

Visualization of Publication Impact

Eamonn Maguire¹, Javier Martin Montull¹, & Gilles Louppe²

¹CERN, Geneva, Switzerland

²New York University, New York, USA



Figure 1: A) A publication impact graph uses time on the X axis, and citations on the Y axis to visualise a publication in context with the references for that paper, and the papers that have cited it. B) Glyphs representations are a more compact version of A) with the citation histogram distinguish by its own section in the glyph. C) Visualise the publication space for an author, institution, or subject area by placing many impact glyphs in 2D space. Citation lines are truncated and expandable with interaction.

Abstract

Measuring scholarly impact has been a topic of much interest in recent years. While many use the citation count as a primary indicator of a publications impact, the quality and impact of those citations will vary. Additionally, it is often difficult to see where a paper sits among other papers in the same research area. Questions we wished to answer through this visualization were: is a publication cited less than publications in the field?; is a publication cited by high or low impact publications?; and can we visually compare the impact of publications across a result set? In this work we address the above questions through a new visualization of publication impact. Our technique has been applied to the visualization of citation information in INSPIREHEP (www.inspirehep.net), the largest high energy physics publication repository.

1. Introduction

While a publications impact is currently viewed at the level of the number of citations, this metric can be misleading, with self-citations for instance artificially driving up an articles perceived importance. Moreover, the weight of a citation (how many times that paper has been cited) can vary but is not immediately available from any existing user interface or visualization tool. Having a way to represent the impact of a publication, not only in the number of citations it received, but how important each of those citations

was would provide an opportunity to assess a publications impact in context with related papers in the domain.

In this work we address the challenge of visualizing publication impact through a novel visualization that can exist in three states as shown in Fig. 1: as a standalone graph (see Fig. 1A); as a glyph design to provide overview level information about a publications impact (see Fig. 1B); and as an informative publication landscape composed of a set of the aforementioned glyphs [BKC*13] (see Fig. 1C).

We utilize a comprehensive citation dataset from the largest curated high energy physics repository, INSPIREHEP. INSPIREHEP maintains its own high quality citation engine to determine accurate citation counts for publications in high energy physics.

Our visualizations have been implemented in D3.js and are open source and available at <https://git.io/vayY0>.

The remainder of this paper will be organized in to five sections: **related work** in Section 2 where we detail the most relevant related work; **design** in Section 3 outlines the design processes involved in creating the visualizations in Fig. 1; the **implementation** in Section 4; **future work** in Section 5; and conclusions in Section 6.

2. Related Work

Related work is sub divided in to: 1) visualization of citation networks; and 2) visualization of publication impact.

2.1. Visualising Citation Networks

As a natural fit to the data, network and more recently matrix-based techniques dominate the approaches used to visualize this type of data [GFV13]. Publications are generally represented as nodes in a graph which can be colored by their subject area or publication venue for example. Directed edges indicate when a publication references another. The citation count for a paper is computed through consideration of the number of incoming edges to a node.

CitNetExplorer by van Eck and Waltman [vEW14] is an example of a citation visualization tool that utilized more graph-based approaches. *Citeology* by Matejka *et al* [MGF12] provides a context driven approach to visualizing a papers citations by arranging each reference and citation along the X axis by year of publication. On the Y axis, each paper is represented by its title. On hovering over a publication title, all references and citations can be viewed as a pathway. Additionally, Noel *et al* [NCR03] devised a technique using minimal spanning trees to visualize co-citations and correlations between authors.

As is often the case with network visualization techniques, when the network becomes large, what is termed a ‘hairball’ can form where nothing is visible anymore. Other techniques have been developed to navigate this issue. *CiteVis* [SCH*13] for example uses a matrix to view papers. *CiteRivers* [HHK16] is a powerful tool for the visual exploration of citation patterns that features the use of streams [HHWN02].

Finally, *Hive Plots* [KBJM11] are a technique that could be used to reduce the visual complexity of large networks to make it easier to view within discipline/field citations (e.g. publications within high energy physics) or citations from external fields (e.g. citations from papers in high energy physics to mathematics).

