Network Analysis of the American Geophysical Union's Fall Meetings

Tom Narock[1], Sarah Hasnain[1], and Ronie Stephan[1]

1. Notre Dame of Maryland University, Baltimore, MD, USA
Author contact: Tom Narock, tnarock@ndm.edu

1. Introduction

The American Geophysical Union (AGU) is an Earth and space science professional society based in the United States. AGU publishes scientific journals, sponsors meetings, and supports education and outreach efforts to promote public understanding of science. Research conducted by AGU members ranges from the Earth's deep interior to the outer planets of our solar system. Despite the *American* in its name, roughly 40% of the AGU's membership comes from outside of the U.S.[1]

Each year, the AGU hosts a Fall Meeting that draws tens of thousands of participants. The research presented at these meetings has been discussed and debated extensively. However, little research exists on the AGU meeting itself. In this work, we apply network analysis and scientometrics to seventeen years of AGU Fall Meetings. We model the AGU Fall Meetings as graphs in which presentation co-authors are connected nodes and analyze these graphs to ascertain their structure and properties. We are interested in what the structure and network properties can tell us about the scientometrics of the AGU.

Scientometrics is the science of measuring and analyzing science itself, such as a discipline's structure, growth, change, and interrelations (Hood and Wilson, 2001). Vassily Nalimov first coined the term in the 1960s and subsequent work has focused on a discipline's methodologies and principles as well as individual researchers' scientific output (Braun, Glänzel, and Schubert, 2006; Hirsh, 2005). Here, we are interested in how science collaboration and networking are taking place and how the procedures of the AGU Fall meeting could be enhanced to facilitate better scientific communication and collaboration.

2. Datasets, Assumptions, and Limitations

2.1 Dataset

The data in this study came from the AGU Abstract Browser[2]. The Abstract Browser is a publicly available database of historical abstracts presented at AGU meetings. This database contains abstracts from meetings other than the Fall Meeting, such as the Ocean Sciences Meetings; however, we limited our study to Fall Meetings only. The Fall Meetings are the largest of the AGU-hosted meetings and are multi-disciplinary. Restricting our study to Fall

---

[1] Based on data from the AGU's membership page: https://membership.agu.org/
[2] http://abstractsearch.agu.org/about/

Meetings only provides the most data and also ensures equal coverage of the sub-domains covered by AGU. Our study includes 17 years of data and covers the Fall Meetings from 2000 to 2017.

The AGU is divided into sections representing the subdisciplines of Earth and space science. As science evolves over the years, new sections are formed, and older ones can be merged or dissolved. The sections on which we had data to perform our analysis are listed in Table 1.

Table 1. The AGU sections covered in this study.

Abbreviation	Full Name
A	Atmospheric Sciences
AE	Atmospheric and Space Electricity
B	Biogeosciences
C	Cryosphere
DI	Study of the Earth's Deep Interior
ED	Education and Human Resources
EP	Earth and Planetary Surface Processes
G	Geodesy
GC	Global Environment Change
GP	Geomagnetism and Paleomagnetism
H	Hydrology
IN	Earth and Space Science Informatics
MR	Mineral and Rock Physics
NG	Nonlinear Geophysics
NH	Natural Hazards
NS	Near Surface Geophysics
OS	Ocean Sciences
P	Planetary Sciences
PA	Public Affairs
PP	Paleoceanography and Paleoclimatology
S	Seismology
SA	SPA-Aeronomy
SH	SPA-Solar and Heliospheric Physics
SM	SPA-Magnetospheric Physics
T	Tectonophysics
U	Union
V	Volcanology, Geochemistry, Petrology

Data were retrieved by programmatically querying the AGU Abstract Browser's Linked Open Data interface[3,4]. Linked Open Data (LOD, Berners Lee, 2006; Bizer et al., 2009) is part of the methods and tools collectively known as the *Semantic Web* (Hitzler et al., 2010), which aim to bring machine-readable meaning to the Web through common data formats, exchange protocols, and computational reasoning. The LOD methodology has become a widely adopted data sharing format and at last count (Hogan et al., 2011), roughly thirty billion semantic statements were available on the emerging "Web of Data". In 2012 the AGU's historical abstracts were converted to LOD (Narock, Rozell, and Robinson, 2012; Rozell, Narock, and Robinson, 2012) with new meeting data being added each year.

2.2 Limitations and Assumptions

The Abstract Browser contains Fall Meeting data such as sessions held, presentations given in each session (including title, authors, affiliations, and an abstract), and the AGU section in which the session was held. However, the author data contains only email address, last name, and initials. Moreover, the same author sometimes has only a first initial while other times having a first and middle initial. The first author of this study is a prime example. He appears in the abstract database as both: T. W. Narock and T. Narock. This raises significant challenges for autonomously disambiguating people. Further complicating this issue is the case where authors change institutions. For example, T. Narock appears with his graduate school email address and later with the email address of his affiliation post-graduation. Each author does have an organizational affiliation provided; however, this data is also messy and difficult to use for disambiguation. There is no standard naming convention and the same institution often appears with multiple names. For example, the NASA Goddard Space Flight Center is listed as NASA/Goddard, NASA/GSFC, and NASA/Goddard Space Flight Center. Ideally, authors would be listed with their ORCID (Haak et al., 2012); however, at present, such data is not available via any public AGU interface that we are aware of. Lacking the means to perform a large-scale crowdsourced disambiguation project, we sought other means to disambiguate authors.

We considered email address to be a unique and distinguishing feature. Our disambiguation efforts consisted of finding all cases where email address and last name were the same, but initials only partially matched. For example, [T. Narock, tom.narock@gsfc.nasa.gov] was considered the same person as [T. W. Narock, tom.narock@gsfc.nasa.gov]. This approach identified 56,155 matches, which we corrected in our dataset. Yet, there are likely many other authors who were not disambiguated. We identified an additional 19,896 cases where last names matched, initials were a partial match, and email addresses differed (e.g. [T. W. Narock, tom.narock@gsfc.nasa.gov] and [T. Narock, tnarock@ndm.edu]. Many of these people are likely the same (the example given here is known to be the same); yet, in the vast majority of cases we have no means of knowing for sure and have chosen not to claim these authors as identical. Thus, our results have an inherent

---

uncertainty to them. Specifically, the network graphs we construct from the AGU data likely have multiple nodes representing the same person. As such, we consider the network analysis portion of our study an upper limit. We know that the actual values for network density and connected components are not higher than the values reported here, and they would likely be a bit smaller had we been able to uniquely identify all authors in our dataset. Despite this limitation, we feel our analysis can still provide useful insights into the AGU meetings.

All networks are comprised of nodes (also called vertices) and edges (connections between the nodes). Networks also come in multiple types ranging from directed to undirected. Twitter is an example of a directed network. Edges have directionality in a directed network. For example, Twitter user A can follow user B; however, user B is not obligated to follower user A back. The edge between users A and B would have directionality. In an undirected network all edges are bidirectional by default. This is how "friending" works in Facebook. Both users (nodes) must agree to the "friendship" and a link (edge) is created. There are no directed edges allowed in an undirected network.

We model each AGU section as an undirected network based on co-authorship. If A co-authored a presentation with B and C, then A, B, and C become nodes in the network with bidirectional links between each (e.g. A-B, A-C, B-C). We do not apply any weighting to the edges. If authors A and B co-authored a presentation at the 2000 Fall Meeting and then again at the 2010 Fall Meeting this adds no new information to the graph. We also consider edges to be eternal when studying the temporal evolution of the network. For example, if authors A and B co-authored a presentation at the 2000 Fall Meeting these nodes and edges persist in 2017 even if those authors never co-authored another presentation. We also note that we are measuring co-authorship and not necessarily collaboration. Our dataset does not contain references and acknowledgements used in presentations. There may be secondary connections (e.g. citing a paper or acknowledging a discussion) that do not show up as edges in our graphs.

2.3 Open Source Software

The analysis software used in this study is freely and publicly available at: https://github.com/narock/agu_analytics. The graph data generated from our software is available at: https://figshare.com/articles/AGU_Network_Analysis/6625673

3. Network Analysis

3.1 Network Density

Network density is defined as the ratio of actual connections to possible connections. Possible values for network density range from 0 (no connections at all) to 1 (everyone is connected to everyone else). Figure 1 illustrates the concept of network density on sample networks. In 1.) there are three nodes and three potential connections. These three potential connections are realized as all nodes are connected to each other. This is representative of the

AGU case in which A, B, and C have co-authored presentations with each other; although, not necessarily the same presentation. The network in 1.) has a density of 3/3 = 1.

The network shown in 2.) has the same three potential connections. However, only two of the nodes are actually connected. In this example, A has co-authored a presentation with B and B has co-authored a presentation with C; yet, A has not co-authored a presentation with C. The network in 2.) has a density of 2/3 = .67.

It's unlikely that a real-world network such as the AGU would have network density of 1. Given the diversity of research topics it's unlikely that the network would be completely connected. But, what are the actual density values and how do they change over time?
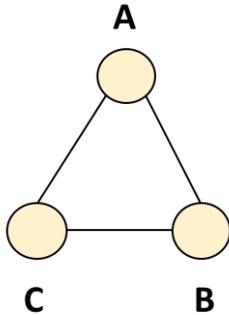
1.)   **2.)**
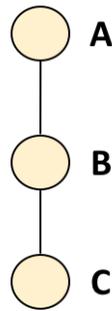
Figure 1. Example networks and network density.

To answer these questions, we first considered each AGU section to be its own network. Yearly network graphs were then created for each section using the Abstract Browser data. Next, we computed the percentage change in network density for each section. We note that percentage change values do not always encompass the whole 17 years of the data. For example, the Earth and Space Science Informatics (IN) section did not come into existence until 2005. Percentage change was computed using the first year in which we had data and 2017. Results are shown in Figure 2.
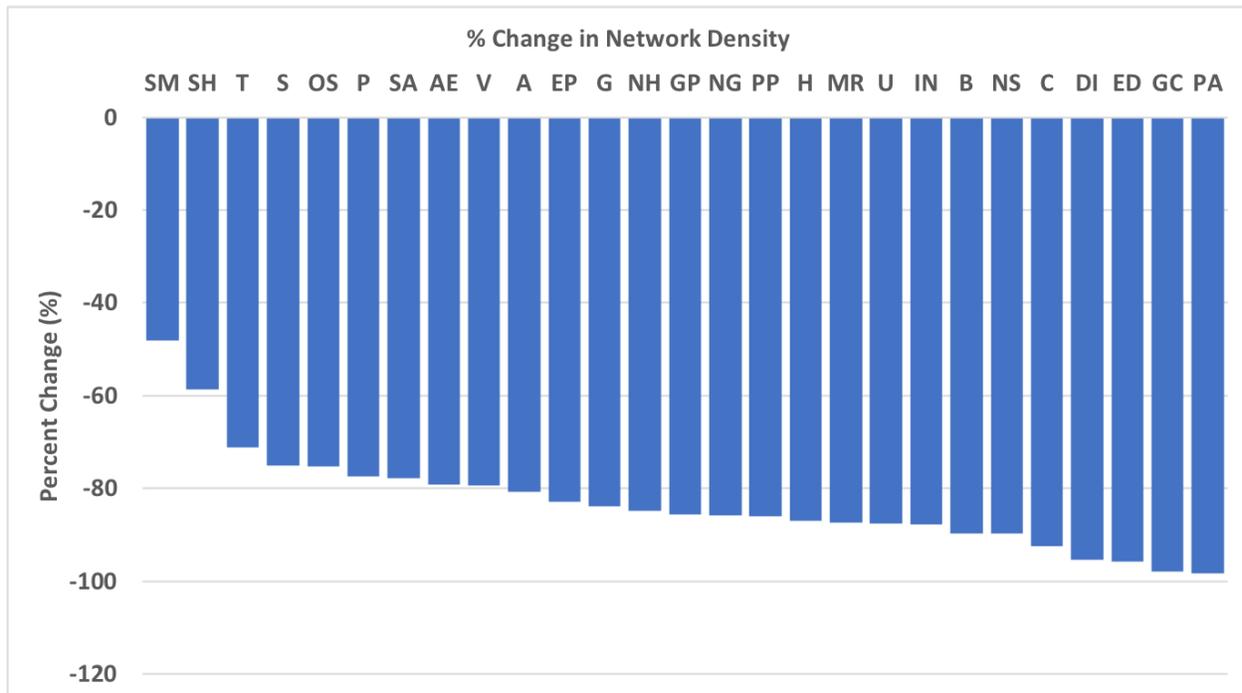
Figure 2. Percentage change in network density.

