



Toward a deeper understanding of Visualization through keyword analysis

Petra Isenberg, Tobias Isenberg, Michael Sedlmair, Jian Chen,
Torsten Möller

**RESEARCH
REPORT**

N° 8580

August 2014

Project-Team Aviz



Toward a deeper understanding of Visualization through keyword analysis

Petra Isenberg^{*}, Tobias Isenberg[†], Michael Sedlmair[‡], Jian Chen[§],
Torsten Möller[¶]

Project-Team Aviz

Research Report n° 8580 — August 2014 — 24 pages

Abstract: We present the results of a comprehensive analysis of visualization paper keywords supplied for 4366 papers submitted to five main visualization conferences. We describe main keywords, topic areas, and 10-year historic trends from two datasets: (1) the standardized PCS taxonomy keywords in use for paper submissions for IEEE InfoVis, IEEE Vis-SciVis, IEEE VAST, EuroVis, and IEEE PacificVis since 2009 and (2) the author-chosen keywords for papers published in the IEEE Visualization conference series (now called IEEE VIS) since 2004. Our analysis of research topics in visualization can serve as a starting point to (a) help create a common vocabulary to improve communication among different visualization sub-groups, (b) facilitate the process of understanding differences and commonalities of the various research sub-fields in visualization, (c) provide an understanding of emerging new research trends, (d) facilitate the crucial step of finding the right reviewers for research submissions, and (e) it can eventually lead to a comprehensive taxonomy of visualization research. One additional tangible outcome of our work is an application that allows visualization researchers to easily browse the 2600+ keywords used for IEEE VIS papers during the past 10 years, aiming at more informed and, hence, more effective keyword selections for future visualization publications.

Key-words: Keyword analysis, research themes, research topics, taxonomy, visualization history, theory.

^{*} Petra Isenberg is with Inria, France. E-mail: petra.isenberg@inria.fr.

[†] Tobias Isenberg is with Inria, France. E-mail: tobias.isenberg@inria.fr.

[‡] Michael Sedlmair is with the University of Vienna, Austria. E-mail: michael.sedlmair@univie.ac.at.

[§] Jian Chen is with the University of Maryland, Baltimore County, USA. E-mail: jichen@umbc.edu.

[¶] Torsten Möller is with the University of Vienna, Austria. E-mail: torsten.moeller@univie.ac.at.

**RESEARCH CENTRE
SACLAY – ÎLE-DE-FRANCE**

1 rue Honoré d'Estienne d'Orves
Bâtiment Alan Turing
Campus de l'École Polytechnique
91120 Palaiseau

Vers une meilleure compréhension de la visualisation à travers l'analyse de mots-clés

Résumé : Nous présentons les résultats d'une analyse exhaustive de mots-clés pour 4366 articles de visualisation soumis à cinq principales conférences de visualisation. Nous décrivons les mots clés principaux, domaines thématiques, et les tendances historiques sur 10 ans pour deux jeux de données: (1) les mots-clés standardisés de la taxonomie PCS actuellement utilisés pour la soumission d'articles à IEEE InfoVis, IEEE Vis-SciVis, IEEE VAST, EUROVIS, et IEEE PacificVis depuis 2009 et (2) les mots-clés choisis par les auteurs pour les articles publiés dans la série de conférences de visualisation IEEE (appelée IEEE VIS depuis 2004). Notre analyse des sujets de recherche en matière de visualisation peut servir de point de départ pour (a) aider à créer un vocabulaire commun pour améliorer la communication entre les différents sous-groupes du domaine de la visualisation, (b) faciliter la compréhension des différences et points communs entre ces différents sous-groupes, (c) mieux comprendre les nouvelles tendances de recherche émergentes, (d) faciliter l'étape cruciale consistant à trouver les bons experts pour la relecture de soumissions, et (e), à terme, conduire à une taxonomie complète de la recherche en visualisation. Un résultat supplémentaire et tangible de notre travail est une application qui permet aux chercheurs en visualisation de parcourir facilement les mots-clés utilisés dans plus de 2600 articles de la conférence IEEE VIS au cours des 10 dernières années. Cette application vise à faciliter la sélection mieux informée et par conséquent plus efficace de mots clés pour les futures publications en visualisation.

Mots-clés : analyse de mots-clés, thèmes de recherche, taxonomie, l'histoire de la visualisation, la théorie.

1 Motivation

One of the main reasons why the field of visualization is such a fascinating field of research is due to its diversity. We not only refer to the diversity of applications, but the diversity of research methods being employed, the diversity of research contributions being made, as well as the diversity of its roots.

Diversity of roots: The term *visualization* can be understood very broadly, expressing a long history of its use in common language. Therefore, it is not surprising that concepts of visual thinking have penetrated many areas of science, engineering, and philosophy. The field of modern (computer-based) visualization has been greatly influenced by research methods from the fields of numerics and computer graphics, which have given it its birth in 1990. The impact of human-computer interaction affected the birth of the InfoVis community in 1995 and the influence of applied statistics (such as data mining) and cognition has led to the establishment of VAST in 2006.

Diversity of research methods: Given its diverse roots, visualizations remains a highly inter-disciplinary field that borrows and extends research methods from other fields. Methods come from fields as diverse as the broader computer science, mathematics, statistics, machine learning, psychology, cognitive science, semiotics, design, or art.

Diversity of contributions and applications: Based on these diverse influences, the results of visualization research can be manifold: from engineering solutions to dealing with large data sources (such as real-time rendering solutions, distributed and parallel computing technologies, novel display devices, and visualization toolkits) to understanding design processes (as in perceptual guidelines for proper visual encodings and interaction or facilitating collaboration between different users through visual tools) to scientific inquiries (such as improved understanding of perceptual and cognitive processes).

While all these diverse influences make the field of visualization research an exciting field to be a part of, they also create enormous challenges. There are different levels of appreciation for all aspects of visualization research, communication challenges between visualization researchers, and the challenge of communicating visualization as a research science to the outside. These issues lead, in particular, to the frequently asked question “what is visualization?”—among funding agencies or even between colleagues. Given our field’s broad nature, we need to ask how we can comprehensively describe and summarize all on-going visualization research. These are not just theoretical and philosophical questions, but the answer to these questions has many real-world impacts—from such simple (but career-deciding) questions as finding the right reviewers during peer-review to administrative strategic decisions on conference and journal structures and foci.

So while “what is visualization?” is a fundamental question, it is little discussed within our community. In fact, thus far the approaches have mostly focused on understanding some sub-field of visualization (e. g., [16, 29, 34]) but the question for the broader community has rarely been tackled beyond general textbook definitions (e. g., [7]). Those who have approached the problem, did so in a top-down approach. For example, several taxonomies were suggested by experts based on tasks, techniques, or data models (e. g., [8, 29, 35]). Another way of splitting visualization into more focused areas has been through specific application foci (e. g., VisSec, BioVis, SoftVis, etc.).

What is missing in this picture is a bottom-up analysis: What types of visualization research are actually happening as expressed by single research contributions in the visualization conferences and journals. Our paper is one of the first steps in this direction. We analyze author-assigned keywords from the three IEEE VisWeek/VIS conferences of the past ten years as well as author-selected taxonomy entries in the submission system for three IEEE VisWeek/VIS conferences, EuroVis, and PacificVis of the past six years. Based on this analysis, we make the following contributions:

Mapping visualization research: In Sect. 4, through the vehicle of keyword analysis, we build a conceptual map of all visualization work as indexed by individual authors. Our main assumption here is that, while each single keyword might be understood in a slightly different way by different researchers, their co-occurrence with other keywords clarifies their meaning, especially when aggregated over many

different usages (i. e., many research papers in a major publication venue). This co-occurrence analysis is the basis for deriving clusters and, therefore, research sub-fields. The use of keywords seen over the past ten years also allows us to understand historical trends and we report on the most prominent declining and rising keywords within all of visualization.

keyvis.org: In Sect. 5 we describe a Web-based search tool that (a) makes the keyword meta-data available to a broad set of people, that (b) helps researchers/users of visualizations to quickly find the right papers that relate to a given topic, and that (c) helps visualization researchers find descriptive keywords for their publications.

Terminology: Visualization research is influenced by a diverse set of application domains. The vocabulary of visualization is influenced by all these application areas. However, the resulting diversity of terms are only understood by a small set of visualization researchers and, therefore, hinders the dissemination of research results and insights across all visualization sub-fields. This is very well articulated in the collection of keywords throughout all of visualization in the past ten years. We are the first that collected and “cleaned” this data, making it available to the benefit of our community, allowing its systematic analysis. In Sect. 6 we argue for a unification of different vocabulary.

