Schütze, *Foundations of Statistical Natural Language Processing*, MIT Press, Cambridge, MA, 1999.

[13] D. W. Oard, "A Whirlwind Tour of Automated Language Processing for the Humanities and Social Sciences," in *Working Together or Apart: Promoting the Next Generation of Digital Scholarship*, Washington DC: Council on Library and Information Resources, 2009, pp. 34-42.

[14] N. Phillips and C. Hardy, *Discourse Analysis: Investigating Processes of Social Construction* vol. 50, Sage Publications, Thousand Oaks CA, 2002.

[15] E. M. Rogers, *Diffusion of Innovations*, 5th ed., Free Press, New York, NY, 2003.

[16] E. B. Swanson and N. C. Ramiller, "Information Systems Research Thematics: Submissions to a New Journal, 1987-1992," *Information Systems Research*, vol. 4, no. 4, 1993, pp. 299-330.

[17] Y. Takane, F. W. Young, and J. De Leeuw, "Nonmetric Individual Differences Multidimensional Scaling: An Alternating Least Squares Method with Optimal Scaling Features," *Psychometrika*, vol. 42, no. 1, 1977, pp. 7-67.

[18] C.-J. Tsui, P. Wang, K. R. Fleischmann, D. W. Oard, and A. Sayeed, "Understanding It Innovations through Computational Analysis of Discourse," in *30th International Conference on Information Systems*, Phoenix, AZ, 2009.

[19] P. Wang, "Popular Concepts Beyond Organizations: Exploring New Dimensions of Information Technology Innovations," *Journal of the Association for Information Systems*, vol. 10, no. 2, 2009, pp. 1-30.

[20] P. Wang and N. C. Ramiller, "Community Learning in Information Technology Innovation," *MIS Quarterly*, vol. 33, no. 4, 2009.

[21] R. H. Whittaker and S. A. Levin, *Niche: Theory and Application*, Dowden, Hutchinson & Ross, Inc., Stroudsburg, PA, 1975.

**Table 1. List of information and communication technology innovations**

AI	Artificial Intelligence	MP3	MP3 player
ASP	Application Service Provider	MySpace	MySpace
ATM	Automated Teller Machine	OLAP	Online Analytical Processing
BI	Business Intelligence	OSS	Open Source Software
Blog	Blog	Outsource	Outsourcing
Bluetooth	Bluetooth	PDA	Personal Digital Assistant
CAD	Computer Aided Design	RFID	Radio Frequency Identification
CRM	Customer Relationship Management	SmartCard	Smart Card
DigiCam	Digital Camera	SCM	Supply Chain Management
DLearn	Distance Learning	SFA	Sales Force Automation
DSL	Digital Subscriber Line	SocNet	Social Networking
DW	Data Warehouse	SOA	Service-Oriented Architecture
eBiz	eBusiness	Telecommute	Telecommuting
eCom	eCommerce	TabletPC	Tablet PC
EDI	Electronic Data Interchange	UtiComp	Utility Computing
ERP	Enterprise Resource Planning	Virtualization	Virtualization
GPS	Global Positioning System	VPN	Virtual Private Network
Grpware	Groupware	Web2.0	Web 2.0
IM	Instant Messaging	WebServ	Web Services
iPhone	iPhone	WiFi	Wi-Fi
iPod	iPod	Wiki	Wiki
KM	Knowledge Management	Wikipedia	Wikipedia
Linux	Linux	YouTube	YouTube
Multimedia	Multimedia		