Network density decreases for all sections. This is telling us that nodes are being added faster than edges. In practical terms, the rate at which new people (nodes) are attending AGU sessions is greater than the rate at which continuing attendees (nodes) are making new connections. Again, these percentage change values should be considered as upper limits due to our inability to completely disambiguate the authors in our data. We know that the decline in density for each section is no more than what is shown in Figure 2. Yet, it is likely a bit smaller for each section.

3.2 Connected Components

In graph theory, a connected component of an undirected graph (also referred to as a component) is a subgraph within the whole graph. Figure 3 shows an example. The network in the figure is comprised of three connected components. Although not shown here, an isolated node not connected to any other nodes in the network is also considered a connected component. Analysis of connected components within the AGU networks gives us an indication of how fragmented the networks are.
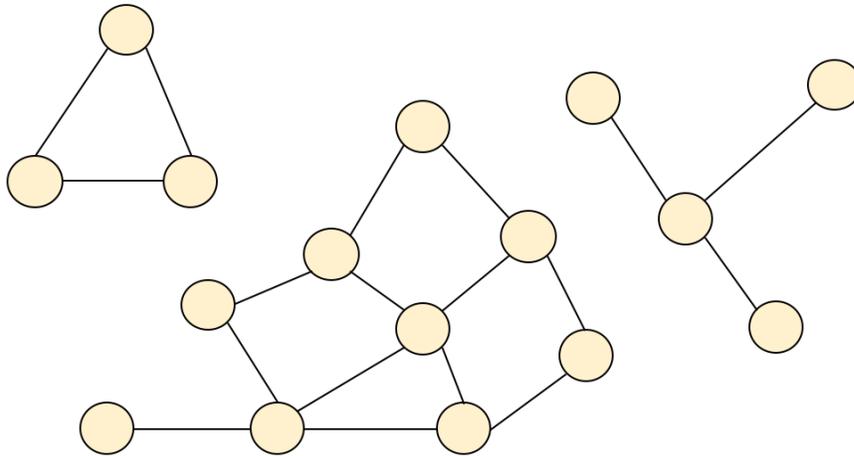
Figure 3. An illustration of connected components. This graph has three connected components.

Table 2 lists the connected components of the AGU section graphs. Specifically, we combined all 17 years of data for each section and computed the number of connected components for the section, the number of nodes in the largest component, the number of components comprised of only one node, and the percentage of each section network that is single node components.

Table 2. Connected Components of AGU Section Graphs.

Section	Number of Connected Components	Nodes in Largest Component	Number of Single Node Components	% of Whole Network that is Single Node
A	1139	29431	439	39%
AE	108	1812	31	29%
B	1407	28888	471	33%
C	508	11148	160	31%
DI	267	3500	57	21%
ED	1970	7689	830	42%
EP	883	7748	177	20%
G	393	7471	136	35%
GC	1770	15936	581	33%
GP	297	5194	98	33%
H	2060	35304	793	38%
IN	884	8074	283	32%
MR	481	4179	109	23%
NG	884	2127	248	28%
NH	1125	5766	283	25%
NS	501	2175	80	16%
OS	1235	16136	504	41%
P	495	11259	202	41%

PA	1154	581	441	38%
PP	530	13932	172	32%
S	607	13192	271	45%
SA	202	5526	88	44%
SH	262	6052	136	52%
SM	169	7071	95	56%
T	814	19208	354	43%
U	1426	6652	630	44%
V	889	17477	348	39%

The diversity of research topics likely guarantees that we are going to have some fragmentation of the network. Not everyone is working on the same topic and we would expect to see the number of connected components greater than 1. Moreover, there's nothing wrong with working by oneself and single node components are to be expected. Yet, the numbers in Table 2 seem too large to us. Each connected component can be thought of as a cluster (or clique) of presenters. Connected components have no link (edge) between them as shown in the Figure 3 examples. AGU attendees may be seeing new presentations and having useful discussion across connected components; however, it does not appear to be the case that these discussions are stimulating organic growth and connecting the components. We return to this issue in our discussion in section 4.

3.3 Multi-Disciplinary Authors

We define a multi-disciplinary author as anyone who appears in the network graph of more than one AGU section. We looked at all pair-wise comparisons of sections and obtained the results in Figure 4, which shows the number of unique authors who have appeared in both sections over the 17 years of data. For instance, the nearly 8,000 occurrences between the Biogeosciences (B) and Hydrology (H) indicates that 8,000 individuals have appeared on presentations in both Biogeoscience and Hydrology sessions during the time period 2000 to 2017. The top half of Figure 4 shows multi-disciplinary authors across all pair-wise comparisons. The bottom half of Figure 4 is an enlargement of the top 50 section pairs. The arrow from the bottom half of the figure to the top half indicates the area of enlargement.
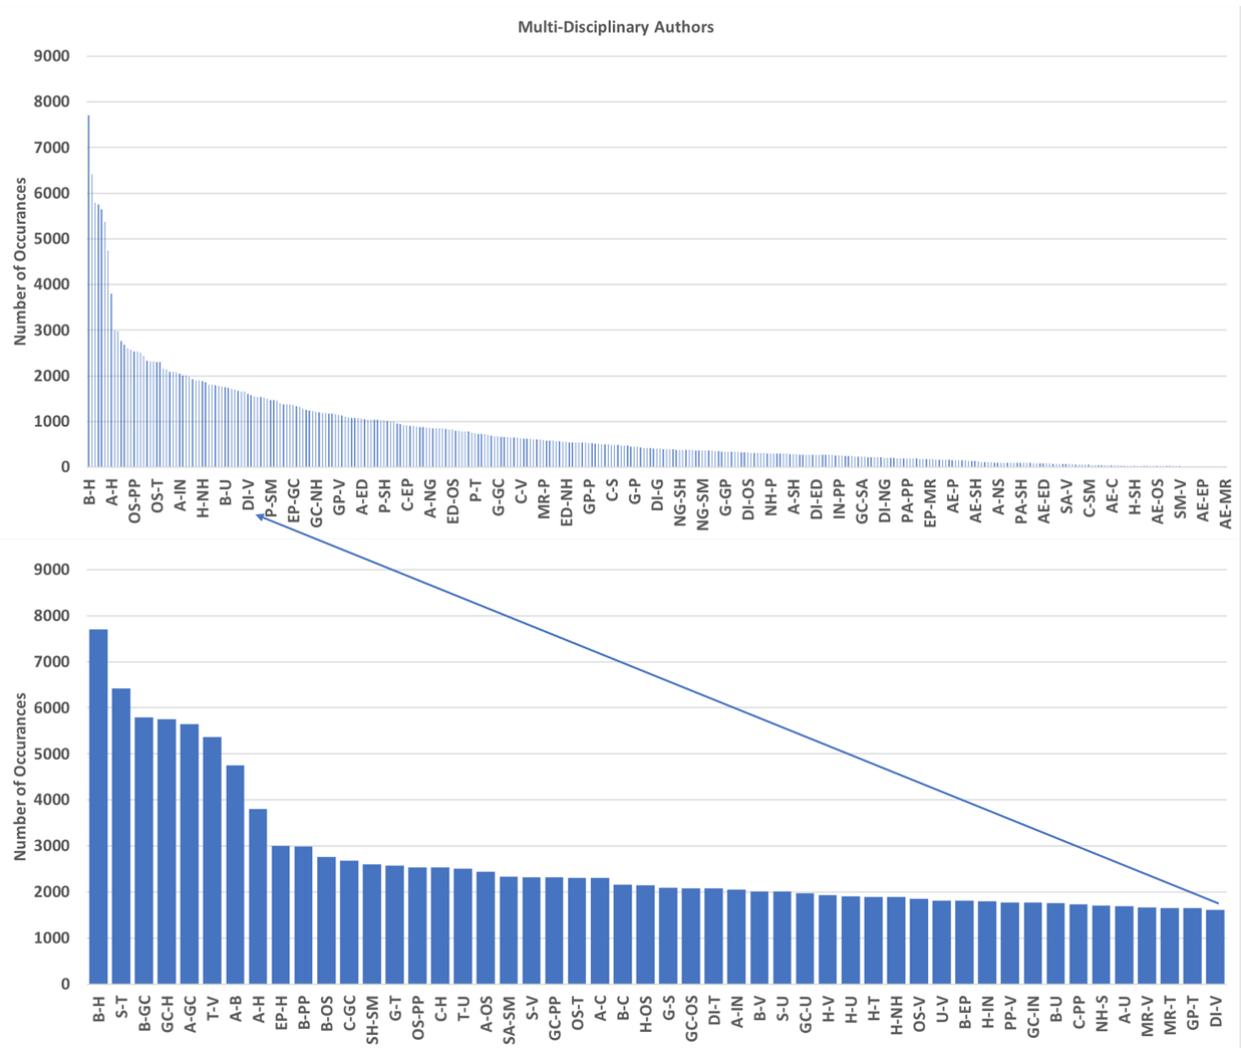
Figure 4. Number of occurrences of authors presenting in more than one section over the years 2000-2017. The bottom half of the figure is an enlargement of the top 50 multi-disciplinary sections.

3.4 Keyword Usage Across Sections

Authors submitting to the Fall Meeting are asked to tag their abstracts with keywords from the AGU's keyword hierarchy[5]. We computed counts of each keyword category for each year of our dataset across all sections. For instance, *Post-secondary Education* and *Teaching Methods* are sub-topics within the higher-level *Education* section of the keyword hierarchy. If the Hydrology section had an abstract tagged with *Post-secondary Education* in 2005 and an abstract tagged with *Teaching Methods* in 2005 then this would be counted as two *Education* abstracts for the year 2005. We note that abstracts are not exclusive to one keyword group. Authors are free to self-tag their abstracts with multiple keywords that may span multiple parts

---

[5] https://publications.agu.org/author-resource-center/index-terms/

of the keyword hierarchy. This is reflected in our analysis where the same abstract may contribute to keyword usage counts in multiple parts of the keyword hierarchy.

For clarity of display, we filtered out keyword groups that did not reach 100 occurrences during the 17 years in which we had data. Figures 5 through 8 highlight specific trends in keyword usage that were observed in our data. The full set of images showing keyword usage from all keyword categories is included in the Appendix.