2.2. Visualising Publication Impact

Publication impact is typically visualized by looking at the number of citations received by the paper. Visualization tools such as

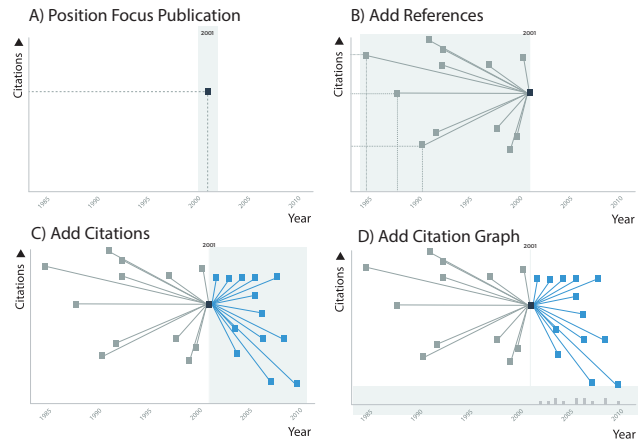


Figure 2: A) First we position the publication of focus by its publication date and citation count. B) We add each reference, again by its publication date and citation count and connect each node with an edge. C) We repeat the process in B for citations. D) A histogram is added to show the number of citations per year/month.

Paperscape[†] use the area of a circle to represent publication importance. Citation networks typically represent the impact of a publication by its connectedness in the graph. An alternative, but effective representation is to plot a publication by its citation count (y axis) and time (x axis). *Altmetric*[‡], a publication impact tracking service that considers social media shares, views, addition to citation management tools, and so on uses 2D plots to communicate the impact of a publication.

3. Design

There are numerous tools and techniques already available for the visualization of citation networks. However, they focus primarily on authors or research subjects, and don’t make it possible to compare the impact between publications. The motivation of this work is based on INSPIREHEP user requests to devise a solution that answers the following questions: 1) is a publication cited less than publications in the field?; 2) is a publication cited by high or low impact publications?; and 3) can we visually compare the impact of publications across a result set?

The aim of our design is to take into consideration the questions posed in Section 3 to provide a visualization that can deliver important information to users across different resolutions. The design takes the form of three interconnected parts: 1) impact graphs (detailed information for one paper); 2) impact glyphs (compact versions of the impact graph); and 3) impact overviews (where we position many impact graphs for a subject area, author, institution, etc.).

[†] Paperscape <http://paperscape.org/>

[‡] Altmetric <http://www.altmetric.com>

3.1. Impact Graphs

Given the questions the users wished to answer, and the information available, we started with the idea of an impact graph that can provide a way of visualizing a focus publication, its references, and citations. The impact graph is composed in the step wise way shown in Fig. 2.

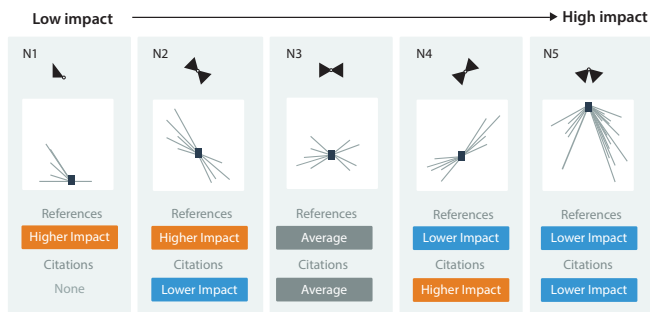


Figure 3: Motifs showing general patterns that can be used to identify publications of varying impact.

Through the topological arrangement of a publications references and citations, it should be possible to define motifs, or frequently occurring patterns that correspond to publications of varying impact within their sphere of influence.