2 Related Work

We are not the first to have made an effort to summarize a large set of visualization papers in order to understand topics or trends. One of the earliest such efforts was a summary and clustering of visualization research papers by Voegele [36] in 1995 in the form of a two-dimensional clustering of all visualization papers up to this point. Other efforts have focused on specific aspects of visualization research. Sedlmair et al. [28], for example, did a thorough analysis of all design study papers to summarize practices and pitfalls of design study approaches. Further, Lam et al. [22] studied the practice of evaluations in Information Visualization papers which was then extended to include all visualization papers by Isenberg et al. [20]. Others have surveyed, for instance, the literature on interactive visualization [21, 39], on tree visualizations [27], on quality metrics in high-dimensional data visualization [2], on human-computer collaborative problem-solving [11], or on visualization on interactive surfaces [19].

In other disciplines, specific techniques have been used to analyze the scientific literature more broadly: to get a better sense of global research trends, links and patterns within the scientific literature. Co-word analysis is one approach among others (e. g., co-citation analysis) that has tackled the problem by analyzing the scientific literature according to the co-occurrence of keywords, words in titles, abstracts, or even in the full texts of scientific articles [5, 9, 13, 15, 23, 38]. Callon et al. [6], in particular, wrote a seminal book on the topic that provides several methods that others have used and extended upon.

Co-word analysis has been used in different research areas, e. g., polymer chemistry [5], acid rain research [23], or education [26]. Others further restricted the scope of the literature to specific countries, such as Hu and Liu et al.’s [18, 24] co-word analysis on library and information science in China. The closest co-word analysis studies to our work are Coulter et al.’s [10] work on the software engineering community, Hoonlor et al.’s [17] general investigation of the computer science literature, and, most recently, Liu et al.’s [25] analysis of the human-computer interaction literature. Liu et al. examined papers of the ACM CHI conference from 1994–2013, identified research themes and their evolution, and classified individual keywords as popular, core, or backbone topics. We employ similar approaches as used in particular in Liu et al.’s work. Naturally, however, we differ as our focus is on a different research community, visualization, with different keywords, trends, and patterns and a different historical evolution.

In the visualization and data analysis literature, the closest work to ours is Chuang et al.’s [9] machine learning tool for topic model diagnostics, and Görg et al.’s [14] visual text analysis using Jigsaw, and the CiteVis tool [30]. These lines of work are not per se co-word analyses. However, their data sources

Table 1: Keyword datasets we collected.

conference	author	IEEE	INSPEC	years	PCS	years
IEEE InfoVis	✓	✓	✓	2004–2013	✓	2008–2013
IEEE Vis/SciVis	✓	✓	✓	2004–2013	✓	2008–2013
IEEE VAST	✓	✓	✓	2006–2013	✓	2008–2013
EuroVis					✓	2008–2013
IEEE PacificVis					✓	2009–2013

also include visualization research papers. In contrast, we primarily focus on the results of our analysis of themes and trends in the visualization literature rather than on the description of any specific tool or algorithm.

3 Co-Word Analysis of the Visualization Literature

For our analysis of the visualization research literature we considered several data sources and analyzed two in detail.

3.1 Datasets

We collected the following datasets:

Author-assigned keywords: Keywords freely assigned by the authors to their research paper. We collected this data manually from paper PDFs as the IEEE digital library contained partially incorrect data.

IEEE terms: These index terms are manually assigned by an IEEE-hired “university expert” [33] using the IEEE taxonomy [31, 32].

INSPEC terms: These index terms are automatically derived [33] either using the same IEEE taxonomy (controlled) or without constraints (non-controlled).

PCS taxonomy keywords: These keywords chosen from a pre-defined visualization keyword taxonomy and are assigned by the authors during the submission of their research paper.

Details of the data collected can be found in Table 1. Table 2 provides an overview of the top ten keywords from each data source for IEEE VisWeek/VIS papers for 2004–2013. After a first inspection of the IEEE and INSPEC terms, we found these datasets to be unsuitable to answer our research questions and chose not to analyze them further.

3.2 Keyword Analysis

Our general approach was first to extensively clean the data and then to analyze similarities and differences of both keyword sets. We largely followed Liu et al.’s [25] process for analyzing keyword data for the ACM Conference on Human Factors in Computing Systems (CHI), extending it with expert coding and an analysis of the taxonomy use.

3.2.1 Author-Assigned and Expert-Coded Keyword Datasets

Our collection of papers from the three IEEE VisWeek/VIS conferences (see Table 1) contained 1097 published papers (excluding posters). Out of these, 58 contained no author-assigned keywords, yielding a set of 1039 papers we considered in our analysis. These papers contained a total of 2823 unique

Table 2: Top ten keywords *only* of IEEE InfoVis, IEEE Vis/SciVis, and IEEE VAST for the different classification schemes; 2004–2013.

#	author keywords	PCS taxonomy (lower level; only 2008–13)	IEEE terms	INSPEC controlled	INSPEC non-controlled
1	information visualization	95 visual knowledge discovery	296 data visualization	685 data visualisation	540 data visualization
2	visual analytics	86 graph/network data	273 visualization	272 data visualization	393 information visualization
3	visualization	81 volume rendering	237 rendering (comp. graph.)	176 rendering (comp. graph.)	227 interactive visualization
4	volume rendering	72 coordinated and multiple views	233 visual analytics	127 data analysis	135 visual analysis
5	flow visualization	48 user interfaces	221 computer graphics	117 interactive systems	118 data analysis
6	interaction	38 biomedical and medical visualization	218 data analysis	107 computational geometry	111 volume rendering
7	evaluation	26 time-varying data	204 computational modeling	104 graph theory	67 GPU
8	isosurfaces	25 visualization system and toolkit design	194 data mining	103 pattern clustering	65 visual representation
9	volume visualization	24 high-dimensional data	191 displays	95 medical image processing	63 direct volume rendering
10	parallel coordinates	23 interaction design	185 humans	95 feature extraction	59 visual exploration

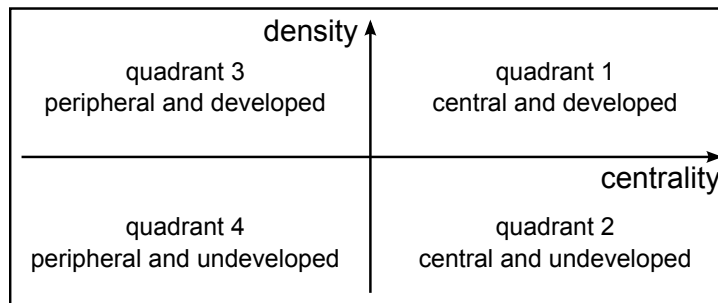


Figure 1: Strategic diagram to characterize the topic clusters (after [15]).

keywords. Next, we engaged in an extensive manual cleaning pass in which we consolidated keywords that were the same but presented either as singulars/plurals, with spelling mistakes, or as acronyms. This yielded a cleaned dataset that contained 2629 unique keywords that occurred a total of 4780 times across all papers.

Next, we engaged in a manual expert coding of these keywords in order to find higher-level clusters of keyword topics in this dataset. All authors of this paper participated in a multi-pass coding of the cleaned keywords and assigned one or more higher-level cluster codes from a freely evolving, personal code set. Next, we calculated the clusters that emerged from each expert’s coding, encoded them in individual keyword-cooccurrence matrices, and performed a hierarchical clustering (e. g., [1]) on the joint matrix using average linkage with the Bray-Curtis distance metric [4]. This clustering yielded an ordering of keywords that most strongly matched the consensus of each expert’s clustering but did not yet yield satisfactory final clusters. We, thus, engaged in a final manual clustering pass in which we labeled the clusters that emerged or broke them apart if the clustering had not provided meaningful distinctions. The resulting set of keywords contained 156 unique higher-level keywords that occurred a total of 4026 times across all papers (this number is lower than for the cleaned data as potentially generated duplicates per paper were removed). Throughout our paper we refer to this keyword set as “expert keywords.”

3.2.2 PCS Taxonomy Keyword Dataset

The PCS taxonomy keyword data consists of the authors’ classification of their papers according to a visualization taxonomy (called “PCS taxonomy” from now on). This PCS taxonomy was created in its present form in 2009¹ and was used for the different visualization venues starting that year (for PacificVis starting 2010). For the year 2008 (and 2009 for PacificVis) we mapped keywords from a previous to the current taxonomy. We analyzed anonymized data on 3,927 paper submissions that included the submission IDs and keywords, but the titles only of papers that were finally accepted to the conferences. The dataset included 1070 accepted and 2857 rejected papers. We did not have access to author names or any other identifying information for rejected papers. Considering data from both accepted and rejected papers allowed us to analyze the types of topics that the community was working on in a given year. In total, the PCS taxonomy includes 14 higher-level categories and 127 lower-level keywords (4–17 per higher-level category). In addition, the dataset also included the paper submission types (one of *technique paper*, *system paper*, *application/design study paper*, *evaluation paper*, or *theory/model paper*) for IEEE InfoVis, IEEE Vis/SciVis, IEEE VAST, and EuroVis, but not for PacificVis.