Scenario 1 - Two (or more) seemingly unrelated groups use the same topics

The Earth and Space Science Informatics (IN) section self-describes[6] itself as being "concerned with evolving issues of data management and analysis, technologies and methodologies, large-scale computational experimentation and modeling, and hardware and software infrastructure needs". These concerns span many areas of geoscience and one might expect IN related keywords to appear in several computationally intensive domains. This does in fact occur as evidenced in Figure 5. Yet, we also see a sharp rise in the Natural Hazards section's usage of IN keywords from 2016 to 2017.
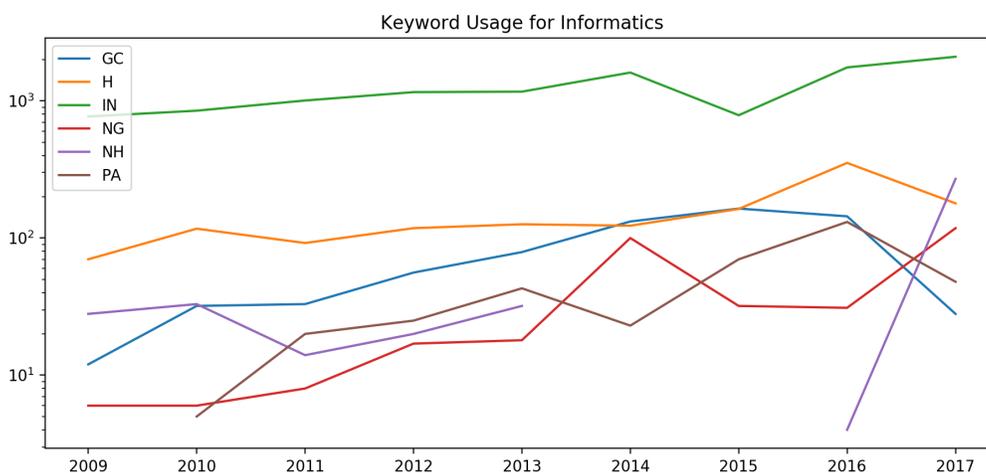
Figure 5. Informatics themed keyword usage

Prior to 2016, the Natural Hazard section never had a year in which they received more than 33 abstracts tagged with IN keywords (this occurred in 2010). Suddenly, in 2017 they received 270 abstracts tagged with IN keywords. This is up from 4 such abstracts in 2016. The bulk of these 270 abstracts in 2017 can be attributed to the *Data Assimilation, Integration, and Fusion* and *Forcasting* topics. These two keyword categories accounted for 87% of the Natural Hazard IN-related abstracts in 2017. To us, this is indicative of the power of simple scientometric visualizations. By simply counting keywords we can begin to identify emerging collaborations, which, as we discuss further in the next section, can be exploited by meeting and section leadership to better structure future Fall Meetings.

Scenario 2 – Increase in Volume

The Planetary Science section is the primary user of Astrobiology keywords as shown in Figure 6. Usage from 2005 to 2010 was more or less consistent. However, beginning in 2011 a sudden increase in usage is seen that continues to today. A similar trend is seen with Education keywords in Figure 7. In 2015, Public Affairs and Union sessions saw an increase in abstracts tagged with Education keywords.

Keyword Usage for Planetary Sciences: Astrobiology
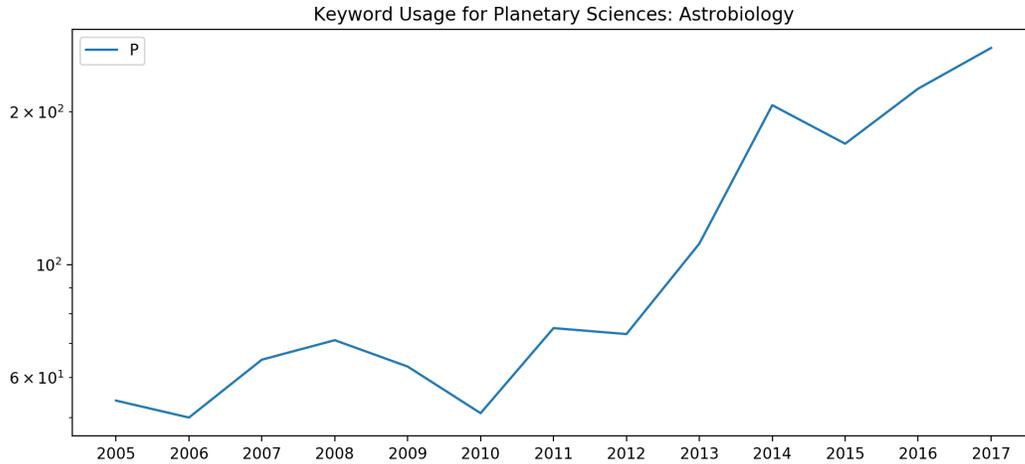
Figure 6. Astrobiology keyword usage.

It may not be surprising that planetary scientists are using astrobiology terms to tag their abstracts. Meeting attendees may even have anecdotal evidence of observing this themselves. Yet, had someone been tracking this data in 2012 and 2013 we could have seen this trend emerging. This information could have gone into meeting planning and potentially led to more physical space at the meeting venue, joint sessions, increased public outreach, and other initiatives that could have maximized the dissemination of astrobiology science.

A related trend is shown in Figure 7 where Public Affairs and Union sessions show an uptick in Education-related abstracts from 2014 to 2017. A scientometrics and data driven AGU could leverage this information in being proactive with joint sessions and when/where presentations are given at the Fall Meeting. We explore this in more detail in the next section.
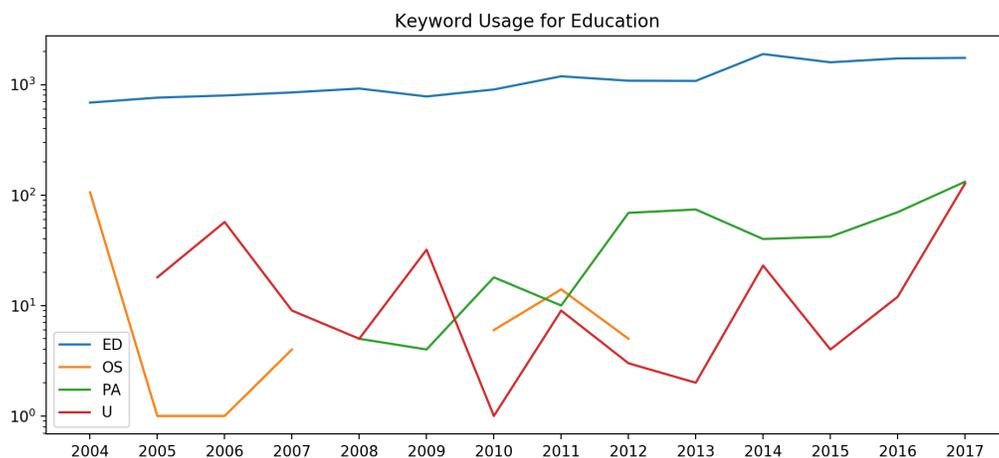
Figure 7. Education keyword usage.

Scenario 3 - Keyword Usage May Indicate New Science

The Earth and Space Science Informatics section was formed in 2005. From 2005 until 2008 this section did not have any section-specific keywords in the aforementioned AGU keyword hierarchy. In 2009 IN-specific keywords were introduced. We see this clearly in Figure 8 where IN's usage of *General or Miscellaneous* keywords decreased significantly between 2008 and 2011 as IN-specific keywords began to be used. Yet, we also see a steady increase in *General or Miscellaneous* from 2011 to 2015. Further analysis of this keyword group reveals steady usage of *General or Miscellaneous: Instruments useful in three or more fields* and *General or Miscellaneous: Techniques applicable in three or more fields* during the time period 2011 to 2015. This is suggestive to us that emerging computational approaches and collaborations are not adequately reflected in the AGU keyword hierarchy. This may be more than just the frustration of not finding an appropriate keyword to tag one's abstract. New science may be emerging that could be capitalized on in subsequent Fall Meetings if we are watching the evolution of the AGU network. Further exploration of this particular trend would involve more data than we currently have available and is outside of our current scope.
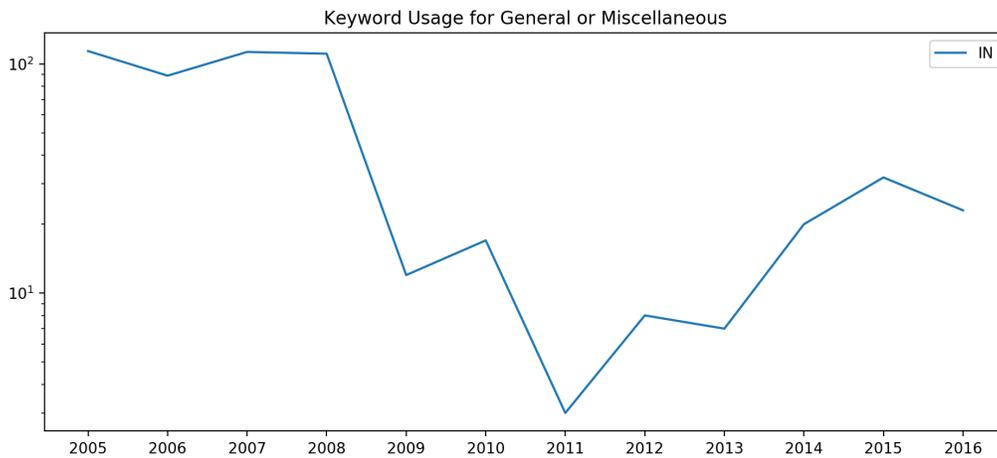
Figure 8. General or Miscellaneous keyword usage.

4. Scientometrics

AGU Fall Meetings are already very busy. Figure 9 shows the number of presentations given each year from 2000 to 2017. We see a steady increase in presentations with the 2017 Fall Meeting having over 20,000 accepted presentations. Fall Meeting attendees are already hard-pressed to see everything of interest. Using network analysis and having section leaders be proactive prior to a meeting can improve efficiency of science communication and collaboration.
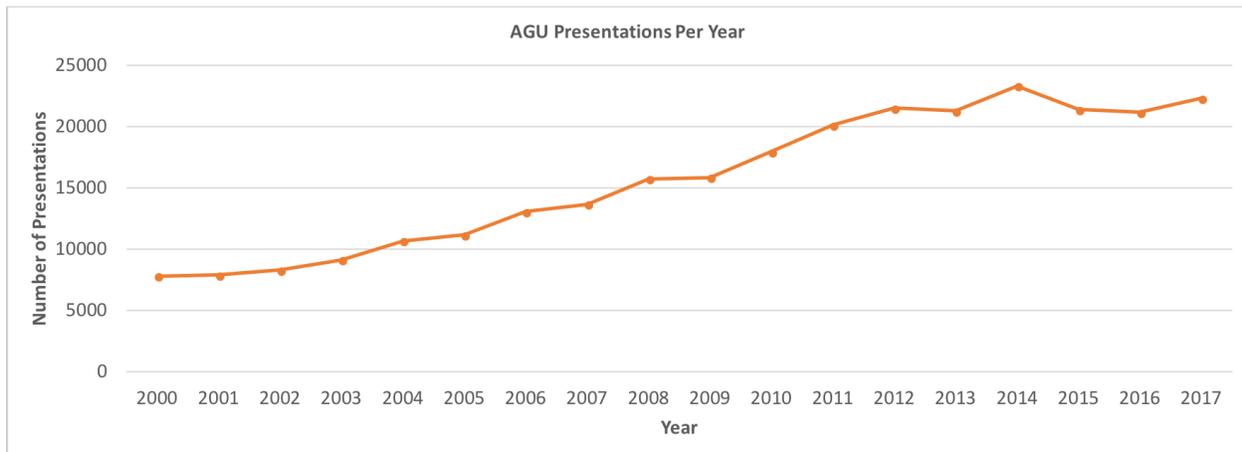
Figure 9. Number of presentations given at AGU Fall Meetings each year.

In regard to network density and connected components, there is no optimal network clustering value. However, lower density networks comprised of many loosely connected clusters have been shown to be beneficial (Burt, 2004). In these networks, everyone doesn't already know each other, and multiple clusters leads to new and unique perspectives. On the contrary, when everyone knows everyone else (density=1) you're more likely to repeatedly

hear the same ideas (Burt, 2004). Moreover, the number of connected components and single author presentations (Table 2) is worrisome given that analysis of scientific publications (Dong et al., 2017) has revealed a trend towards team science and increased connections.

In order for information to spread across a network there needs to be connections between the clusters. We want to avoid the scenario depicted in Figure 3 and have at least one connection between each connected component in an AGU section. By knowing how many connected components there are, what is the primary research topic of each (most used keyword), and whom the components are comprised of, can be beneficial for meeting planners and section leadership. For the AGU Fall Meeting, session proposal is open to any self-organized group of up to four AGU members. Authors then opt to have their submission assigned to a particular session. We could make this process more proactive by providing section leadership with connected component data and encouraging connections between specific AGU members. This could range from informal networking events to suggesting session co-conveners.