In Fig. 3 we show a selection of topologies defined from the observation of a large corpora of citation data. We have identified five motifs that we believe adequately represent the common citation patterns for publications. Each identified topological arrangement can point to papers of various levels of importance in their field depending largely on the citation counts of references and citations. For example, a publication may have a low number of citations, but the impact of that paper could be considered greater if those citing papers had high citation counts of their own (see Fig. 3 N4).

Conversely, a paper with a high number of citations may appear to be a high impact impact publication, however if all those citing papers have been cited less or if there are a large number of self citations, then the actual impact of the publication should be considered lower (see Fig. 3 N2).

3.2. Impact Glyphs

As shown in Fig. 4, *Impact graphs* can be condensed in to *impact glyphs* to show the general importance of a paper, number of citations and references (and their impact). Their design considered the requirement to show the important features of an impact glyph even at low resolutions.

We tested our design to ensure that key information such as the topological arrange of citations and references, citation density, and self-citations, could be seen at low resolutions. Crush tests as introduced by Maguire *et al* [MRSS* 12] allow for such comparisons in glyph designs. Shown in Fig. 5, our glyph design has been subjected to crush tests from 80 pixels down to 20 pixel wide glyphs. At 80 to 40 pixels, all important information is available. Even high spatial frequency information such as that encoded in the citation

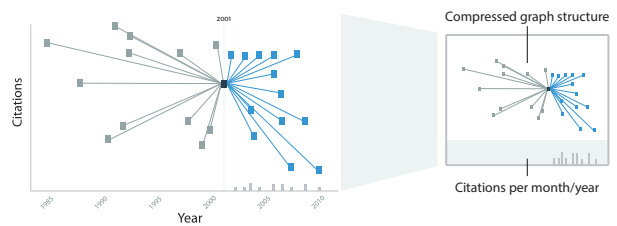


Figure 4: Impact glyphs are impact graphs but without the axes.

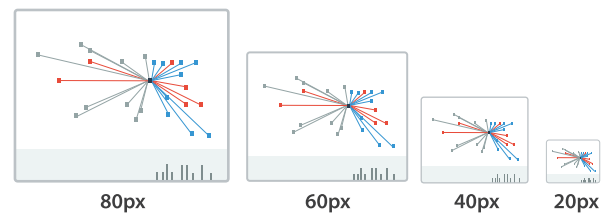


Figure 5: Crush tests are a way of checking to ensure the key information is displayed even at low resolutions.

graph is visible down to 40 pixels. At 20 pixels, the topological arrangement is still evident showing a fairly average publication impact among the scope of related papers.

3.3. Impact Overviews

Finally, impact overviews provide a way of viewing many papers from a subject area, author, collaboration, institution, etc. in a condensed view. Illustrated in Fig. 6, they take the core concepts from the impact graph and impact glyphs, but provide a way to layout the glyphs in 2D space in context with other publication impact glyphs.

These overviews are constructed through the use of a modified glyph that shows much the same information as for impact glyphs, however the edges aren't always drawn to the exact point in a graph that a reference or citation exists. Instead, we draw a line that matches the citation count of the publication, however the time element is scaled in an attempt to avoid overlaps with other glyphs. We are aware that with a large number of publications that there could be overcrowding. To avoid this, we provide the option to change the transparency of glyphs so that the effect of overlaps is reduced.

4. Implementation

Our designs have been implemented in D3.js [BOH11] and use a simple JSON data format. All three modes of operation, to create publication impact graphs, impact glyphs, and impact overviews can be accessed from one library, and use the same overarching data format (multiple network definitions are consumed for impact overview visualizations). The library is easily installable through *bower* via the *impact-graphs* package.