¹The first version of the taxonomy was created by the 2007 EuroVis paper chairs, Ken Museth, Torsten Möller, and Anders Ynnerman. This version was iterated by several people over the next two years. After a broad consultation within in the visualization community, the final version was assembled by Hans-Christian Hege, Torsten Möller, and Tamara Munzner. This effort was supported by the VGTC—the IEEE Visualization and Graphics Technical Committee.




Figure 2: The frequency of the author-assigned keywords (2004–2013, published VIS papers), sorted by their rank (both axes log-transformed).

3.2.3 Analysis Process

To analyze these keyword datasets (cleaned *author*-assigned keywords, *expert* keywords, *PCS* keywords) we first filtered each dataset. We removed keywords that occurred less than a minimum threshold and also excluded higher-level terms as outlined in Table 3. Next, we generated document-keyword matrices for each dataset with the keywords as variables (rows) and documents as observations (columns). Each cell contained a 0 if a keyword was not present in a paper and a 1 if it was. On each matrix, we performed a correlation computation using Python’s NumPy `corrcoef` function that yielded a correlation matrix holding keyword-keyword correlation coefficients. On each correlation matrix we performed a hierarchical clustering using Ward’s method [37] and a squared Euclidean distance metric. We also generated a keyword network in which two keywords were linked if their correlation was > 0 and each link was assigned its respective correlation value. From this network we computed the density of each cluster with the median of all inter-cluster links and the centrality by computing the sum of square of all links from within a cluster to another cluster. We plotted centrality and density in *strategic diagrams* [5, 10, 15, 18, 24, 25]. These diagrams distinguish four quadrants (Fig. 1) that characterize the different clusters based on their centrality within the research domain and on how developed they are. The diagram axes are centered on the median of the observed density and centrality values.

We found that the frequency distribution of the author-assigned keywords followed a power law ($\alpha = 1.66$, $R^2 = 0.80$) as is evident from the linear shape of the log-transformed frequency over rank plot (Fig. 2). Similar to the CHI keywords [25], this indicates that the research structure within visualization is a scale-free network. This means that the research network includes a small number of popular nodes that represent topics on which researchers in the field concentrate. These central nodes serve as hubs to which other research clusters are connected [24, 25]. In contrast to the author-assigned keywords, the distributions of expert keywords and PCS taxonomy keywords do not follow a power law distribution. For the PCS taxonomy, for example, this means that we do *not* observe that the majority of keywords would rarely get used at all, which indicates that the present taxonomy does not contain a majority of topics which would be irrelevant.

4 Results

Two main research questions drove our analysis of the data. In particular, we were interested in understanding major research themes and their relationship to other themes (Sect. 4.1) and the importance and evolution of individual keywords (Sect. 4.2).

Table 3: The analyzed keyword datasets with their occurrence thresholds applied to keywords, number of remaining keywords analyzed, and target number of clusters set for the cluster analysis.

dataset	occur. thresh.	excluded terms	remaining keywords	target clusters
author	6	visualization, information visualization, scientific visualization, visual analytics	101	16
expert	10			
PCS	0	none	127	15

Table 4: Cluster result for author-specified keywords. Keywords are sorted by frequency with the two most frequent keywords highlighted in bold.

ID	keywords (InfoVis, Vis/SciVis, VAST; 2004–2013)	N	#	cw-#	centr.	dens.
A1	isosurfaces, direct volume rendering , transfer function, gpu, time-varying data, molecular visualiz., raycasting, sampling	8	11.5	1.2	0.286	0.085
A2	volume visualization, illustrative visualization , focus+context techniques	3	16.0	3.3	0.182	0.219
A3	volume rendering, graphics hardware , interpolation, visualization systems, level-of-detail, vector field visualization, unstructured grids, astronomy, programmable graphics hardware	9	7.0	0.6	0.318	0.047
A4	focus+context visualization, coordinated & multiple views , interactive visualization, interactive visual analysis	4	7.0	0.7	0.185	0.094
A5	interaction, sensemaking , multiple views, network visualization, time series data, social networks, exploratory visualization	7	9.0	0.8	0.316	0.063
A6	geovisualization, spatio-temporal data	2	9.5	4.0	0.165	0.449
A7	graph visualization, clustering , treemaps, node-link diagrams, hierarchies	5	15.0	1.8	0.278	0.143
A8	parallel coord., multi-variate data , user interfaces, high-dimensional data, scatterplots, visual data mining, user interaction	7	12.0	1.3	0.218	0.091
A9	evaluation, design , framework, animation, knowledge discovery, taxonomy, experiment	7	8.0	1.0	0.196	0.081
A10	focus+context, multi-variate visualization , uncertainty, classification, dimensionality reduction, time-varying, visual analysis, principal component analysis	8	9.0	0.5	0.239	0.047
A11	user study, human-computer interaction , visual knowledge discovery, design study, medical visualiz., collaboration, diffusion tensor imaging, linked views, graph layout, text visualiz., non-photorealistic rendering, tiled displays, social data analysis, bioinformatics, applications of visualiz., glyphs, illustrative rendering, intelligence analysis, geographic visualiz.	19	7.0	0.1	0.283	0.011
A12	feature extraction, unsteady flow visualization	2	10.0	3.0	0.143	0.308
A13	uncertainty visualiz., perception , graph drawing, multi-dimen. visualiz., multi-dimen. scaling, data mining, qualit. eval.	7	9.0	0.5	0.213	0.049
A14	flow visualization, vector fields , streamlines, 3d vector field visualization	4	8.0	2.3	0.264	0.155
A15	isosurface extraction, hardware acceleration , point-based visualization	3	7.0	1.3	0.126	0.176
A16	topology, comparative visualization , morse-smale complex, data exploration, contour tree, feature detection	6	7.5	0.5	0.148	0.066

4.1 Analysis of Major Topic Areas

To understand major themes we analyzed the results of the clustering and the generated network graphs. We first report on the results per keyword set and then discuss how these results relate with each other.

4.1.1 Author-Assigned Keywords

CLUSTER ANALYSIS: We created 16 clusters from the 101 author-assigned keywords. The number of clusters we chose was based on manual inspection of the content validity of the hierarchical clustering result. Table 4 summarizes the clusters (A1–A16) that were created. In the table we report the keywords assigned to each cluster, the size of the cluster, the average amount of times each keyword occurred in the

whole corpus, the average amount of times the keywords in each cluster co-occurred (CW-frequency), the centrality of the subnetworks created from each cluster, and the density of the subnetworks. Keywords are sorted by number of occurrences in the dataset and the top two keywords are bolded and are used to refer to the clusters in the text.

By inspecting of the clusters we found that several of them seemed to form more coherent themes than others; many keywords made intuitive sense together, while some other clusters seemed to include outliers. A11, in particular, contains a large number of diverse keywords around applications, encodings, and specific visualization techniques. There were only few exceptions where individual keywords seemed slightly out of place. For example, the keywords *astronomy* in A3 or *multi-dimensional scaling* in A13, at first, may seem misplaced. The correlation matrix, however, reveals that *astronomy* was highly positively correlated with *interpolation*, the second most frequent keyword in this cluster, same as *multi-dimensional scaling* with *graph drawing* and *multi-dimensional visualization*. We similarly checked other keywords that seemed out of place and confirmed their placements in the respective clusters to be meaningful with respect to the correlation scores. The network graph in Fig. 4 can also further serve as an analysis aid to check the strongest correlations among keywords.

Table 4 also includes several clusters with similar keywords but using slightly different wording. For example, A1 (*isosurfaces, direct volume rendering*), A2 (*volume visualization, illustrative visualization*), and A3 (*volume rendering, graphics hardware*) center around volumes. Similarly, Clusters A12 and A14 relate to flow visualization. Both volume and flow visualization have been central topics in the IEEE Vis/SciVis conference, illustrated by the fact that both *volume rendering* and *flow visualization* were the top two non-generic keywords in this dataset. Thus, the emerged clusters around these two general topic areas suggest that the two topics are established and sub-areas of research have emerged. It was more surprising to us that other related terms emerged under different terminology. For example, three focus+context-related keywords were not clustered together: *focus+context techniques* (A1), *focus+context visualization* (A4), and *focus+context* (A10). Similarly, *graph visualization* was placed in A7 and *graph drawing* in A13. This suggests that sub-communities within visualization may use slightly different terminology to label similar or related topics.