4.1 Steps Towards Optimizing Meeting Space

One potential means of enhancing the AGU Fall Meeting is to optimize the physical layout of the event. Historically, oral presentations are arranged by section with a section having all of its talks grouped in the same part of the building. The poster hall is organized alphabetically by section. What if we leveraged what we're seeing in Figures 5 and 7 to physically place related sections next to each other? For example, the 2018 Fall Meeting could place Natural Hazards posters next to Informatics posters to stimulate more discussion. Similarly, Public Affairs and Union sessions could be physically located near Education sessions and, having identified the trend in Figure 7, attendees could be encouraged to visit related presentations they may not otherwise be aware of.

Another option is to facilitate navigation of the meeting via analytics tools built on top of the AGU's historic meeting data. A simple example is shown in Figure 10. This so-called force directed graph adds additional information to a standard network graph. In a force directed graph the distance between two nodes is indicative of the strength of the connection. For instance, in Figure 10 we are showing the 10 AGU members who most used the oceanographic Aerosols keyword. R. Weber has used this keyword the most over the 17-year period 2000 to 2017. This is indicated in the figure where the R. Weber node is closest to the central Aerosols node.
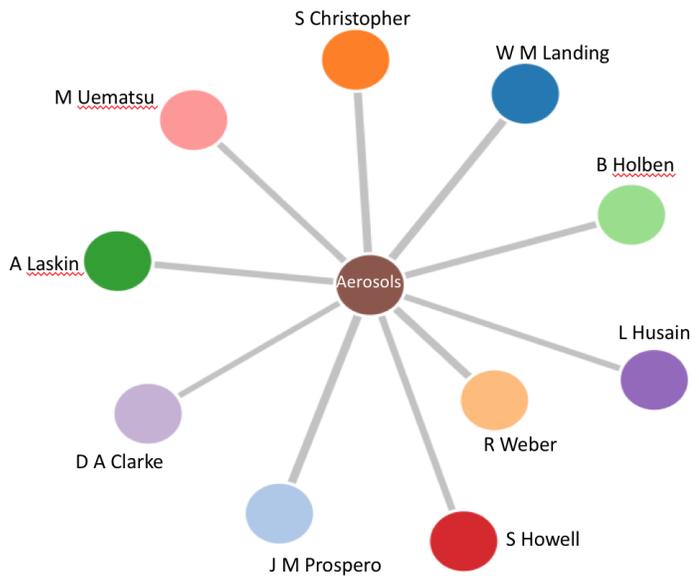
Figure 10. Force directed graph of oceanographic aerosols keyword usage.

We want to be clear that we are ***not*** advocating for any sort of new metric. We do not need to rank researchers nor do we need to rank the value of their work based on where its presented. The journal impact factor already does a poor enough job of this already (Shanahan, 2016). Rather, we are advocating for tools that would help attendees, especially early-career and new attendees, identify whom they might want to seek out based on their research interests. Figures 11 through 13 show an example tool we build for the AGU Open API Challenge[7],[8]. After identifying a researcher, possibly through a visualization like Figure 10, the user is guided through finding that researcher in the historical abstract database (Figures 11 and 12). The co-authorship network is then leveraged to identify all AGU presenters who have co-authored a presentation with the researcher of interest. Figure 13 shows an example for our colleague Peter Wiebe. For brevity, only the 2018 co-authors are shown in the figure. The Abstract column in Figure 13 lists the year of presentation, the section of the presentation, and the presentation ID. Each row in the Abstract column is a clickable link that will take the user to a web page displaying the presentation title, keywords, and abstract. In this manner, AGU attendees can *follow the network* to explore existing connections amongst nodes and topics. At present, Fall Meeting data is not available in the Abstract Browser until after the Fall Meeting concludes. Making this data available prior to the meeting could lead to new tools and apps. AGU does appear headed in this direction with its recent Open API Challenge.

---

[7] https://developer.agu.org/projects/
[8] http://apiprojects.agu.org/project1/

Author Last Name

Wiebe

Submit

Figure 11. Step one of the author search tool.

Name	Affiliation
Wiebe, B	University of the Fraser Valley, Abbotsford, BC, Canada
Wiebe, P H	Woods Hole Oceanographic Inst, Woods Hole, MA, United States
Wiebe, K	International Food Policy Research Institute, Washington, DC, United States

Figure 12. Step two of the author search tool. The system returns all matching authors.

Person	Abstract
Ake, H	2018, OS, OD34B-2759
Copley, N J	2018, OS, OD34B-2759
Saito, M A	2018, OS, OD34B-2759
Switzer, M	2018, OS, OD34B-2759
Biddle, M	2018, OS, OD34B-2759
Kinkade, D	2018, OS, OD34B-2759
Rauch, S	2018, OS, OD34B-2759
York, A	2018, OS, OD34B-2759
Shepherd, A	2018, OS, OD34B-2759
Gjster, H	2018, OS, ME34A-0957
Ingvaldsen, R B	2018, OS, ME34A-0957
Knutsen, T	2018, OS, ME34A-0957
Hobson, B	2018, OS, IS43B-02
Risi, M	2018, OS, IS43B-02
Yoerger, D R	2018, OS, IS43B-02
Breier, J A Jr	2018, OS, IS43B-02
Fujii, J	2018, OS, IS43B-02

Figure 13. The result of our author search tool is a web table with links to everyone who has ever co-authored a presentation with the author of interest. Users can explore the abstracts and network connections of the those co-authors – and their co-authors.

4.2 Steps Toward Gender Equality

Ford and colleagues (Ford et al., 2018) have identified a gender imbalance in AGU presentations. Women are invited and assigned oral presentations less often than men. It was found that male primary conveners allocate invited abstracts and oral presentations to women less often and below the proportion of women authors.
This trend was apparent regardless of the male primary conveners being students or at more senior career stages. Ford et al. (2018) also identified that women elect for poster only presentations more so than men.

The dataset used in this study has a longer timespan than the one used by Ford et al. (2018). However, our dataset does not include gender or career stage information. We cannot add any new information on the gender imbalance. Yet, we can suggest that new tools leveraging the AGU network, such as those mentioned above, could be helpful in addressing the gender imbalance. We would recommend that AGU members have the option of making limited personal information (e.g. gender and career stage) publicly available. Ideally, we'd recommend that AGU members have ORCID profiles that could also be linked to and be used for disambiguation. This would have two benefits. First, it would allow for public tracking of imbalances (gender and/or career stage imbalance) on a section by section basis. Ford and colleagues (2018) did an amazing job of identifying a gender imbalance and bringing to the community's attention. Yet, it required requesting non-public data. We shouldn't have to periodically request data and check in on our progress toward equality. We would like to see an open sharing of data and a continual open evaluation of progress over time. Second, network exploration tools can help identify whom to invite for panels and invited presentations. We could be collectively working toward presenter recommendation systems that leverage gender, career stage, and keyword usage. Network analysis won't solve the gender imbalance at AGU, but it may provide a step in the right direction.

4.3 Steps Toward Connections to Other Networks

GeoLink (Narock et al., 2014; Krisnadhi et al., 2015; Cheatham et al., 2018) is a collection of Linked Open Data that addresses scholarly discovery and collaboration in the geosciences. GeoLink leverages the Semantic Web to publish open data regarding data centers, digital repositories, libraries, and professional societies. One component of the GeoLink knowledge graph (Cheatham et al., 2018) is a collection of all National Science Foundation (NSF) funded projects. Figure 14 (reproduced from Narock and Wimmer, 2017) illustrates what can be done when one network is connected to another. This figure is produced by subsetting the GeoLink NSF funded projects by people who have presented at AGU. In particular, we are looking at *Semantic web and semantic integration* - a keyword in the Informatics portion of the AGU keyword hierarchy. Combining these two open datasets allowed us to identify which AGU authors had active funded grants at the time of their AGU presentation. We define "active funded grant" as the AGU presentation date falling between the NSF grant's start and end date. We then looked at the distribution of funding sources. Figure 14 shows the NSF divisions and

offices that have funded an AGU author's semantic project. This is only one example and specific to one topic area. Yet, if illustrates the potential of open science and cross-organizational network analysis. We can begin to see how this research topic is funded by the NSF. In addition, we can start to see the scientific results (AGU presentations) attributable to each NSF division. In this regard, AGU scientometrics can go beyond optimizing Fall Meetings to more general enhancements of open science and science communication. Exponential growth is being observed with the amount of available Linked Open Data roughly doubling each year. Corporations (e.g., the BBC and BestBuy), governments (e.g., the U.S. and U.K. governments), Wikipedia, social networking sites (e.g.. Flicker, Facebook and Twitter), and various academic communities are all contributing to the movement (Hogan et al., 2011). We encourage AGU to do the same.
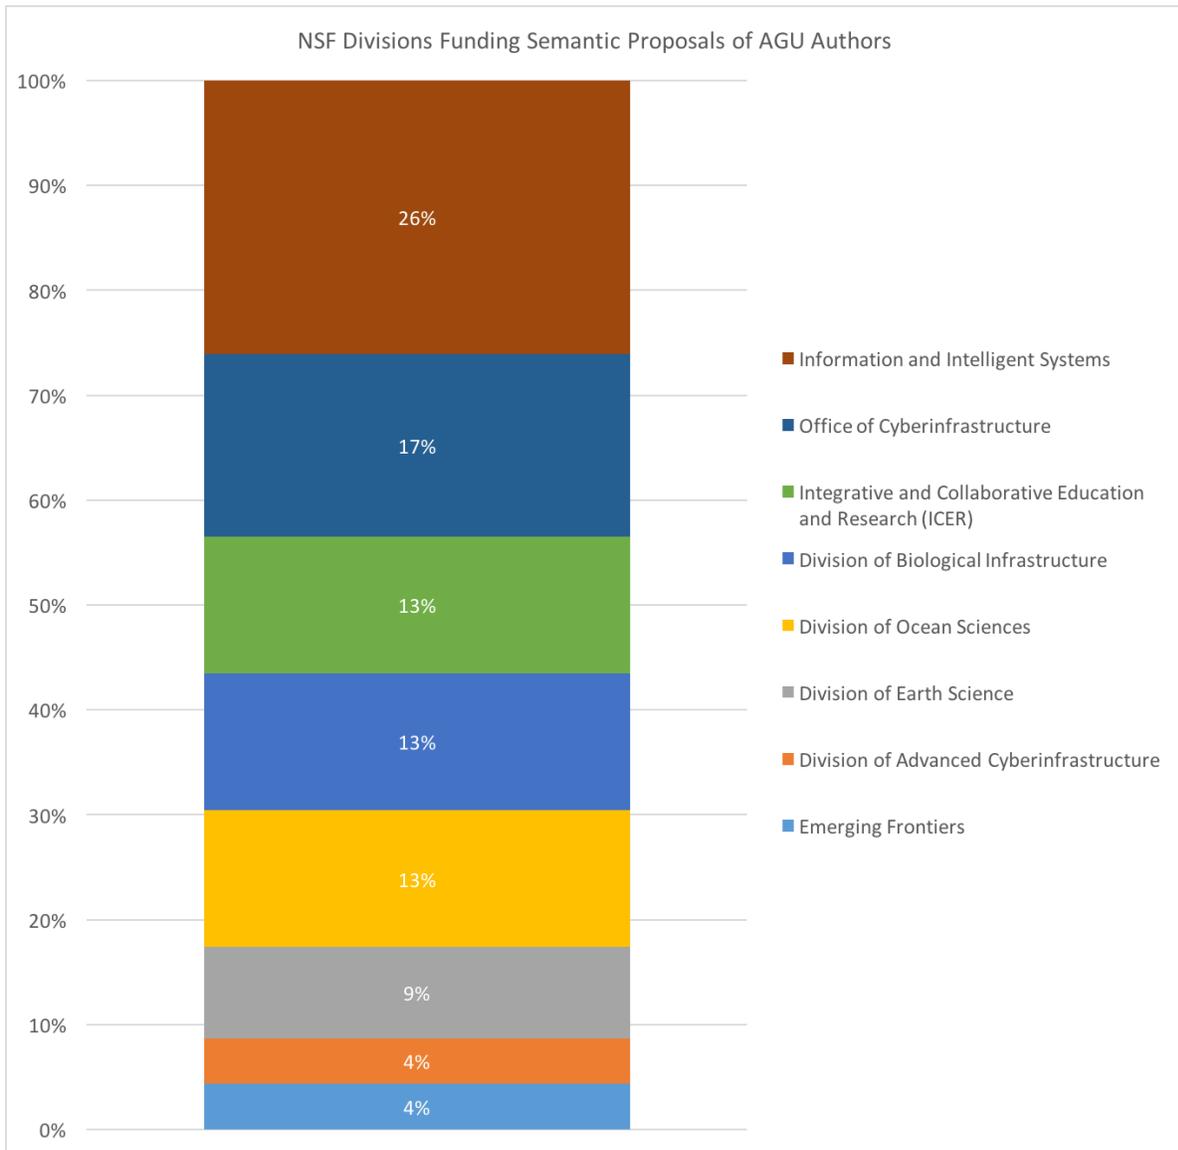
Figure 14. An example of combining network data. Here, AGU and NSF networks are merged to identify where AGU presenters are receiving their funding.