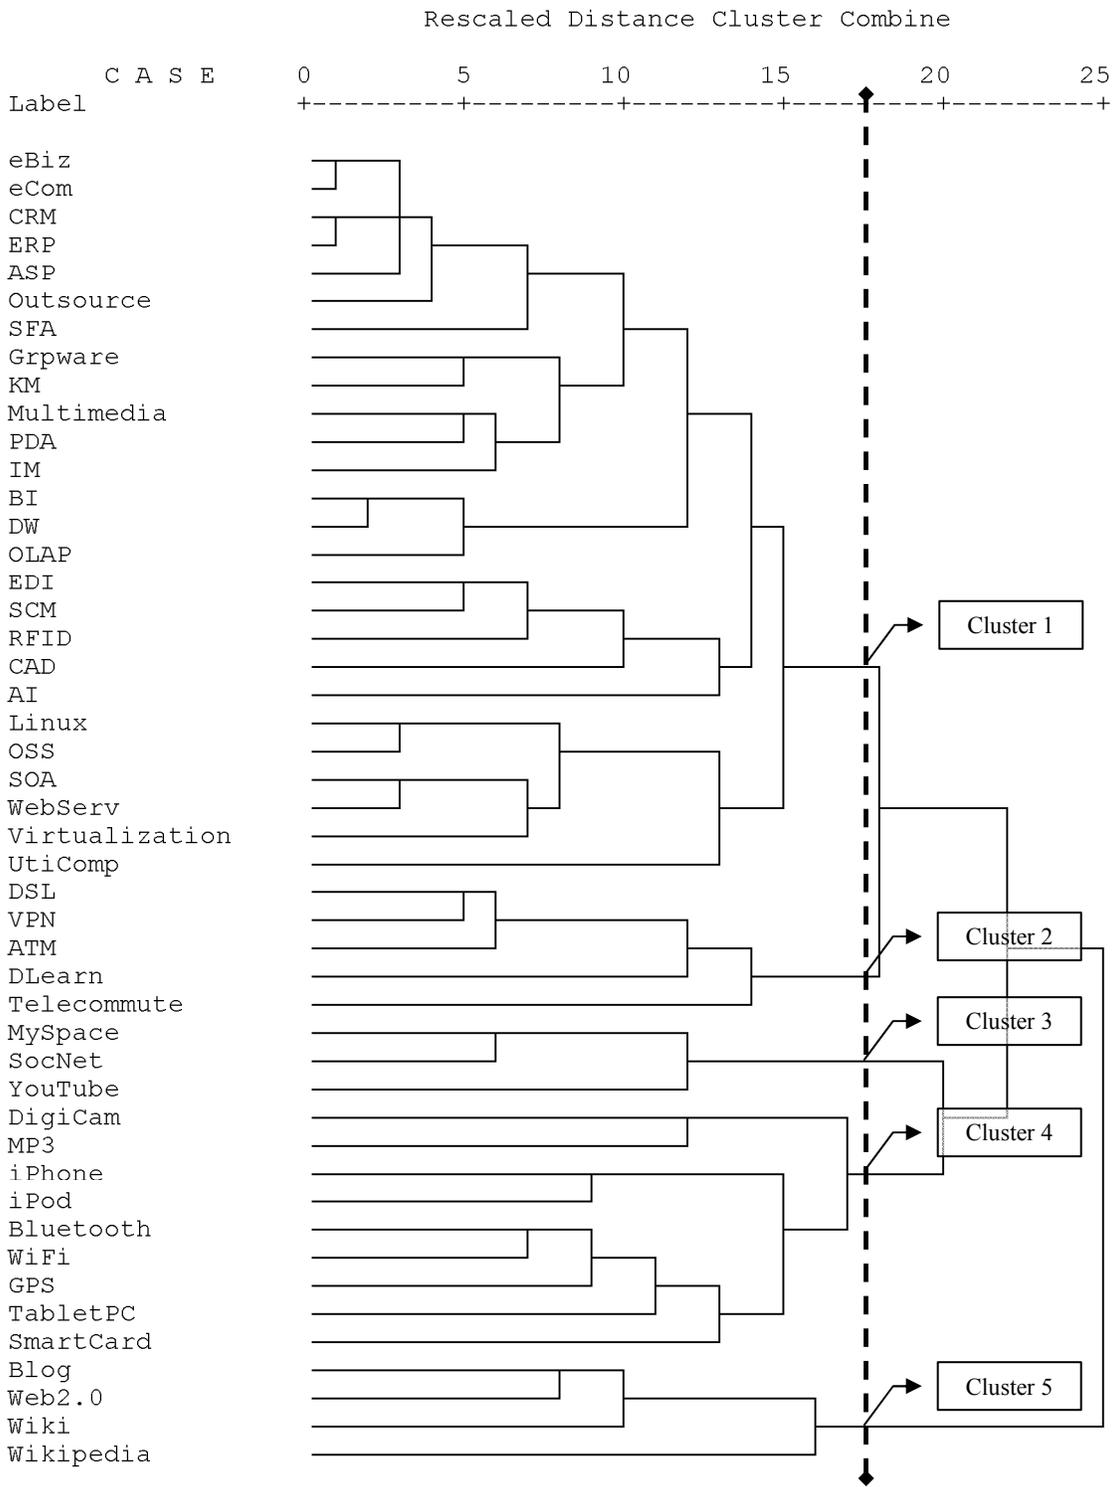


Figure 1. Dendrogram produced by hierarchical clustering analysis of the *InformationWeek* data