NETWORK ANALYSIS: Next, we analyzed the strategic diagram in Fig. 3(a) and the keyword cluster map in Fig. 4. Fig. 3(a) shows that topic areas emerge in each of the previously mentioned quadrants. In the following we do not discuss clusters A1, A4, A8, A9, and A13 as their density or centrality measures were very close to the median and, thus, their membership in specific quadrants was not very strong.

The topics in Quadrant I of the strategic diagram (top right) are considered motor themes or “main-stream” topics as they are both internally coherent and central to the research network. In our network, clusters A7 (*graph visualization, clustering*) and A14 (*flow visualization, vector fields*) are considered motor themes.

Topics in Quadrant II (bottom right) are A3 (*volume rendering, graphics hardware*), A5 (*interaction, sensemaking*), A10 (*focus+context, multi-variate visualization*), and A11 (*user study, human-computer interaction*). These themes are weakly linked together but present strong ties to other themes (clusters) within the network. They are considered basic and transversal themes.

Quadrant III (top left) includes developed but isolated themes. These are the smaller sub-networks A2 (*volume visualization, illustrative visualization*), A6 (*geovisualization, spatio-temporal data*), and A12 (*feature extraction, unsteady flow visualization*), and A15 (*isosurface extraction, hardware acceleration*). Their high density scores indicate that they are relatively strongly tied within themselves—also evident in the network graph in Fig. 4 where these clusters remain mostly tied together through strong (thick) links. While all four subnetworks also had the low centrality scores, they remained connected to the larger central network in Table 4.

Finally, Quadrant IV (lower left) contains emerging or declining topics as they are neither very dense nor central to the network. In particular, A16 (*topology, comparative visualization*) lies in this quadrant.




Figure 3: Strategic diagrams of the cluster results for (a) author-assigned, (b) expert, and (c) PCS keywords. Black lines indicate the medians.

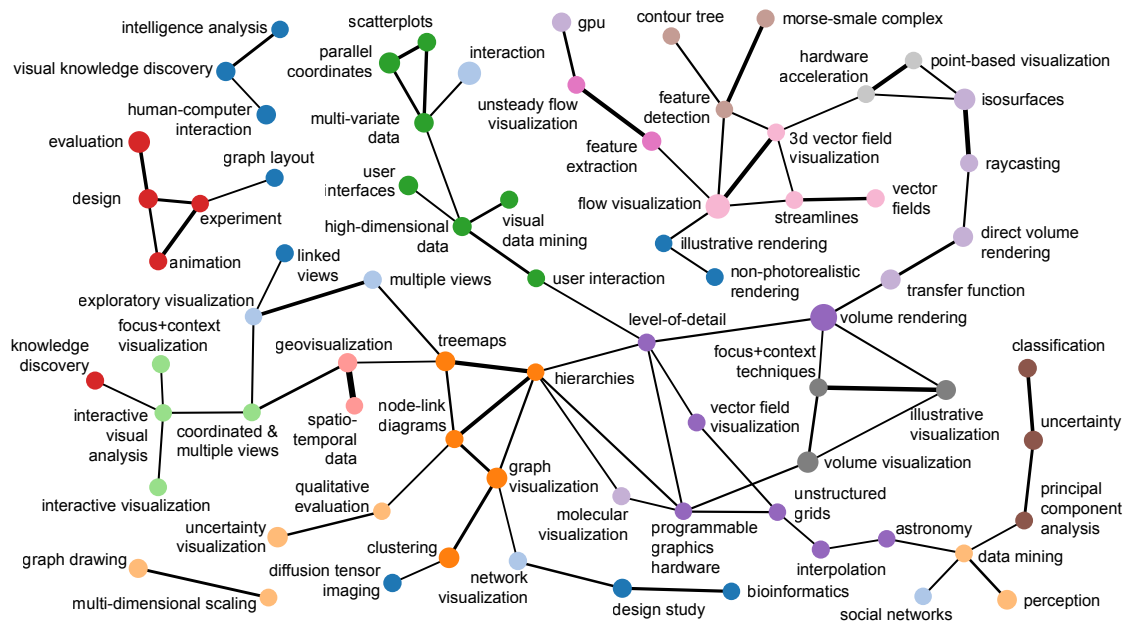


Figure 4: Keyword map from hierarchical clustering of author-assigned keywords; showing only connected nodes with correlation strength ≥ 0.13 . Circle areas correspond to number of occurrences of the keyword in the dataset, link width corresponds to correlation strength, and color distinguished node clusters. We chose a link threshold based on visual inspection of the resulting graph to generate a manageable and readable layout. Isolated nodes were removed from the image.

4.1.2 Expert-coded Keywords

CLUSTER ANALYSIS: We created 16 clusters from the 101 expert-coded keywords (after removing 53 whose frequency was < 10). Again, we chose the cluster number based on a manual inspection of the hierarchical clustering result. Table 5 gives an overview of the created clusters (E1–E16) as well as the same cluster metrics as in Table 4.

It is important to note that many of the keywords contained in these 101 expert keyword clusters capture topics that were not covered in the cluster analysis of the author keywords. Specifically, 46 of

Table 5: Cluster result for expert-coded keywords. Keywords are sorted by frequency with the two most frequent keywords highlighted in bold.

ID	keywords (InfoVis, Vis/SciVis, VAST; 2004–2013)	N	#	cw-#	centr.	dens.
E1	graphs, visualiz. techniques & tools , clustering, hierarchies, layout, charts/diagrams/plots, networks, scatterplots, matrices	9	34.0	3.9	0.355	0.063
E2	evaluation, perception , design & guidelines, visual encoding, color, human factors/hci, art & aesthetics, provenance & history	8	28.5	3.2	0.377	0.068
E3	geography/geospatial/cartography, motion , spatiotemporal, maps	4	20.5	4.0	0.276	0.152
E4	displays, toolkits/systems/environments , interfaces, theory/models/methods, knowledge discovery, cognition, metrics, video & multimedia, virtual environments	9	29.0	1.4	0.394	0.040
E5	interaction, analysis process , multidimensional/multivariate, multiple & linked views, queries & search	5	83.0	11.5	0.350	0.081
E6	data & data management, dimensionality reduction , social data/media/vis, databases & data mining, collaboration, programming, internet/web, scale & multiscale, curves & curvature	9	25.0	1.5	0.413	0.038
E7	applications, text & documents , physics, patterns & pattern detection, semantics/semiotics, storytelling	6	20.5	2.2	0.317	0.052
E8	machine learning & statistics, images & image/signal processing , uncertainty, segmentation/classification, simulation	5	34.0	3.0	0.295	0.067
E9	interpolation, sampling , grids & lattices	3	14.0	4.0	0.267	0.261
E10	surfaces, numerical methods/math , isosurfaces, meshes & grids, geometric modeling, point based data, raytracing/raycasting	7	32.0	4.0	0.384	0.097
E11	flow, topology , vectors, features & attributes, 3d visualization, streamlines/pathlines/streaklines	6	42.0	6.6	0.380	0.122
E12	volume rendering, hardware & computation , rendering & illumination, textures	4	68.5	14.8	0.387	0.159
E13	large scale data, cameras & views , resolution/multiresolution, astronomy/astrophysics, computer networks, particles, information theory, level of detail	8	15.0	1.1	0.403	0.048
E14	medical, tensors , tractography, glyphs, climate/environment, biomedical	6	20.0	1.9	0.358	0.058
E15	design studies & case studies, biology , molecules, bioinformatics	4	19.0	2.2	0.302	0.107
E16	time, focus+context , illustrative visualiz., comparison/comp. visualiz., spatial, events/trends/outliers, animation, comp. graph.	8	22.0	1.6	0.391	0.039

the expert keywords in Table 5 do not contain any of the 101 most frequent author-assigned keywords in Table 4 that we used for author keyword clustering. In other words, these 46 expert keywords capture topics in which none of the contained author keywords occurred more than 5 \times . Out of these 46, the following ones were among the top two most common themes in one of the expert keyword clusters: *motion* (E3), *data and data management* (E6), *machine learning & statistics* (E8), *images and image/signal processing* (E8), *surfaces* (E10), *numerical methods/mathematics* (E10), *large-scale data* (E13), *cameras & views* (E13), *tensors* (E14), *biology* (E15). The top three of these keywords occurred more than 50 \times in our corpus and would not have emerged as important without the expert coding.

NETWORK ANALYSIS: The strategic diagram for expert keywords can be seen in Fig. 3(b). From the following discussion we exclude clusters E1, E2, and E8, as their centrality and density scores were very close to median centrality or density and, thus, quadrant membership was not strong for these clusters.