5. Conclusion

AGU is on the cusp of an incredible milestone. Founded in 1919, the AGU will celebrate its centennial in 2019. There is a lot we can learn from the past 100 years. Network analysis, scientometrics, and data science can help us quantify what we're doing right and identify paths toward improvement. Let's leverage open data and open science to improve how we present our science over the next 100 years.

Acknowledgements

References

Berners-Lee, T. (2006). Linked Data - Design Issues. Retrieved May 20, http://www.w3.org/DesignIssues/LinkedData.html

Bizer, C., Heath, T., Berners-Lee, T., (2009), Linked Data – The Story So Far. *International Journal of Semantic Web and Information Systems*, 5(3): 1-22

Braun, T., Glänzel, W., & Schubert, A. (2006). A Hirsch-type index for journals. Scientometrics, 69(1), 169-173.

Burt, R. S., (2004). Structural Holes and Good Ideas, American Journal of Sociology, Volume 110, Number 2, September 2004.

Cheatham, M., Krisnadhi, A., Amini, R., Hitzler, P., Janowicz, K., Shepherd, A., Narock, T., Jones, M., and Ji, P., (2018), The GeoLink knowledge graph, Big Earth Data, Published online: May 18 2018

Dong,Y., Ma, H., Shen, Z., and Wang, K., (2017), A Century of Science: Globalization of Scientific Collaborations, Citations, and Innovations, KDD'17 Applied Data Science Paper, August 13-17, Halifax, NS, Canada.

Ford, Heather L., Brick, C., Blaufuss, K., and Dekens, P. S., (2018), "Gender Representation of Speaking Opportunities at the American Geophysical Union Fall Meeting." EarthArXiv preprint, January 2018. Available online: https://doi.org/10.17605/OSF.IO/6QHVD

Haak LL, Fenner M, Paglione L, Pentz E, Ratner H. ORCID: a system to uniquely identify researchers. Learned Publishing. 2012. 25, pp. 259–264. doi: 10.1087/20120404.

Hirsch, J. E. (2005). An index to quantify an individual's scientific research output. Proceedings of the National academy of Sciences of the United States of America, 102(46), 16569-16572.

Hitzler, P., Krötzsch, M., Rudolph, S. (2010), Foundations of Semantic Web Technologies. CRC Press, 2010.

Hogan, A., Zimmermann, A., Umbrich, J., Polleres, A. and Decker, S., (2011), Scalable and distributed methods for entity matching, consolidation and disambiguation over linked data corpora, Web Semantics: Sci. Serv. Agents World Wide Web, doi:10.1016/j.websem.2011.11.002

Hood, W. and Wilson, C. (2001). The literature of bibliometrics, scientometrics, and informetrics. Scientometrics, 52(2), 291-314.

Krisnadhi, A., Hu, Y., Janowicz, K., Hitzler, P., Arko, R., Carbotte, S., Chandler, C., Cheatham, M., Fils, D., Finin, T., Ji, P., Jones, M., Karima, N., Mickle, A., Narock, T., O'Brien, M., Raymond, L., Shepherd, A., Schildhauer, M., & Wiebe, P. (2015). The GeoLink Modular Oceanographic Ontology, In: Proceedings of the International Semantic Web Conference 2015, Volume 9367 of the series Lecture Notes in Computer Science, pp 301-309

Narock, T. W.; Rozell, E. A.; and Robinson, E. M., (2012), Facilitating Collaboration Through Linked Open Data, Abstract ED44A-02 presented at 2012 Fall Meeting, AGU, San Francisco, Calif., 3-7 Dec

Narock, T., Krisnadhi, A., Hitzler, P., Cheatham, M., Arko, R., Carbotte, S., Shepherd, A., Chandler, C., Raymond, L., Wiebe, P. & Finin, T., (2014), The OceanLink Project, International Workshop on Challenges and Issues on Scholarly Big Data Discovery and Collaboration, 2014 IEEE International Conference on Big Data, 27 October 2014, Washington DC, USA.

Narock, T. W. and Wimmer, H., (2017), Linked data scientometrics in semantic e-Science, Computers & Geosciences, Volume 100, March 2017, pp 87-93

Rozell, E. A.; Narock, T. W.; and Robinson, E. M., (2012), Creating a Linked Data Hub in the Geosciences, Abstract IN51C-1696 presented at 2012 Fall Meeting, AGU, San Francisco, Calif., 3-7 Dec

Shanahan D. R. (2016) Auto-correlation of journal impact factor for consensus research reporting statements: a cohort study. *PeerJ* 4:e1887https://doi.org/10.7717/peerj.1887