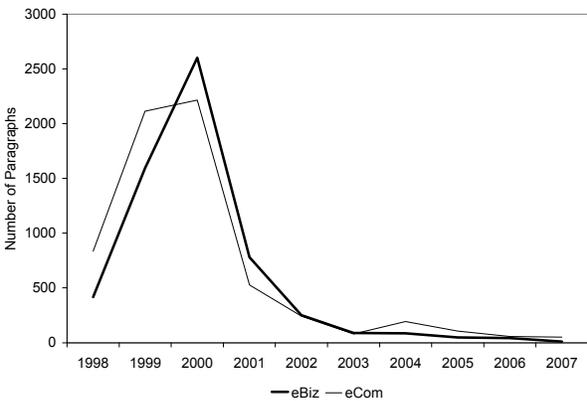


Figure 2. Popularity of e-business and e-commerce

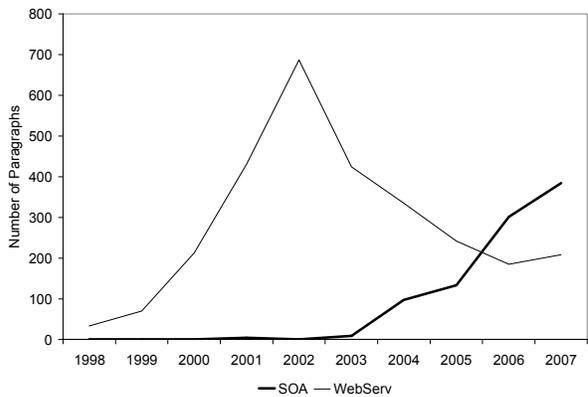


Figure 3. Popularity of SOA and web services

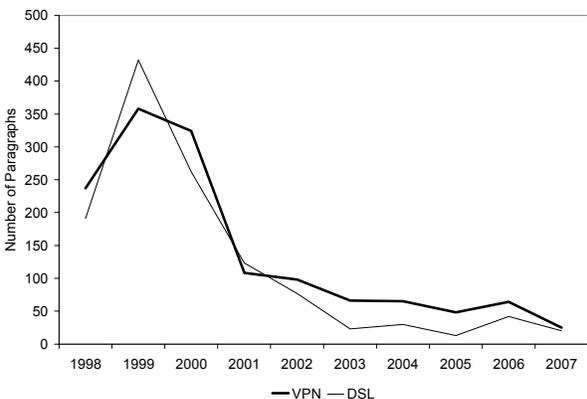


Figure 4. Popularity of DSL and VPN

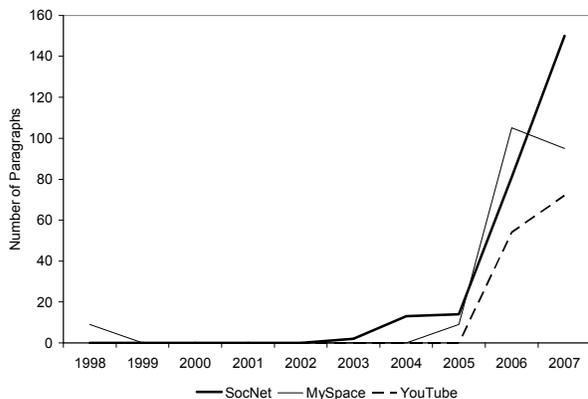


Figure 5. Popularity of social networking innovations

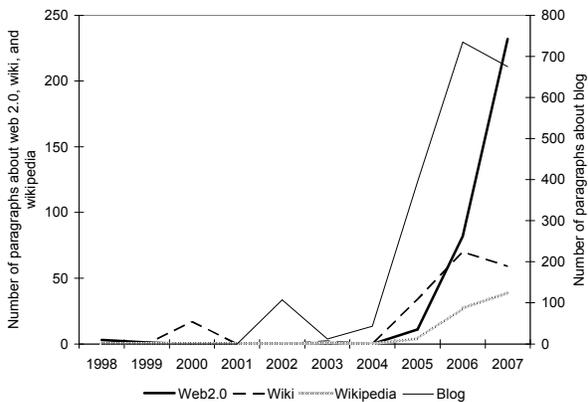