Three central and developed motor themes emerged in Quadrant I (top right): E10 (*surfaces, numerical methods/mathematics*), E11 (*flow, topology*), E12 (*volume rendering, hardware and computation*). All three themes are traditional topics of IEEE Vis/SciVis which also has the longest history of the three IEEE VIS conferences. Hence it is perhaps not surprising that these topics emerged as central and dense. *Flow visualization* and *vector fields* also proved to be a motor theme in the author keywords, further emphasizing E11’s place as motor theme.

As undeveloped yet still central themes in Quadrant II (bottom right), clusters E4 (*displays, toolkits/systems/environments*), E6 (*data and data management, dimensionality reduction*), E13 (*large scale data, cameras & views*), and E16 (*time, focus+context*) emerged most clearly.




Figure 5: Keyword map from hierarchical clustering of the expert-coded keywords; showing only connected nodes with correlation strength ≥ 0.11 .

Themes in Quadrant III (top left; peripheral to the general research network) were E3 (*geography/geospatial/cartography, motion*), E5 (*interaction, analysis process*), E9 (*interpolation, sampling*), and E15 (*design studies & case studies, biology*). E9 is an extreme case as it exhibits the highest density and lowest centrality score in the network and is also quite small. This theme is thus the most developed but isolated of the themes in this Quadrant. E5, in contrast, was very close to the median for both centrality and density and is thus unspecific.

Among the emerging or declining themes in Quadrant IV (lower left) were clusters E7 (*applications, text and documents*) and E14 (*medical, tensors*). Interestingly, the graph visualization-related cluster A7 for the author keywords was characterized as a motor theme during the author-keyword analysis but exhibited much lower centrality and lower density for the expert-coded cluster E1. The network diagram in Fig. 5 gives a visual overview of how words and clusters correlated and how central certain keywords are placed within the general network.

4.1.3 PCS Taxonomy Keywords

Before discussing the analysis of PCS keywords, it is important to note that their usage in practice is different than for author keywords. Author keywords are generally used to label a paper and to explain in a few words what research topics it covers—for example, with the intention to make the paper discoverable in searches. While the PCS keywords can be used in the same way, the choices are limited and it is not always possible to find the one keyword that exactly expresses a contribution. For the PCS data we analyzed, authors were instructed to choose only a limited subset of PCS keywords (e. g., a primary and up to four secondary keywords). Authors, thus, had to balance their choices among all possible keywords

that describe a paper contribution. A fundamental difference to author keywords, however, is that the PCS taxonomy is used not only by the paper authors but also by the reviewers to describe their expertise levels for each topic. The selection of PCS keywords for papers and by reviewers then informs how reviewers are assigned to papers. Since this process is known to many paper authors, PCS keywords are often carefully selected to rank more highly for reviewers with a certain expertise.

Given these differences in use, PCS keywords may not as accurately reflect the topic of a paper as the author keywords do. Nevertheless, by conducting the same analysis on this keyword set as we did before for author keywords, we can speculate on differences in use as well as on topics that may be missing in each type of keyword set.

Overall, we observed that, similar to author and expert keywords, a small number of PCS keywords appeared very frequently while a larger number of keywords did not. The five least frequently used PCS keywords for the entire period of the dataset were *volume graphics hardware* (8×), *haptics for visualization* (7×), *embodied/enactive cognition* (6×), *sonification* (5×), and *special purpose hardware* (3×).

CLUSTER ANALYSIS: The result of the analysis of the PCS taxonomy is reported in the same way as that of the other two keyword sets. Table 6 has the overview of the generated clusters, their keywords, and metrics. Fig. 6 shows a filtered network graph generated from the keyword correlations, while Fig. 3(c) depicts the strategic diagram.

The cluster analysis includes several interesting observations. First—and perhaps not too surprising—most clusters include keywords that belong to different higher-level taxonomy keywords. For example, *flow visualization* in P1 is considered an application according to the PCS taxonomy, while *vector field data* belongs to the category *spatial data and techniques*. Yet, many clusters include keywords from a predominant subset of the higher level of the PCS taxonomy: P2 (*spatial data and techniques*), P4 (*non-spatial data and techniques*), P6 (*display and interaction technology*), P7 (*visual analytics applications*), P11 (*evaluation*), and P13 (*data handling, processing and analysis*).

NETWORK ANALYSIS: Next, we analyzed the strategic diagram in Fig. 3(c). From the following discussion we exclude P2, P3, P7, and P11, as their density or centrality scores were too close to the media to derive conclusive quadrant memberships.

The most evident motor theme in Quadrant I (top right) was P13 (*data transformation and representation, data filtering*). P14 (*focus+context techniques, zooming and navigation techniques*) also fell within Quadrant I but was relatively close to median centrality. The network diagram in Fig. 6 illustrates the high density and centrality of P13 and its keywords (colored in light orange) well.

The themes in Quadrant II (bottom right; undeveloped topics that are central to the research network because they are well connected) are P5 (*graph/network data, visual knowledge discovery*), P10 (*illustrative visualization, animation*), P8 (*scalability issues, streaming data*), and P15 (*user interfaces, visualization system & toolkit design*).

Quadrant III (top left; developed but isolated themes) contained, in particular, P1 (*flow visualization, vector field data*)—the theme with highest density—and P4 (*multidimensional data, high-dimensional data*)—the theme with the lowest centrality. The other two themes in Quadrant III are P6 (*collaborative and distributed visualization, large and high-res displays*) and P12 (*visualization models, taxonomies*).

Finally, Quadrant IV (bottom left; emerging/declining themes—undeveloped and peripheral to the network) contains P9 (*time-varying data, geographic/geospatial visualization*).

4.1.4 Relationship between Expert and PCS Keywords

As both expert and PCS keywords present a higher-level assessment of the visualization topics, we were interested in analyzing their relationship. Yet, we found it difficult to generate an exact match from PCS to expert keywords and vice versa—in part because PCS and expert keywords were not always on the same levels. For example, many specific application domains in the expert keywords could not be matched

Table 6: Cluster result for PCS taxonomy keywords. Keywords are sorted by frequency with the two most frequent keywords highlighted in bold.

ID	keywords (EuroVis, PacificVis, InfoVis, Vis/SciVis, VAST; 2008–2013)	N	#	cw-#	centr.	dens.
P1	flow visualization, vector field data , feature detection and tracking, topology-based techniques	4	210.0	26.3	0.328	0.216
P2	volume rendering, gpus and multi-core architectures , scalar field data, extraction of surfaces (isosurfaces, material boundaries), point-based data, irregular and unstructured grids, molecular visualization, pde's for visualization	8	143.0	11.5	0.404	0.079
P3	biomedical and medical visualization, visualization in physical sciences and engineering , geometry-based techniques, multi-field, multi-modal and multi-variate data, uncertainty visualiz., mathematical foundations for visualization, glyph-based techniques, tensor field data, volume modeling, visualiz. in mathematics, data registration, sonification	12	124.5	4.1	0.484	0.028
P4	multidimen. data, high-dimen. data , parallel coordinates, dimensionality reduction, statistical graphics, tabular data	6	125.0	11.9	0.314	0.095
P5	graph/network data, visual knowledge discovery , coordinated & multiple views, time series data, text & document data, data clustering, visual knowledge represent., hierarchy data, visualiz. in social & information sciences, bioinformatics visualiz., visual analysis models, integrating spatial & non-spatial data visualiz., machine learning, pixel-oriented techniques, software visualiz., hypothesis testing, visual evidence, hypothesis forming, business & finance visualiz., data segmentation, data fusion & integration, visualiz. in the humanities	21	119.0	8.1	0.704	0.032
P6	collaborative & distributed visualization, large & high-res displays , immersive & virtual environments, stereo displays	4	74.5	3.8	0.376	0.155
P7	intelligence analysis, situational awareness , emergency/disaster management, knowledge externalization, network security and intrusion, time critical applications, privacy and security, distributed cognition	8	29.5	1.4	0.594	0.080
P8	scalability issues, streaming data , data warehousing, database visualization and data mining, parallel systems, distributed systems and grid environments, petascale techniques, cpu and gpu clusters	7	47.0	2.5	0.511	0.073
P9	time-varying data, geographic/geospatial visualization , visualization in earth, space, and environmental sciences, multiresolution techniques, view-dependent visualization, compression techniques, terrain visualization, sensor networks	8	98.5	4.8	0.462	0.048
P10	illustrative visualization, animation , visualization for the masses, aesthetics in visualization, multimedia (image/video/music) visualization, human factors, color perception, mobile and ubiquitous visualization, scene perception, visualization in education, motion perception, texture perception, attention and blindness	13	51.0	1.7	0.617	0.035
P11	quantitative eval., qualitative eval. , laboratory studies, metrics & benchmarks, percept. cognition, cognitive & percept. skill	6	69.5	10.0	0.489	0.116
P12	visualization models, taxonomies , cognition theory, perception theory, embodied / enactive cognition	5	42.0	5.8	0.448	0.117
P13	data transformation and representation, data filtering , data aggregation, data acquisition and management, data smoothing, data editing, data cleaning, volume graphics hardware	8	36.0	5.1	0.714	0.122
P14	focus+context techniques, zooming and navigation techniques , manipulation and deformation, multimodal input devices, haptics for visualization, special purpose hardware	6	44.0	3.6	0.546	0.087
P15	user interfaces, visualiz. system & toolkit design , interact. design, human-computer interact., visual design, design studies, usability studies, design methodologies, task and requirements analysis, presentation/production/dissemination, field studies	11	159.0	11.1	0.703	0.069

well because we had matched many applications into a generic *applications* group due to their relatively infrequent use. On the other hand, in the expert coding we had explicitly labeled specific application domains that only existed as relatively large groups in the PCS taxonomy (e. g., expert keywords *physics, material science, automation, traffic* vs. *visualization in physical sciences and engineering*; or expert keywords *astronomy/astrophysics, climate/environment, earth sciences* vs. *visualization in earth, space, and environmental sciences*).