Appendix

Keyword Usage for Atmospheric Composition

Keyword Usage for Planetary Sciences: Astrobiology

Keyword Usage for Atmospheric Processes

Keyword Usage for Biogeosciences

Keyword Usage for Cryosphere

Keyword Usage for Computational Geophysics
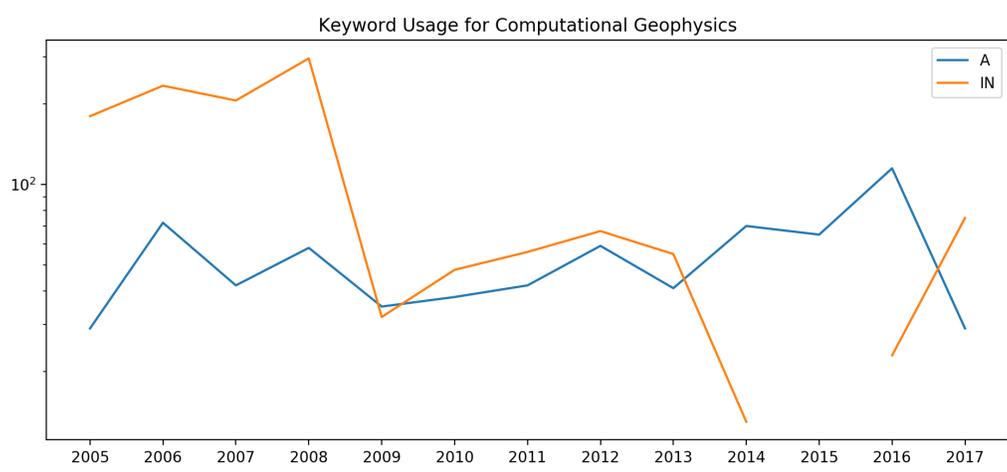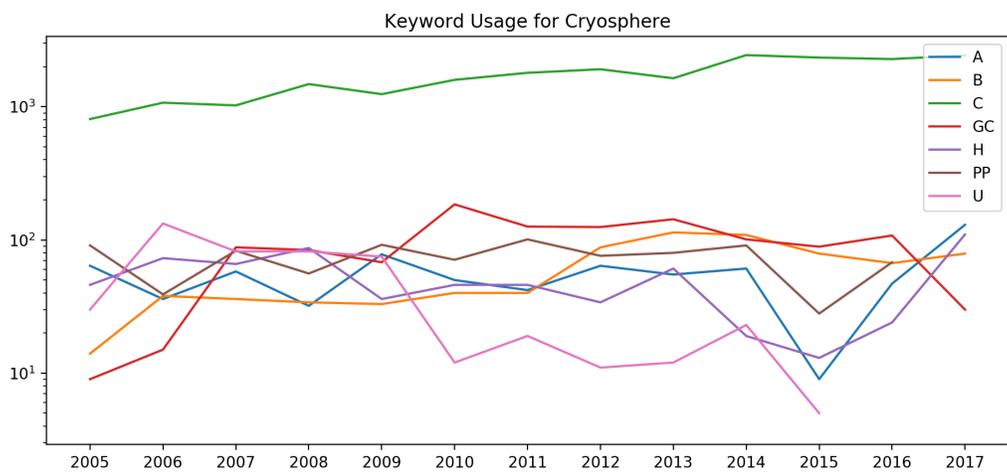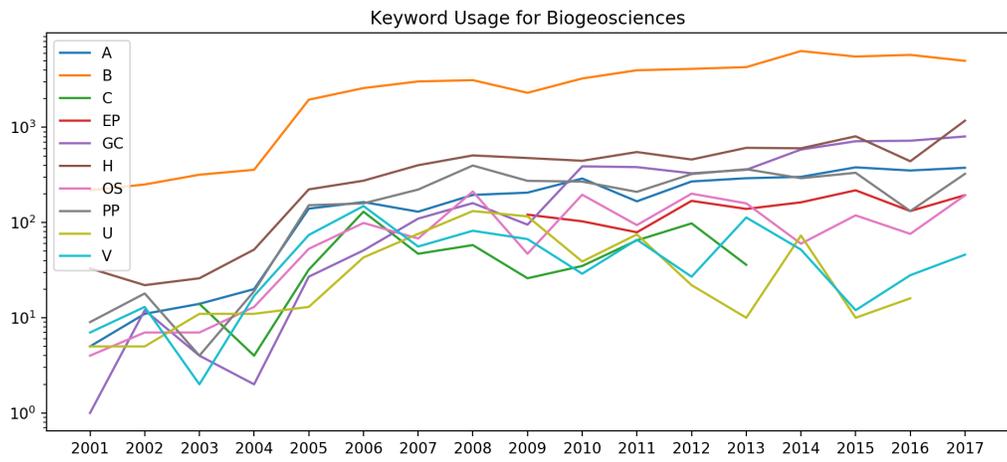
Keyword Usage for Education
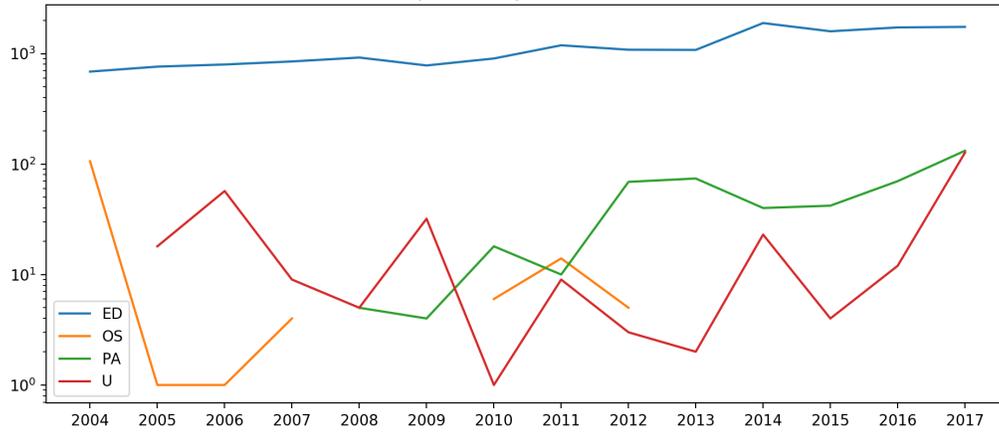
Keyword Usage for Exploration Geophysics
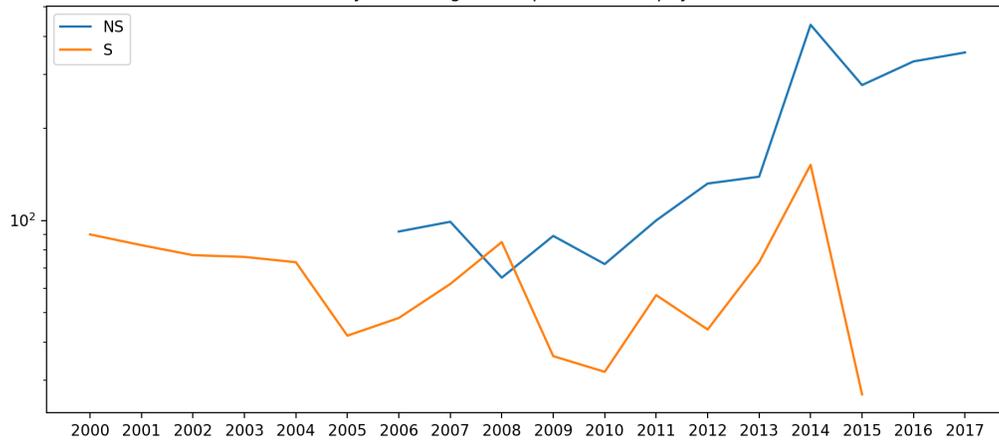
Keyword Usage for Geochemistry
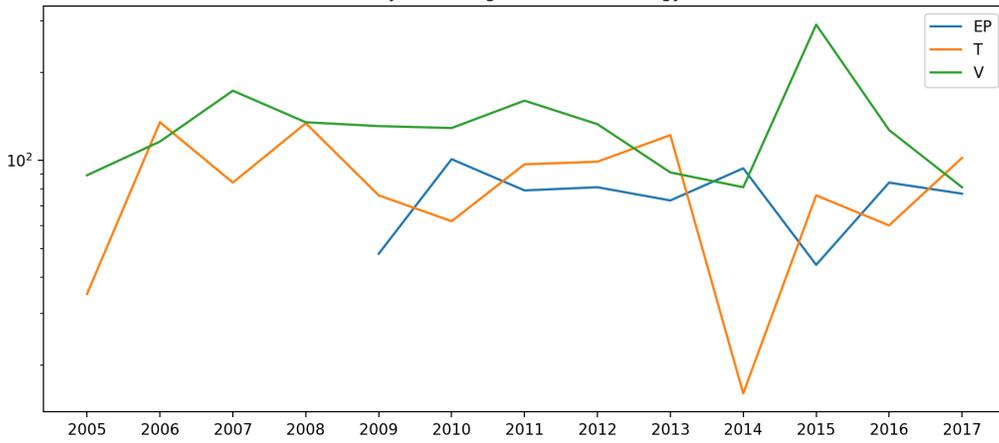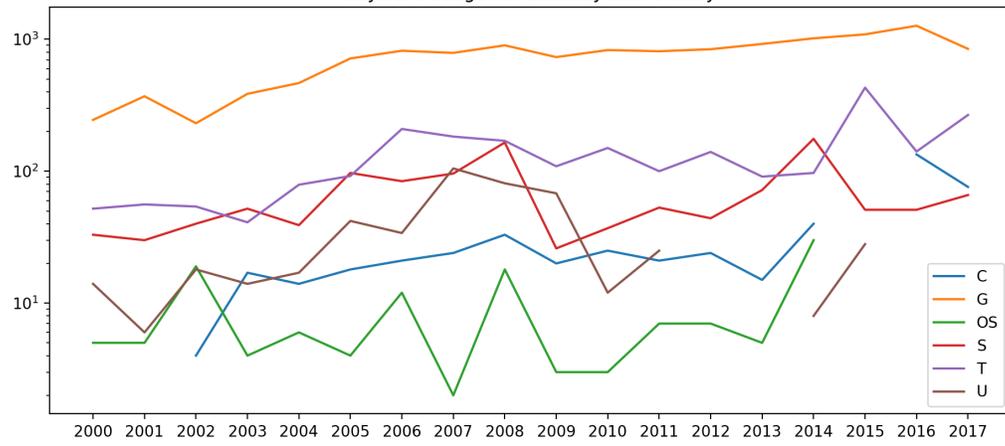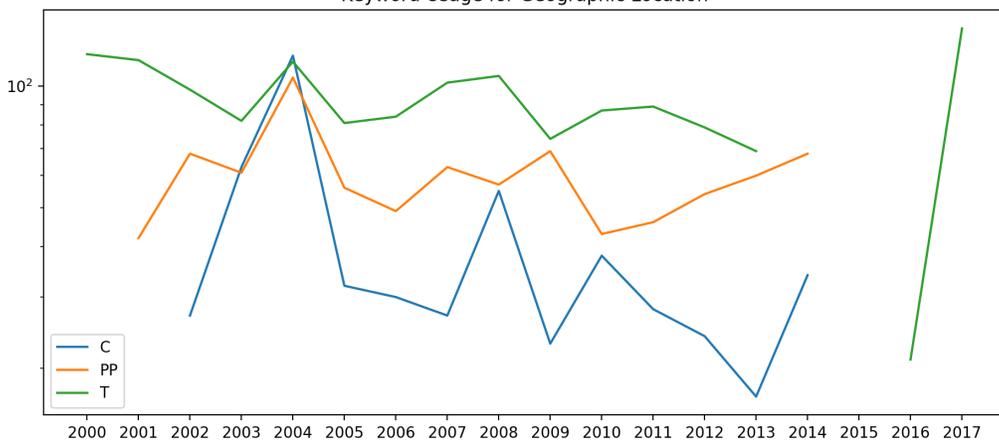
Keyword Usage for Marine Geology and Geophysics
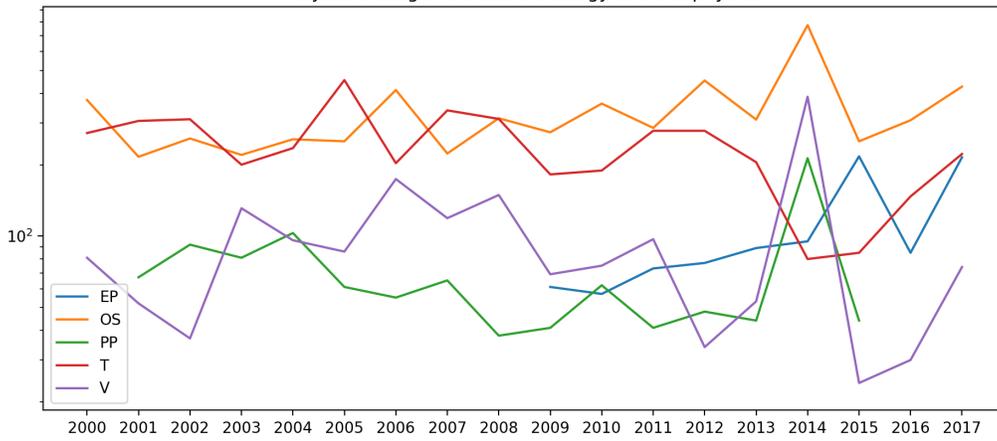

Keyword Usage for Geomagnetism and Paleomagnetism
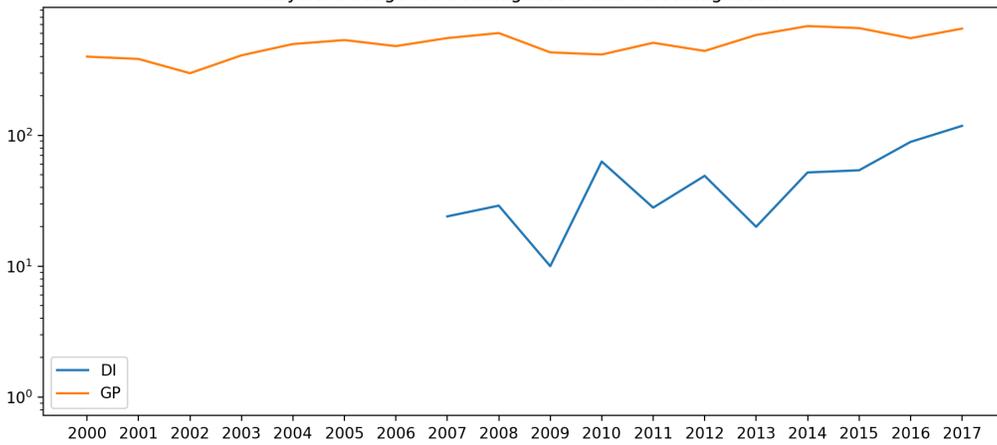
Keyword Usage for Global Change