Our publication impact graphs, glyphs, and overview visualizations are being added to the new INSPIREHEP platform which will be released in the coming months where it will be used to visualize over 1.1 million publications. We have run our visualization

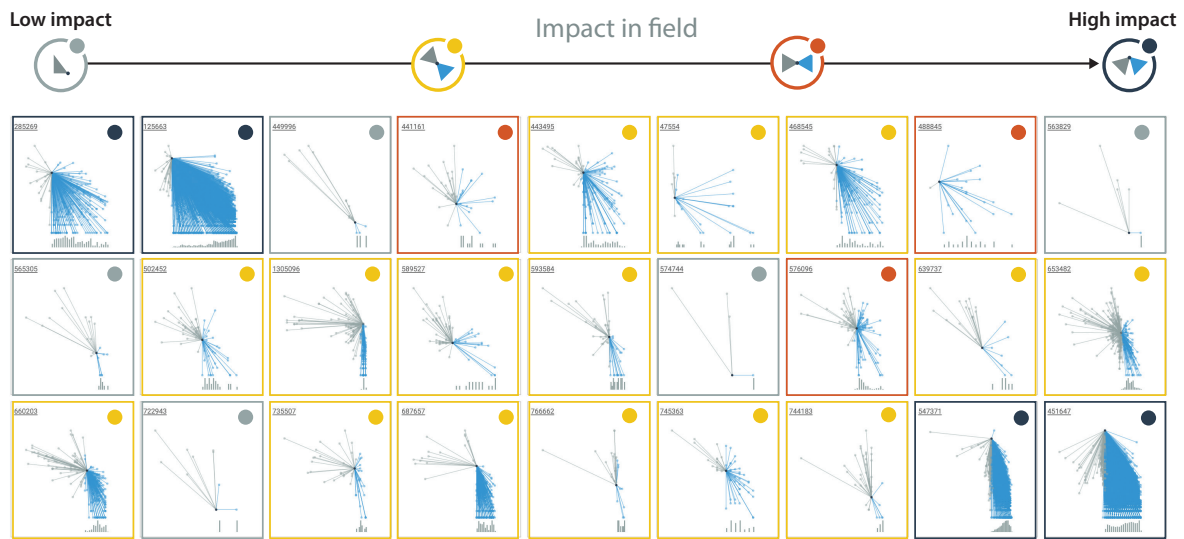


Figure 7: Glyphs created for numerous INSPIREHEP records show a number of impact motifs with low impact compared to their publication sphere, average impact, and high impact.

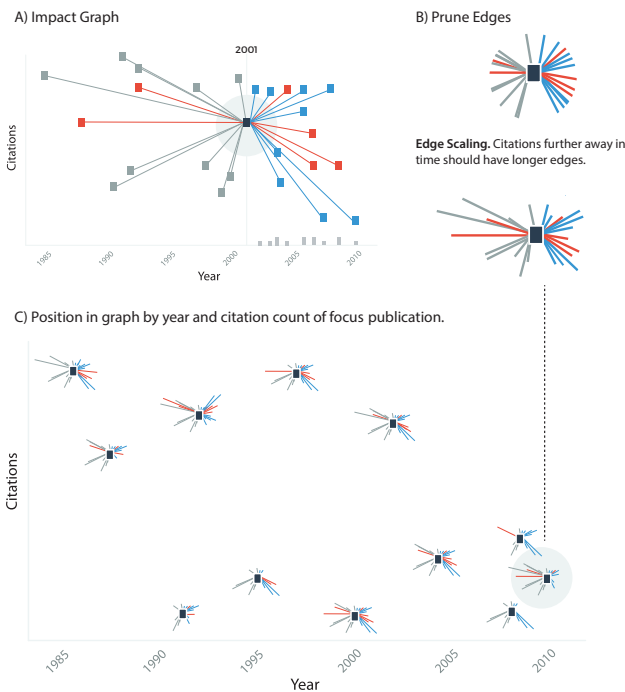


Figure 6: A) We take a standard impact graph as a first step. B) Edges are truncated to avoid overlap with the edges of other publication items. Edge length is scaled to maintain the concept of publication date. C) Each glyph is positioned on a graph area spanning the minimum and maximum publication dates and citation counts.

approach over many thousands of publications to produce output such as that shown in Fig. 7 where even in this this small subset, many of the motifs identified in Fig. 3 can be observed.