Figure 6. Popularity of web 2.0 innovations with user generated contents

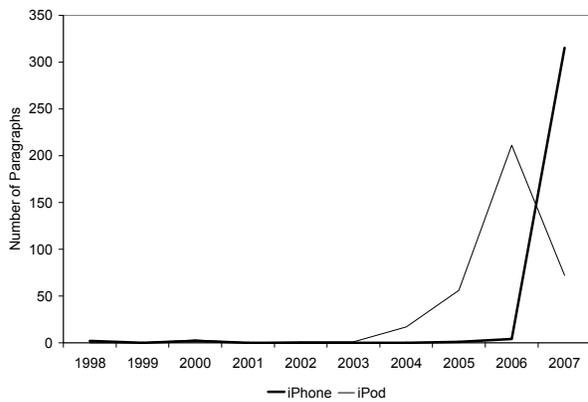


Figure 7. Popularity of iPhone and iPod

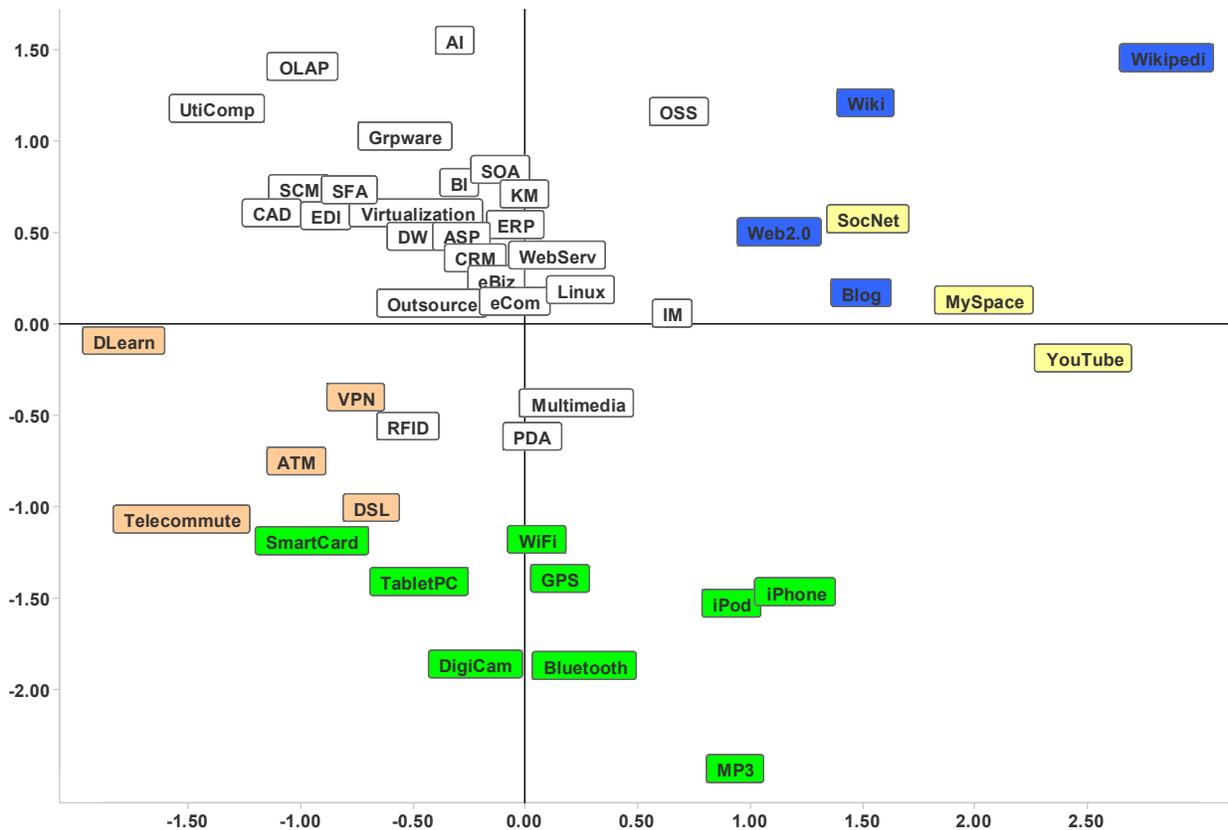


Figure 8. MDS plot of the 47 ICT innovations from 10-year *InformationWeek* data

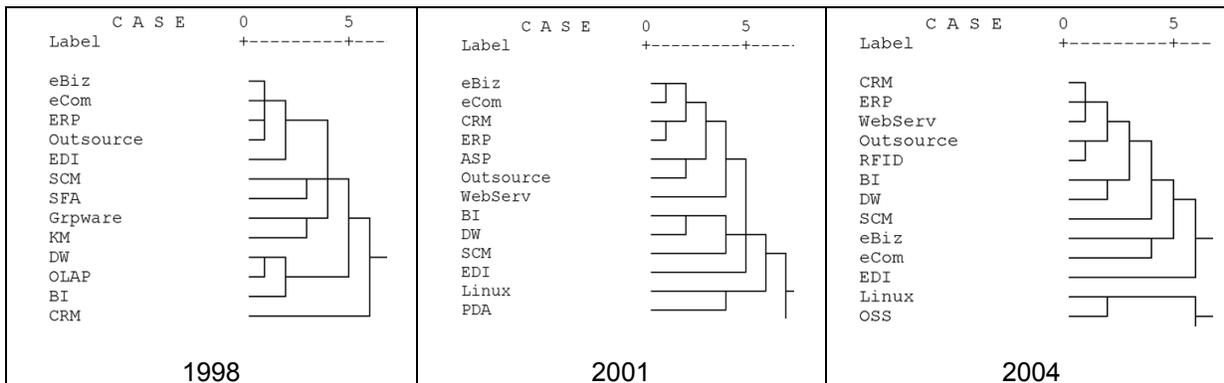


Figure 9. Portions of yearly dendrograms produced by hierarchical clustering analysis of the *InformationWeek* data