Keyword Usage for Hydrology

Keyword Usage for Informatics

Keyword Usage for Interplanetary Physics
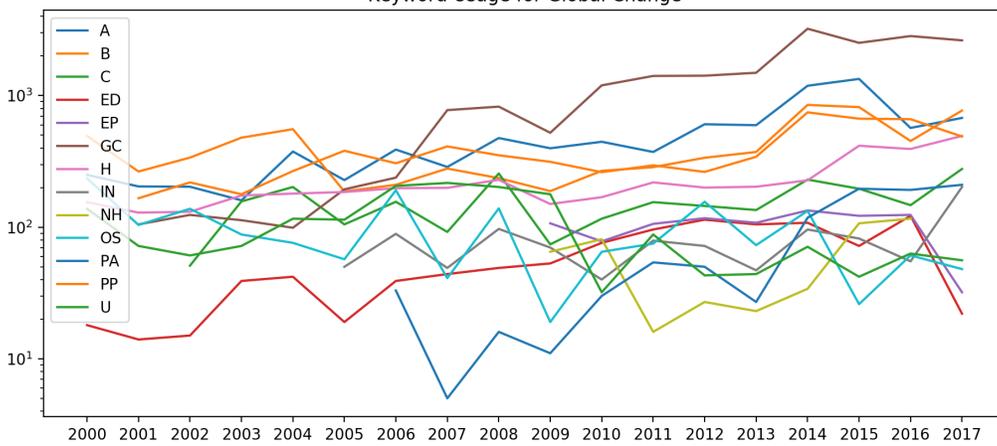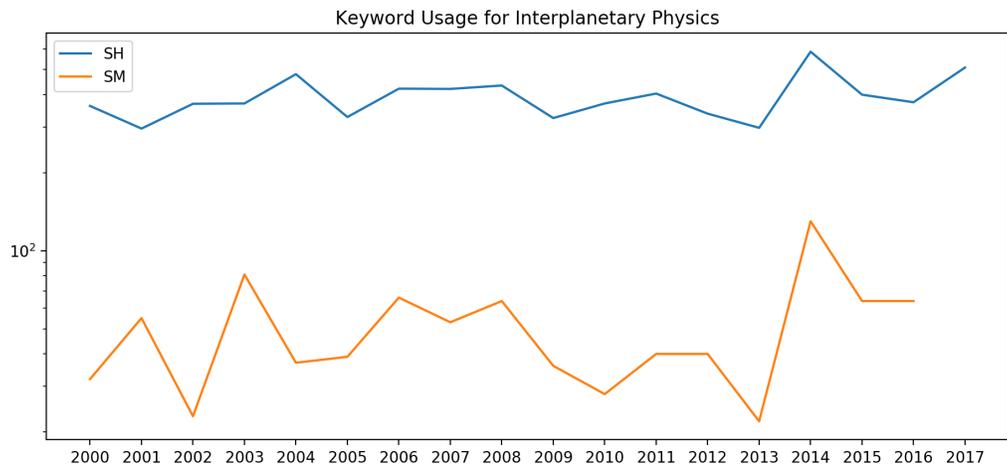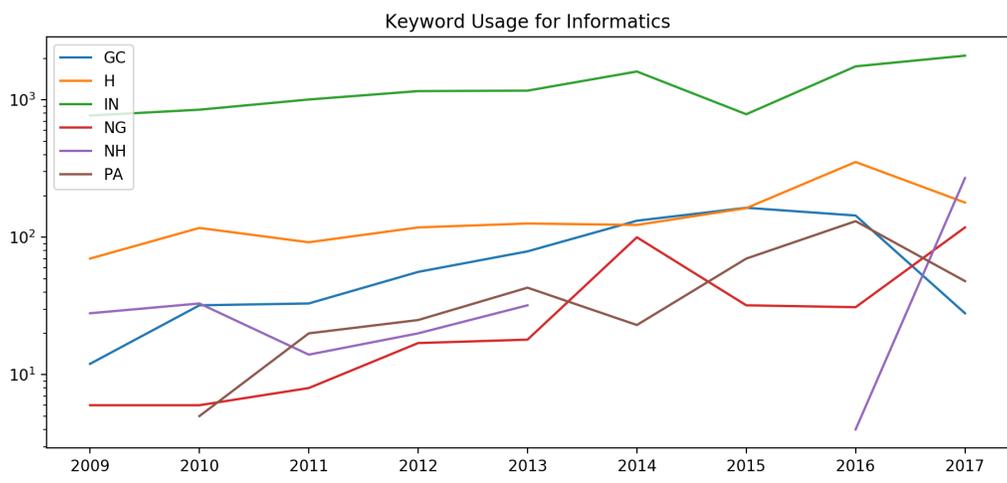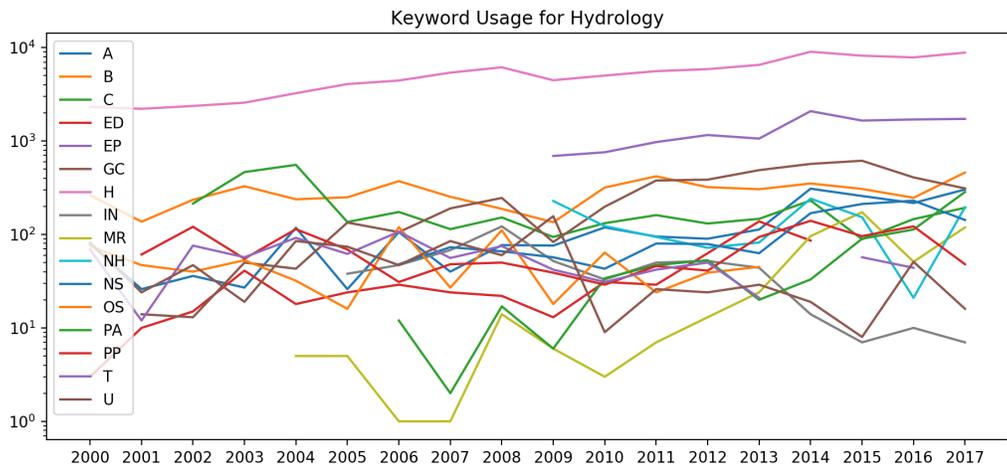
Keyword Usage for Ionosphere
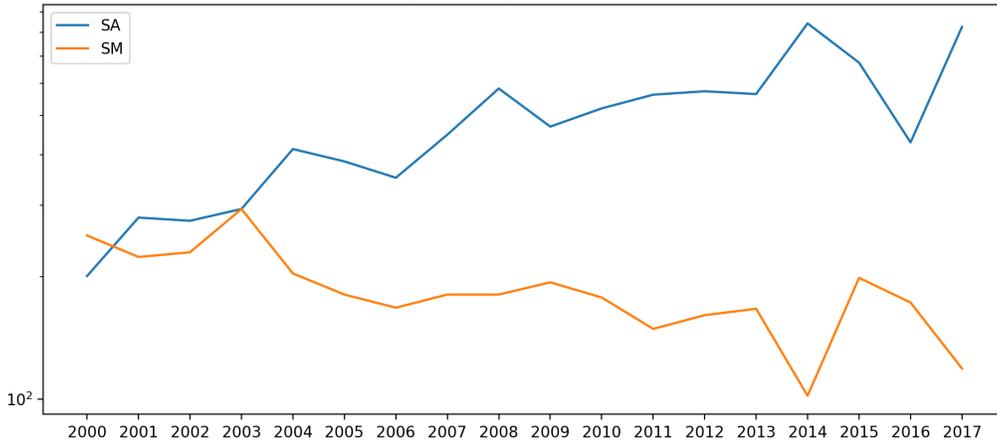
Keyword Usage for Magnetospheric Physics
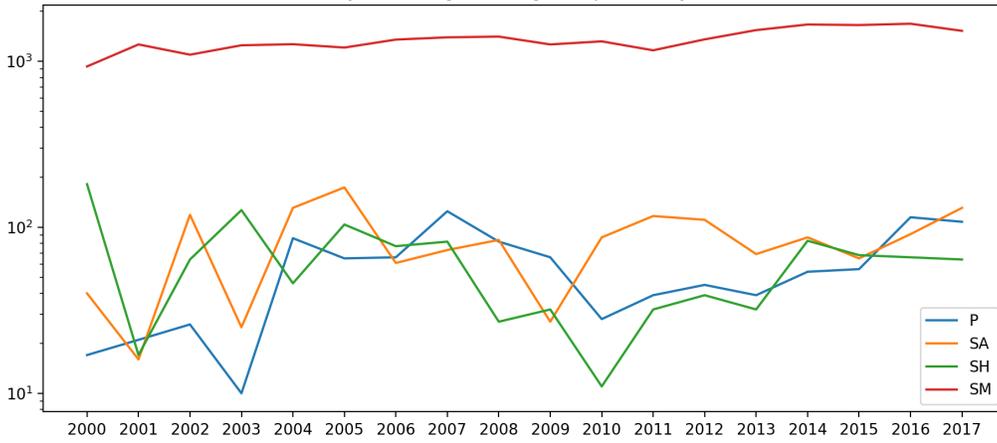
Keyword Usage for Mathematical Geophysics

Keyword Usage for Mineral Physics

Keyword Usage for Mineralogy and Petrology

Keyword Usage for General or Miscellaneous

Keyword Usage for Natural Hazards

Keyword Usage for Nonlinear Geophysics

Keyword Usage for Oceanography: Biological and Chemical

Keyword Usage for Oceanography: General

Keyword Usage for Oceanography: Physical

Keyword Usage for Paleoceanography

Keyword Usage for Physical Properties of Rocks

Keyword Usage for Planetary Sciences: Fluid Planets
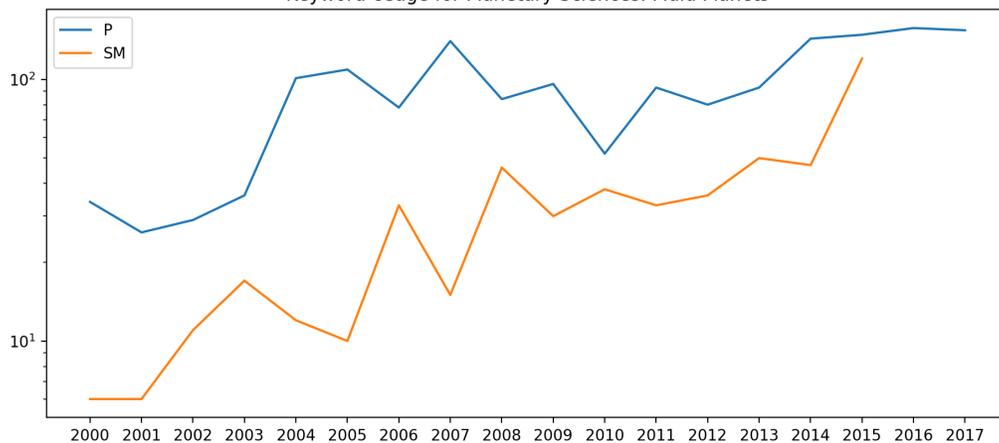
Keyword Usage for Planetary Sciences: Comets and Small Bodies