We also found the higher-level PCS category *data handling, processing and analysis* to contain very




Figure 6: Keyword map from hierarchical clustering for PCS taxonomy keywords; showing only connected nodes with correlation strength ≥ 0.135 .

fine-grained keywords which, in our expert coding, were simply mapped to *data and data management*. Similarly, we had merged all author keywords related to *scalability issues* and *petascale techniques* into one expert-keyword: *large-scale data*. On the other hand, we also found PCS categories that relate to several expert keywords and are, thus, rather broad. For example the keyword *visual design* relates to the expert keywords *layout*, *visual encoding*, and *design & guidelines*.

Despite these differences in categorization, we found a number of expert keywords that we could not match well in the PCS taxonomy but did appear in our dataset $20\times$ or more. These keywords are: *rendering & illumination*, *images and image/signal processing*, *meshes and grids*, *charts/diagrams/plots*, *queries and search*, *comparison/comparative visualization*, *interpolation*, *motion*, *programming*, and *streamlines/pathlines/streaklines*.

4.2 Analysis of Individual Keywords

In this section, we analyze individual keywords in more detail across all three of our data sources.

4.2.1 Common Keywords

The top three most frequently occurring keywords in the author keyword list were community-related keywords: *information visualization*, *visual analytics*, and *visualization* (in this order). These are all keywords that describe a research (sub)community and it is interesting to see that the term scientific visualization does not occur. The terms *information visualization* and *visual analytics* were predominantly present on papers accepted to the respective conferences. We hypothesize that both IEEE VAST and IEEE InfoVis have been the smaller of the three conferences and in the process of becoming more established,

authors may have chosen to tag their papers with these terms to strengthen associations to their respective communities. This hypothesis is also supported by the observation that *information visualization* was the most frequent keyword for the period of 2004–2008, while *visual analytics* was the most frequent keyword for period of 2009–2013 during which IEEE VAST became an established venue.

Next, we analyzed the top individual keywords per visualization conference to get an understanding of main topics and concerns in these subcommunities. We excluded the community-related keywords from this analysis. For author-assigned keywords, the top three for each IEEE VIS conference were:

VAST: *sensemaking* (13), *visual knowledge discovery* (8), *geovisualization* (7)

InfoVis: *interaction* (25), *graph visualization* (19), *evaluation* (17)

Vis/SciVis: *volume rendering* (72), *flow visualization* (48), *isosurfaces* (25)

Among these, no common keywords exist between the conferences. When expanding the analysis to the top 10, VAST and InfoVis share the keyword *interaction* but no other overlaps emerged. This may indicate that the three conferences do focus on different areas of visualization—or use different terminology for related concepts.

In order to understand whether terminology plays a role, we looked at the top three expert keywords. Here, the keyword *interaction* emerges for both IEEE InfoVis and VAST but remains the only shared keyword:

VAST: *analysis process* (53), *interaction* (31), *applications* (24)

InfoVis: *graphs* (73), *interaction* (69), *evaluation* (53)

Vis/SciVis: *volume rendering* (148), *hardware and computation* (89), *flow* (70)

When expanding the analysis to the top 10 keywords, *application* emerges as a keyword common to all three conferences. InfoVis and VAST also share the keyword *evaluation* (ranked 14 in Vis/SciVis) while VAST and Vis/SciVis share the keyword *time* (ranked 12 in InfoVis). Among the top 100 keywords, 17 were used only on Vis/SciVis papers (ignoring a up to one occurrence at each of the other two conferences). The top 3 of these keywords were: *volume rendering*, *isosurfaces*, and *geometric modeling*. We can conclude that all three sub-communities share some joint concerns, e. g. on applications, evaluation, and temporal data but that they do have their respective foci. However, there appears to be more overlap between VAST and InfoVis than we observe between Vis/SciVis and the other two conferences according to the expert keywords.

The top three PCS keywords *only* for the IEEE VisWeek/VIS were:

VAST: *visual knowledge discovery* (160), *text and document data* (78), *coordinated and multiple views* (74)

InfoVis: *graph/network data* (197), *visual knowledge discovery* (99), *user interfaces* (98)

Vis/SciVis: *volume rendering* (235), *biomedical and medical visualization* (193), *flow visualization* (154)

It can be seen that the top three keywords for the three datasets relate to each other in that *graphs* and *volume rendering* remain the top keywords for InfoVis and Vis/SciVis, respectively, and that *visual knowledge discovery* was in the top 3 for VAST in the author keywords and also covered under *analysis process* for the expert keywords. When looking at the top 100 PCS keywords, 6 emerged that were only used by submissions to Vis/SciVis (ignoring up to one occurrence at each of the other two conferences). These keywords all belong to the higher-level category of *spatial data and techniques: vector field data, scalar field data, extraction of surfaces, tensor field data, irregular and unstructured grids, and volume modeling*. Conversely, there were two keywords never or only once used for submissions to Vis/SciVis: *business and finance visualization* and *tabular data*.

4.2.2 Rising and Declining Keywords

Next, we were interested in historic trends for individual keywords. We used Tableau to calculate linear trend lines for the 15 most frequently used keywords in the author, expert, and PCS datasets. Table 7 gives an overview of this data by dataset, ordered by slope from increasing to decreasing. Significant results are highlighted in bold (note that not all significant trends are shown due to the selection by frequency).

RISING KEYWORDS: Across all three IEEE VIS conferences, the use of the term *interaction* is significantly increasing according to author-chosen keywords. If one only considers the expert keywords, the terms *interaction*, *evaluation*, *multidimensional/multivariate*, and *machine learning & statistics* are on the rise. The analysis of the PCS keywords yielded 8 keywords that are significantly increasing: *coordinate and multiple views*, *visual knowledge discovery*, *time-series data*, *data transformation and representation*, *high-dimensional data*, *interaction design*, *multidimensional data*, and *graph/network data*. There is overlap between these increasingly popular topics. We attribute the strongest overall increasing trends to: *interaction*, *evaluation*, *multi-dimensional/high-dimensional/multi-variate data*.

DECLINING KEYWORDS: According to author-chosen keywords, both *volume visualization* and *flow visualization* have been significantly decreasing in frequency of occurrence over the past 10 years. The expert-coded keywords confirm this and add *hardware and computation* as a declining topic. Interestingly, no significantly declining topics were found for the PCS taxonomy keywords (but *GPUs* and *multi-core architectures*—not shown in the table—trended that way, $p = 0.07$). This may be due to the fact that we only had data for six years for PCS but 10 years for the author/expert keyword sets.

For the author/expert keywords, it is interesting to note that two very core and frequent keywords for the IEEE Vis/SciVis conference are in significant decline. This could perhaps be due to the fact the many fundamental research questions have been tackled and that researchers are now using more specific or other keywords.