To exemplify the scalability of the approach, the glyph in the

bottom right of Fig. 7 is visualizing publication 451647 from INSPIREHEP (*The Large N limit of superconformal field theories and supergravity*) which has over 11,000 citations. Rendering speed is also important since we envisage these visualizations being optionally shown in search result pages. With thousands of publications from INSPIREHEP, we have observed rendering speeds of < 10ms for records with less than 500 citations. For 11,000 citations, rendering takes 400ms.

Finally, our library comes with many options to enable configuration of: the Y scales from log to linear; the minimum and maximum citation counts and years (to facilitate easier between-glyph comparison); and automatic anomaly detection to highlight references made after or citations made before the publication date. Such errors can point to issues with multiple versions of the same publication record.

5. Future Work

With much of the functionality already present, future work will focus on an evaluation. Our visualizations have been designed with feedback from day to day users of INSPIREHEP, however we do not assume that the encoding will be immediately familiar to all. So far, our experience shows that users understand the visualization after a short introduction. A full scale user evaluation will help to confirm this across the INSPIREHEP user base.

6. Conclusion

We have presented a new glyph design for the visualization of publication impact. We have provided an implementation that can be immediately incorporated in to existing digital libraries for interactive use either as dedicated visualizations of a papers publication impact (impact graphs), as glyphs to accompany search results, or to be used as mass summarizations of publication impact across a database.

References

- [BKC*13] BORGIO R., KEHRER J., CHUNG D. H., MAGUIRE E., LARAMEE R. S., HAUSER H., WARD M., CHEN M.: Glyph-based visualization: Foundations, design guidelines, techniques and applications. *Eurographics State of the Art Reports* (2013), 39–63. [1](#)
- [BOH11] BOSTOCK M., OGIEVETSKY V., HEER J.: D³ data-driven documents. *Visualization and Computer Graphics, IEEE Transactions on* 17, 12 (2011), 2301–2309. [4](#)
- [GFV13] GIBSON H., FAITH J., VICKERS P.: A survey of two-dimensional graph layout techniques for information visualisation. *Information visualization* 12, 3–4 (2013), 324–357. [2](#)
- [HHK16] HEIMERL F., HAN Q., KOCH S.: Citerivers: visual analytics of citation patterns. *Visualization and Computer Graphics, IEEE Transactions on* 22, 1 (2016), 190–199. [2](#)
- [HHWN02] HAVRE S., HETZLER E., WHITNEY P., NOWELL L.: The-meriver: Visualizing thematic changes in large document collections. *Visualization and Computer Graphics, IEEE Transactions on* 8, 1 (2002), 9–20. [2](#)
- [KBJM11] KRZYWINSKI M., BIROL I., JONES S. J., MARRA M. A.: Hive plots —rational approach to visualizing networks. *Briefings in bioinformatics* (2011), bbr069. [2](#)
- [MGF12] MATEJKA J., GROSSMAN T., FITZMAURICE G.: Citeology: visualizing paper genealogy. In *CHI'12 Extended Abstracts on Human Factors in Computing Systems* (2012), ACM, pp. 181–190. [2](#)
- [MRSS*12] MAGUIRE E., ROCCA-SERRA P., SANSONE S.-A., DAVIES J., CHEN M.: Taxonomy-based glyph design — with a case study on visualizing workflows of biological experiments. *Visualization and Computer Graphics, IEEE Transactions on* 18, 12 (2012), 2603–2612. [3](#)
- [NCR03] NOEL S., CHU C.-H. H., RAGHAVAN V.: Co-citation count vs correlation for influence network visualization. *Information Visualization* 2, 3 (2003), 160–170. [2](#)
- [SCH*13] STASKO J., CHOO J., HAN Y., HU M., PILEGGI H., SADANAAND R., STOLPER C. D.: Citevis: Exploring conference paper citation data visually. *Posters of IEEE InfoVis* (2013). [2](#)
- [vEW14] VAN ECK N. J., WALTMAN L.: Citnetexplorer: A new software tool for analyzing and visualizing citation networks. *Journal of Informetrics* 8, 4 (2014), 802–823. [2](#)