Keyword Usage for Planetary Sciences: Solar System Objects

Keyword Usage for Planetary Sciences: Solid Surface Planets

Keyword Usage for Policy Sciences
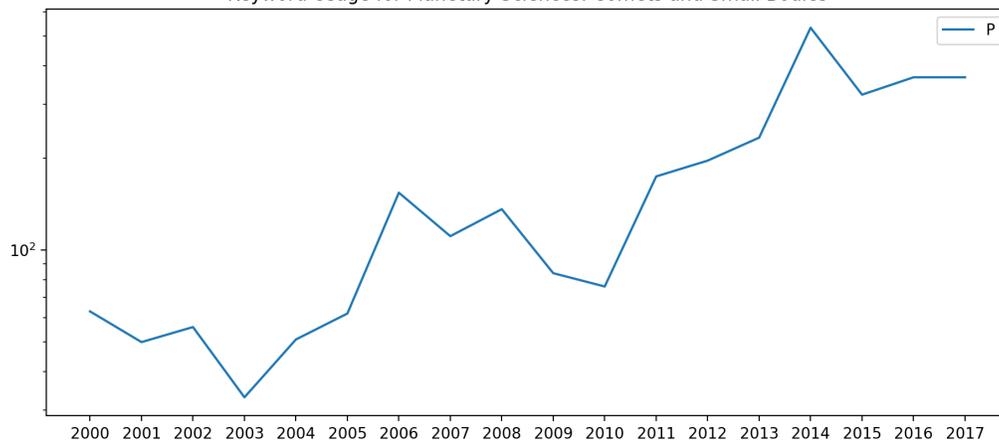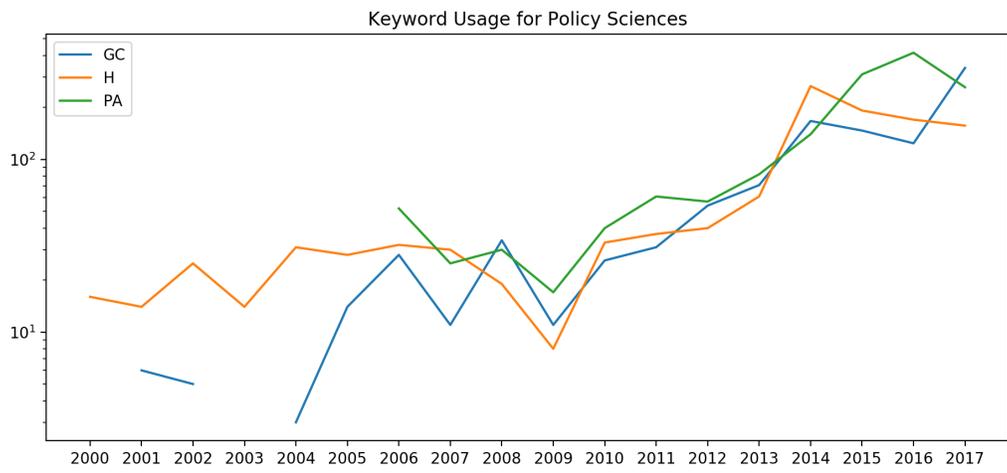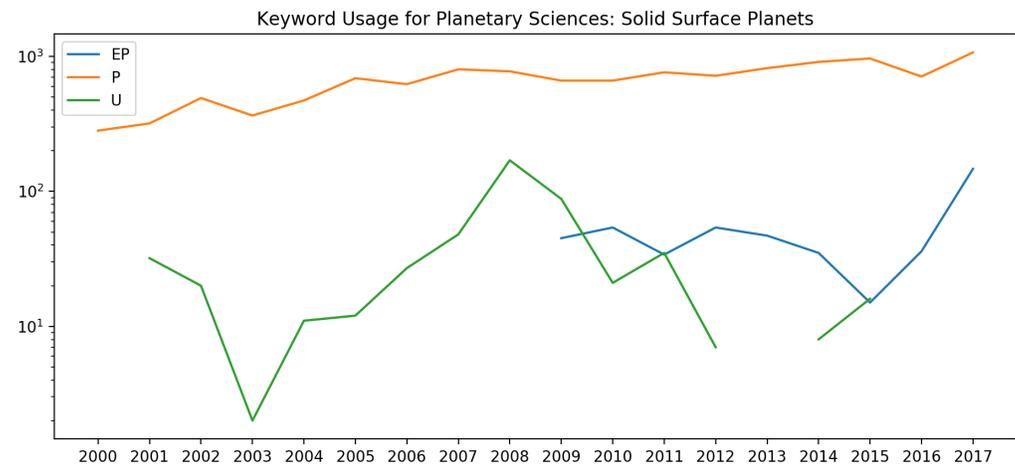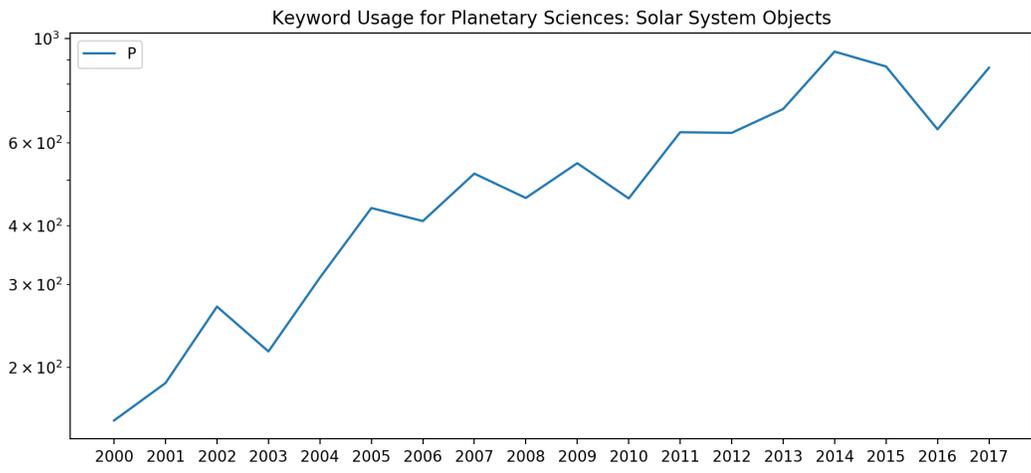
Keyword Usage for Public Issues
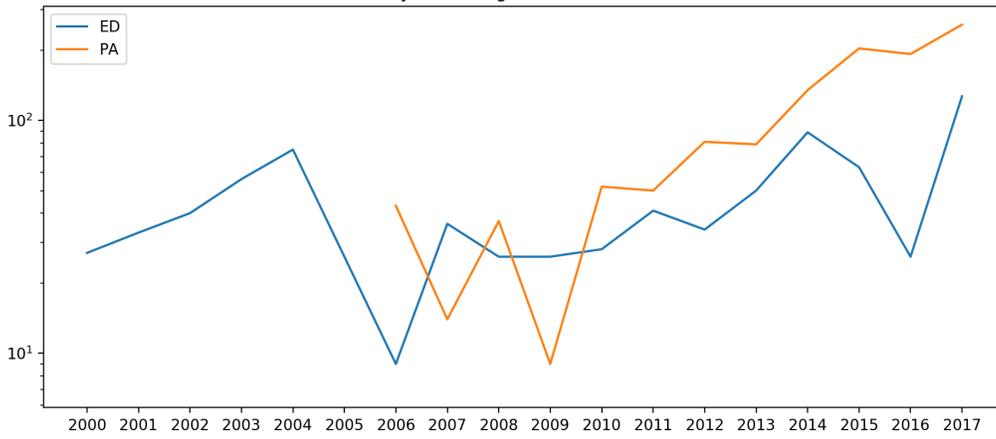
Keyword Usage for Seismology
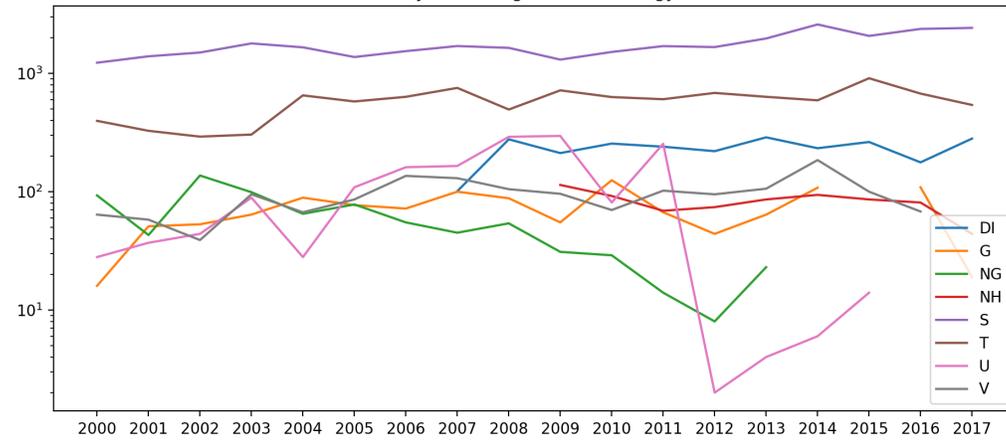
Keyword Usage for Solar Physics, Astrophysics, and Astronomy
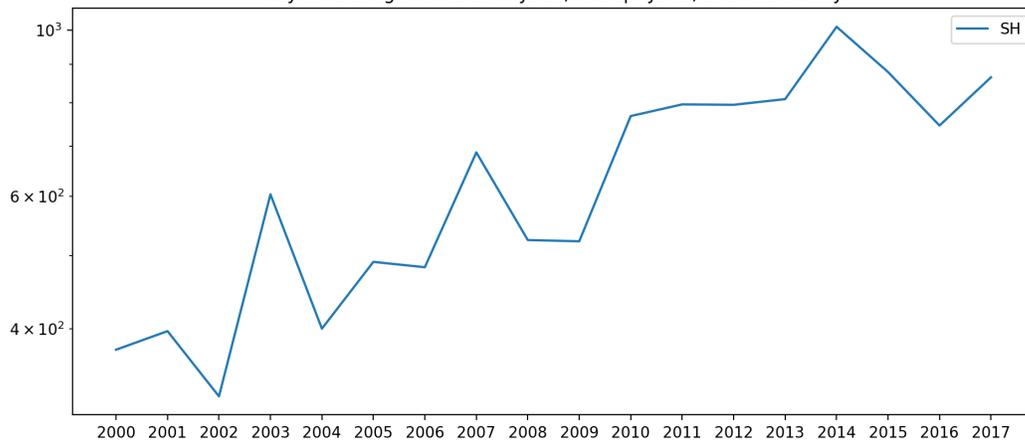
Keyword Usage for Space Plasma Physics
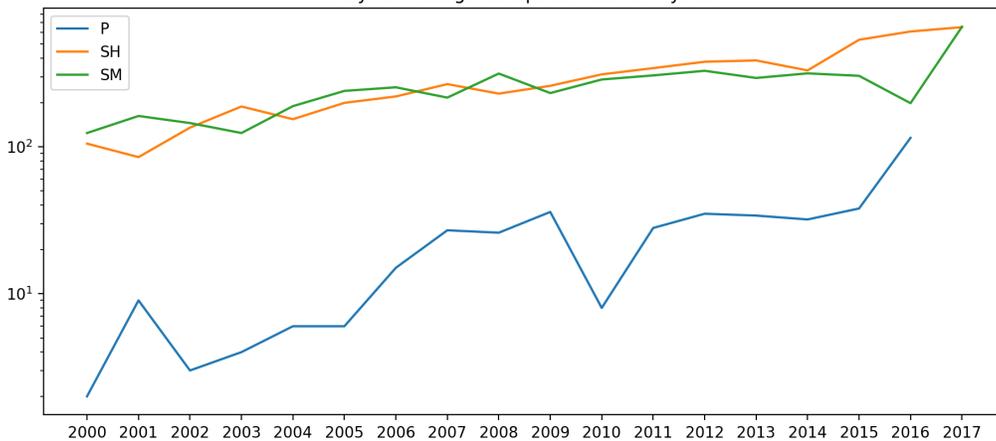
Keyword Usage for Space Weather
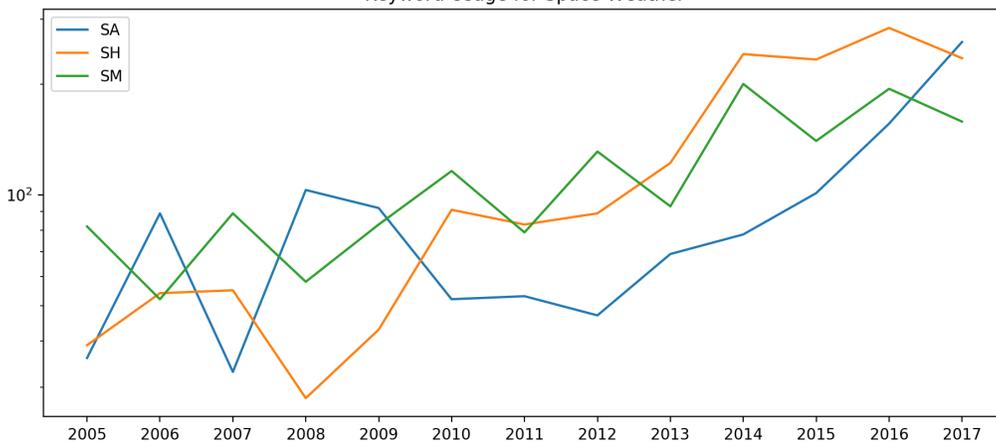
Keyword Usage for Structural Geology
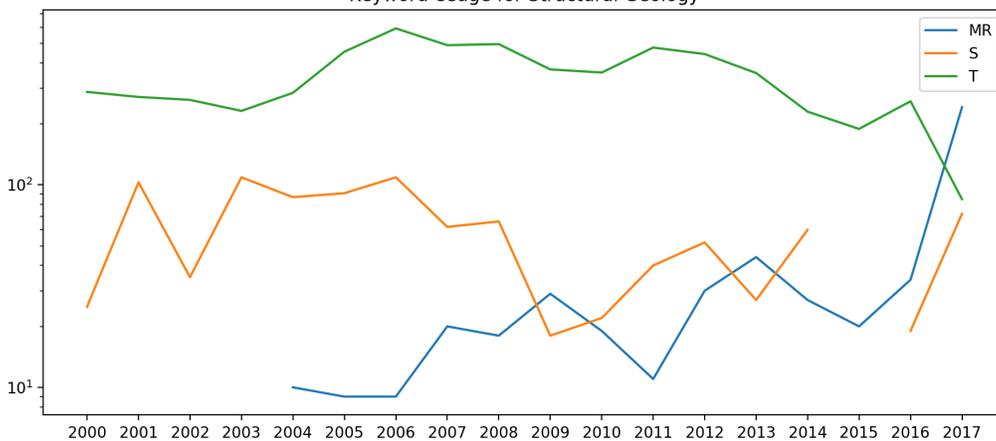
Keyword Usage for Tectonophysics
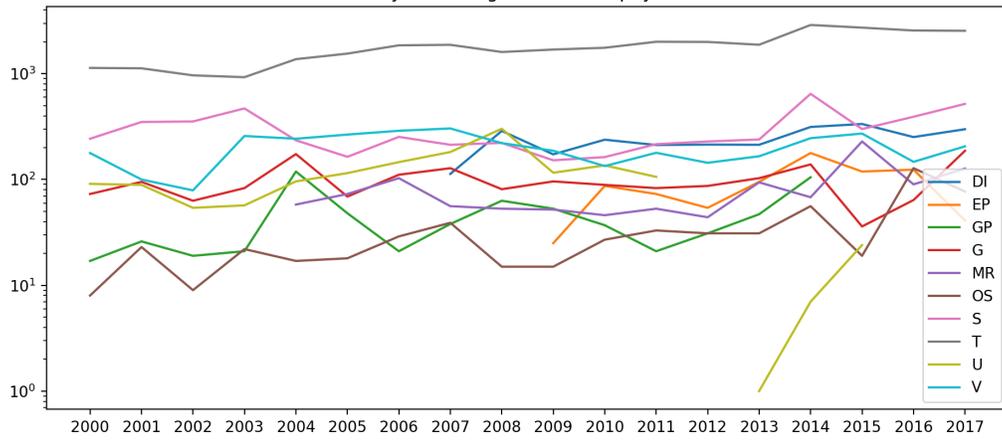
Keyword Usage for Volcanology