4.3 Limitations

While our analysis has revealed a wealth of information, the study results have to be read in light of several analysis limitations. One obvious limitation is, of course, that we only analyzed a subset of publications from the visualization domain. To determine this subset, however, we followed advice from Bradford's law [3] for selecting our data sources. This law states that a small core of publications will account for as much as 90% of the literature in terms of citations received—trying to achieve 100% will add publications at an exponential rate. Thus, it is unreasonable to attempt to cover the whole literature for a field. Given the size and importance of IEEE VisWeek/VIS, we focused on a ten-year dataset from this conference (plus the additional data on the use of the PCS taxonomy). This analysis enabled us to get a rich overview of the field. Yet, compared to several past keyword analysis studies (e. g., [6, 25, 24]), the visualization field is still young and the overall number of keywords was comparably low, in particular for the author-assigned keywords. The low number of overall keywords and the vast difference in number of papers accepted to the IEEE VisWeek/VIS conferences is also one of the main reasons we did not study the difference between IEEE VAST, InfoVis, and Vis/SciVis in depth but looked at the whole field together. By expanding the dataset back to 1995 (the first year of InfoVis) a comparative analysis may be more meaningful. Yet, another peculiarity of the field may further impact such a comparison: for a long period of time it was possible to send information visualization papers to both the InfoVis and Vis/SciVis conference. The Vis/SciVis conference even included information visualization sessions. Thus, the borders between the conferences were not always as clear as their current names may suggest.

For the PCS data we discussed a major limitation earlier that pertains to the different use of these keywords compared to author-assigned keywords. In addition, we also found that several older papers included ACM classifiers and it is possible that authors back then only selected author keywords in addition to these classifiers and did not provide duplicates. As we did not want to speculate on what may have happened, we collected the author keywords as present on each paper.

Table 7: Historical trends for 15 most frequently used keywords for each of the author, expert, and PCS datasets. Significant trends highlighted.

keyword	author (InfoVis, Vis/SciVis, VAST, 2004–2013)		expert (InfoVis, Vis/SciVis, VAST; 2004–2013)		PCS (EuroVis, PacificVis, InfoVis, Vis/SciVis, VAST, 2008–2013)	
	#	SE p-val.	#	SE p-val.	#	SE p-val.
interaction	38	0.59 0.16 0.006	152	1.37 0.25 0.001	323	10.71 3.24 0.030
visual analytics	86	0.58 0.52 0.299	105	1.32 0.42 0.014	414	9.20 1.79 0.007
uncertainty visualization	19	0.33 0.16 0.086	83	0.87 0.34 0.034	293	8.89 2.40 0.021
evaluation	26	0.30 0.17 0.111	86	0.58 0.52 0.299	252	7.77 2.46 0.034
user study	22	0.10 0.12 0.416	113	0.36 0.32 0.296	278	7.54 1.51 0.008
visualization	81	-0.01 0.30 0.984	72	0.34 0.26 0.235	245	6.66 2.28 0.043
parallel coordinates	23	-0.08 0.19 0.680	91	0.26 0.36 0.494	251	5.69 2.23 0.063
clustering	20	-0.12 0.14 0.445	64	0.24 0.31 0.450	291	4.49 2.65 0.166
graph visualization	21	-0.13 0.12 0.321	81	-0.01 0.30 0.984	379	3.80 2.06 0.139
focus+context	18	-0.17 0.12 0.195	85	-0.13 0.35 0.723	306	3.49 1.00 0.025
information visualization	95	-0.24 0.35 0.521	95	-0.24 0.35 0.521	460	3.31 0.93 0.024
isosurfaces	25	-0.42 0.23 0.103	103	-0.27 0.30 0.392	446	2.40 3.88 0.569
volume visualization	24	-0.52 0.20 0.042	74	-0.70 0.24 0.019	274	2.34 3.06 0.487
volume rendering	72	-0.72 0.37 0.091	149	-1.42 0.46 0.014	334	2.29 1.39 0.176
flow visualization	48	-0.82 0.26 0.016	98	-1.59 0.43 0.006	267	1.86 1.33 0.234

The expert coding process that led to the expert keywords also, of course, includes an inherent subjective coding bias. While our team consisted of 5 experts with varying backgrounds in VAST, InfoVis, Vis/SciVis, and HCI, it is entirely possible that with other coders other clusters and cluster names may have emerged. Finally, this keyword analysis was conducted by experts in the field of visualization and not professional social scientists. Our own experience and views of our field have certainly tainted our interpretations of the data—as is common with any kind of qualitative analysis [12].

5 keyvis.org: A Keyword Search Tool

To make our data accessible for others, we created a webpage that makes author and expert keywords and related papers search- and browsable: <http://www.keyvis.org/>. Visitors can search all 2629 unique author-assigned keywords, find out which keywords co-occurred how frequently, which manual expert clusters they belong to, and the actual research papers they appear on.

Our main goal was to generate an easy-to-use, lightweight interface to our keyword data in order to: (a) support visualization researchers in making more informed decisions when picking keywords for their papers, and (b) give a new lens on identifying relevant related work. We have used the site ourselves for choosing keywords and finding related work that we were not aware of before. We hope that others will find it similarly useful. In the long run, our goal is to maintain the website as a platform for visualization keyword access and analysis.

6 Discussion and How to Move Forward

The analysis of the keyword data has revealed several major themes and declining and rising keywords in visualization. Based on the strategic diagrams, we have identified, in particular, motor themes of the field. While we have collected a large amount of author keyword data, invested heavily in clustering related author keywords, and compared this data to a standardized taxonomy, our analysis is only a first step in the direction of two larger research goals:

Creating a common vocabulary: In particular the analysis of raw author keywords has revealed that authors choose many variants of similar keywords based on: singulars and plurals (e.g., glyphs and glyph), abbreviated versions (e.g., DTI vs. diffusion tensor imaging; InfoVis vs. information visualization), spelling (multidimensional vs. multi-dimensional), or specificity (e.g., multi-dimensional vs. multi-dimensional data vs. multi-dimensional visualization). Such a diversity of terms may be a reflection of the diversity of influences on the visualization field—but is not helpful, in particular when one wants to search for keywords or—like us—gain an overview of the research themes of a community. We hope that keyvis.org will help paper authors find common terms and reflect on their keyword usage before submitting a camera-ready paper. In addition, one can think about the problem of creating a common vocabulary for visualization more broadly. By identifying key terms and providing clear definitions, sub-communities in visualization may be able to communicate more clearly about similar approaches and, this, in turn can help to also collaborate more effectively with people outside the community. Finally, a common vocabulary can also facilitate to more easily understand emerging and declining research trends within the field.

Establish a comprehensive taxonomy of visualization research: One of the goals that we had initially set out to accomplish is not yet achieved. Perhaps the holy grail of a keyword analysis is to amount into a taxonomy of the analyzed field, in our case visualization research. This could serve two purposes. One the one hand, a taxonomy will help to better communicate “what is visualization” to other disciplines, i.e., researchers and practitioners not part of the VIS community. On the other hand, we are hoping to be able to facilitate the crucial step of matching reviewers with papers and grants such that the peer review process improves and new contributions is seen in the right context.

Yet, how to exactly create and **maintain** a comprehensive taxonomy of visualization keywords is still an open question. As can easily be seen from an inspection of both the top IEEE and INSPEC terms in Table 2, broader taxonomies are not very successful in capturing the diversity of the visualization field. The PCS taxonomy, in turn, has been developed by a team of dedicated experts with the goals to improve the reviewing system for visualization papers. Yet, what is the right process for maintaining and changing the taxonomy? The visualization field is evolving and, thus, a visualization taxonomy should be regularly updated. Should an analysis such as ours be used to find trends (and keywords representing these trends) that have not been captured? Would it be possible to automate our process without requiring experts to clean and code the data? Should certain keywords be split or merged as sub-areas increase or decrease in popularity? Does it make sense to keep keywords in the PCS taxonomy that are rarely used—or should the taxonomy provide a broad view of the field in order to capture its topic diversity? Even when one has answers to these questions, how would one choose the right level of granularity for keywords? For example, the keyword *evaluation* emerged as a significantly increasing topic for our expert-clustered keywords but did not for the PCS keywords. The reason being that the topic is split into seven sub-topics in PCS, out of which none showed an increasing trend taken alone. Finally, should there be separate taxonomies: one for visualization as a whole (e. g., to conduct analyses on emerging and declining topics, topic coverage across subcommunities, etc.) and one for the submission process for academic research? The former taxonomy could be large and evolving while the latter would have to be reduced in size and remain stable across conferences for given time periods to remain manageable for papers chairs, editors, reviewers, and authors.

In summary, the 25th anniversary of the IEEE Vis/SciVis conference and thus of the whole IEEE VisWeek/VIS conference series is a good moment to take a look back at the diverse and evolving history of our field. We managed to provide only a glimpse at the available data in this paper but hope that our discussion will spark others to discuss and consider the importance of keywords, common vocabularies, and taxonomies for the field as a whole.

Acknowledgments

We would like to thank Vassilis Kostakos for sharing details about the keyword analysis process for CHI keywords. We also thank Nenad Kircanski and Johanna Schlereth for their help with creating keyvis.org.

References

- [1] C. C. Aggarwal and C. K. Reddy, eds. *Data Clustering: Algorithms and Applications*. Chapman & Hall/CRC Press, Boca Raton, FL, USA, 2014.
- [2] E. Bertini, A. Tatu, and D. Keim. Quality metrics in high-dimensional data visualization: An overview and systematization. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2203–2212, Dec. 2011. doi> 10.1109/TVCG.2011.229
- [3] S. C. Bradford. Sources of information on specific subjects. *Journal of Information Science*, 10(4):173–180, Oct. 1985. doi> 10.1177/016555158501000406
- [4] J. R. Bray and J. T. Curtis. An ordination of the upland forest communities of southern Wisconsin. *Ecological Monographs*, 27(4):325–349, Oct. 1957. doi> 10.2307/1942268
- [5] M. Callon, J. Courtial, and F. Laville. Co-word analysis as a tool for describing the network of interactions between basic and technological research: The case of polymer chemistry. *Scientometrics*, 22(1):155–205, 1991. doi> 10.1007/BF02019280

- [6] M. Callon, J. Law, and A. Rip, eds. *Mapping the Dynamics of Science and Technology*. Palgrave Macmillan, Hampshire, UK, 1986.
- [7] S. K. Card, J. D. Mackinlay, and B. Shneiderman. *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufmann, San Francisco, 1999.
- [8] E. H. Chi. A taxonomy of visualization techniques using the data state reference model. In *Proc. InfoVis*, pp. 69–75. IEEE Computer Society, Los Alamitos, 2000. doi> 10.1109/INFVIS.2000.885092
- [9] J. Chuang, S. Gupta, C. D. Manning, and J. Heer. Topic model diagnostics: Assessing domain relevance via topical alignment. *Journal of Machine Learning Research: Workshop and Conference Proceedings*, 28(3):612–620, 2013.
- [10] N. Coulter, I. Monarch, and S. Konda. Software engineering as seen through its research literature: A study in co-word analysis. *Journal of the American Society for Information Science*, 49(13):1206–1223, Dec. 1998. doi> 10.1002/(SICI)1097-4571(1998)49:13<1206::AID-ASI7>3.0.CO;2-F
- [11] R. J. Crouser and R. Chang. An affordance-based framework for human computation and human-computer collaboration. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2859–2868, Dec. 2012. doi> 10.1109/TVCG.2012.195
- [12] N. K. Denzin and Y. S. Lincoln, eds. *Handbook of Qualitative Research*. Sage Publications, Inc., Thousand Oaks, London, New Delhi, 3rd ed., 2005.
- [13] E. Geisler. *The Metrics of Science and Technology*. Quorum Books, Westport, CT, USA, 2000.
- [14] C. Görg, Z. Liu, J. Kihm, J. Choo, H. Park, and J. Stasko. Combining computational analyses and interactive visualization for document exploration and sensemaking in Jigsaw. *IEEE Transactions on Visualization and Computer Graphics*, 19(10):1646–1663, Oct. 2013. doi> 10.1109/TVCG.2012.324
- [15] Q. He. Knowledge discovery through co-word analysis. *Library Trends*, 48(1):133–159, Summer 1999. doi> 2142/8267
- [16] J. Heer and B. Shneiderman. Interactive dynamics for visual analysis. *Communications of the ACM*, 55(4):45–54, Apr. 2012. doi> 10.1145/2133806.2133821
- [17] A. Hoonlor, B. K. Szymanski, and M. J. Zaki. Trends in computer science research. *Communications of the ACM*, 56(10):74–83, Oct. 2013. doi> 10.1145/2500892
- [18] C.-P. Hu, J.-M. Hu, S.-L. Deng, and Y. Liu. A co-word analysis of library and information science in China. *Scientometrics*, 97(2):369–382, Nov. 2013. doi> 10.1007/s11192-013-1076-7
- [19] P. Isenberg and T. Isenberg. Visualization on Interactive Surfaces: A Research Overview. *i-com*, 12(3):10–17, Nov. 2013. doi> 10.1524/icom.2013.0020
- [20] T. Isenberg, P. Isenberg, J. Chen, M. Sedlmair, and T. Möller. A systematic review on the practice of evaluating visualization. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2818–2827, Dec. 2013. doi> 10.1109/TVCG.2013.126
- [21] H. Lam. A framework of interaction costs in information visualization. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1149–1156, Nov./Dec. 2008. doi> 10.1109/TVCG.2008.109

- [22] H. Lam, E. Bertini, P. Isenberg, C. Plaisant, and S. Carpendale. Empirical studies in information visualization: Seven scenarios. *IEEE Transactions on Visualization and Computer Graphics*, 18(9):1520–1536, Sept. 2012. doi> 10.1109/TVCG.2011.279
- [23] J. Law, S. Bauin, J.-P. Courtial, and J. Whittaker. Policy and the mapping of scientific change: A co-word analysis of research into environmental acidification. *Scientometrics*, 14(3–4):251–264, 1988. doi> 10.1007/BF02020078
- [24] G.-Y. Liu, J.-M. Hu, and H.-L. Wang. A co-word analysis of digital library field in China. *Scientometrics*, 91(1):203–217, Apr. 2012. doi> 10.1007/s11192-011-0586-4
- [25] Y. Liu, J. Goncalves, D. Ferreira, B. Xiao, S. Hosio, and V. Kostakos. CHI 1994–2013: Mapping two decades of intellectual progress through co-word analysis. In *Proc. CHI*, pp. 3553–3562. ACM, 2014. doi> 10.1145/2556288.2556969
- [26] A. D. Ritzhaupt, M. Stewart, P. Smith, and A. E. Barron. An investigation of distance education in North American research literature using co-word analysis. *International Review of Research in Open and Distance Learning*, 11(1):37–60, 2010.
- [27] H.-J. Schulz. Treevis.net: A tree visualization reference. *IEEE Computer Graphics and Applications*, 31(6):11–15, Nov. 2011. doi> 10.1109/MCG.2011.103
- [28] M. Sedlmair, M. Meyer, and T. Munzner. Design study methodology: Reflections from the trenches and the stacks. *IEEE Transactions on Visualization and Computer Graphics (Proc. InfoVis)*, 18(12):2431–2440, Dec. 2012. doi> 10.1109/TVCG.2012.213
- [29] B. Shneiderman. The eyes have it: A task by data type taxonomy for information visualizations. In *Proc. Visual Languages*, pp. 336–343. IEEE Computer Society, Los Alamitos, 1996. doi> 10.1109/VL.1996.545307
- [30] J. Stasko, J. Choo, Y. Han, M. Hu, H. Pileggi, R. Sadanaand, and C. D. Stolper. CiteVis: Exploring conference paper citation data visually. In *Posters of IEEE InfoVis*, 2013.
- [31] The Institute of Electrical and Electronics Engineers (IEEE). 1998 IEEE Approved Indexing Keyword List, 1998.
- [32] The Institute of Electrical and Electronics Engineers (IEEE). 2013 IEEE Taxonomy, version 1.0, 2013.
- [33] The Institute of Electrical and Electronics Engineers (IEEE). Thesaurus FAQs, 2014. Web site: <http://www.computer.org/portal/web/tandc/Thesaurus-FAQ>, visited March 2014.
- [34] J. J. Thomas and K. A. Cook, eds. *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. IEEE Computer Society, Los Alamitos, 2005.
- [35] M. Tory and T. Möller. Human factors in visualization research. *IEEE Transactions on Visualization and Computer Graphics*, 10(1):72–84, Jan./Feb. 2004. doi> 10.1109/TVCG.2004.1260759
- [36] K. Voegelé. Annotated bibliography of the visualization conference proceedings. In *Proc. IEEE Visualization*, p. xxii. IEEE Computer Society, Los Alamitos, 1995. doi> 10.1109/VIS.1995.10003
- [37] J. H. Ward. Hierarchical grouping to optimize an objective function. *Journal of the American Statistical Association*, 58(301):236–244, Mar. 1963. doi> 10.1080/01621459.1963.10500845

-
- [38] J. Whittaker. Creativity and conformity in science: Titles, keywords and co-word analysis. *Social Studies of Science*, 19(3):473–496, Aug. 1989. doi> 10.1177/030631289019003004
- [39] J. S. Yi, Y. a. Kang, J. Stasko, and J. Jacko. Toward a deeper understanding of the role of interaction in information visualization. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1224–1231, Nov. 2007. doi> 10.1109/TVCG.2007.70515



**RESEARCH CENTRE
SACLAY – ÎLE-DE-FRANCE**

1 rue Honoré d'Estienne d'Orves
Bâtiment Alan Turing
Campus de l'École Polytechnique
91120 Palaiseau

Publisher
Inria
Domaine de Voluceau - Rocquencourt
BP 105 - 78153 Le Chesnay Cedex
inria.fr

ISSN 0249-6399