This is a working paper. Please do not distribute or share publicly.

Funding the Frontier: Visualizing the Broad Impact of Science and Science Funding

Yifang Wang[1,2,3,4,5], Yifan Qian[1,2,3,4], Xiaoyu Qi[6], Yian Yin[7], Shengqi Dang[6], Ziqing Qian[6], Benjamin F. Jones[1,2,3,4,8], Nan Cao[6*], Dashun Wang[1,2,3,4,9*]

[1]Center for Science of Science and Innovation, Northwestern University, Evanston, IL, USA
[2]Northwestern Innovation Institute, Northwestern University, Evanston, IL, USA
[3]Ryan Institute on Complexity, Northwestern University, Evanston, IL, USA
[4]Kellogg School of Management, Northwestern University, Evanston, IL, USA
[5]Department of Computer Science, Florida State University, Tallahassee, FL, USA
[6]Intelligent Big Data Visualization Lab, Tongji University, Shanghai, China
[7]Department of Information Science, Cornell University, Ithaca, NY, USA
[8]National Bureau of Economic Research, Cambridge, MA, USA
[9]McCormick School of Engineering, Northwestern University, Evanston, IL, USA
[*]Correspondence to: nan.cao@tongji.edu.cn, dashun.wang@northwestern.edu

Abstract

Understanding the broad impact of science and science funding is critical to ensuring that science investments and policies align with societal needs. Existing research links science funding to the output of scientific publications but largely leaves out the downstream uses of science and the myriad ways in which investing in science may impact human society. As funders seek to allocate scarce funding resources across a complex research landscape, there is an urgent need for informative and transparent tools that allow for comprehensive assessments and visualization of the impact of funding. Here we present *Funding the Frontier* (*FtF*), a visual analysis system for researchers, funders, policymakers, university leaders, and the broad public to analyze multidimensional impacts of funding and make informed decisions regarding research investments and opportunities. The system is built on a massive data collection that connects 7M research grants to 140M scientific publications, 160M patents, 10.9M policy documents, 800K clinical trials, and 5.8M newsfeeds, with 1.8B citation linkages among these entities, systematically linking science funding to its downstream impacts. As such, *Funding the Frontier* is distinguished by its multifaceted impact analysis framework. The system incorporates diverse impact metrics and predictive models that forecast future investment opportunities into an array of coordinated views, allowing for easy exploration of funding and its outcomes. We evaluate the effectiveness and usability of the system using case studies and expert interviews. Feedback suggests that our system not only fulfills the primary analysis needs of its target users, but the rich datasets of the complex science ecosystem and the proposed analysis framework also open new avenues for both visualization and the science of science research.

This is a working paper. Please do not distribute or share publicly.

1 Introduction

Scientific progress is crucial to human well-being and prosperity. Advances in science have improved global health [1–3], driven innovation and economic growth [4, 5], and informed policy decisions [6, 7]. From mRNA vaccines to the Internet, scientific progress improves quality of life and standards of living, widely seen as a vital public good [8–10]. Essential drivers of scientific progress are funding and public support. Vannevar Bush highlighted this connection in his landmark 1945 report, *"Science, The Endless Frontier,"* advocating for federal investment in scientific research as *"the fund from which the practical applications of knowledge must be drawn"* [8]. This report led to the creation of the National Science Foundation (NSF) in 1950 and established the basis for modern science policy.

While there is now growing evidence of the strong relationship between science funding, basic science, and its societal impact at a high level [10, 11], a holistic and systematic approach to evaluating the broad outcomes of specific funding decisions remains scarce. Given the growing importance of science funding, the ability to evaluate the multidimensional impacts of this funding is critical to ensure that science policies and investments align with social needs and to maintain public trust in scientific endeavors [12]. These evaluations are of immense interest to a diverse array of stakeholders, including funding agencies seeking effective science investment strategies, national policymakers pursuing opportunities to enhance national competitiveness through science and technology policies, university leaders planning strategic initiatives, and individual scientists formulating impactful research programs. However, the analysis required for this work is complex and challenging due to both the multifaceted nature of funding outcomes and the extended time period that passes between the initiation of projects and their potential societal impact.

Much research has emphasized the importance of science funding [13–15]. Yet, the bulk of existing research concentrates on grants and their resulting scientific publications, which evaluate the impact of funding within science, while ignoring other facets of the impacts. Emerging efforts examine the impacts of scientific research more broadly, but they rarely trace these outcomes back to their funding sources [1,4,7,16]. The visualization community has also long been interested in studying science funding, using, for example, grant data for topic mining and principal investigator (PI) network visualization [17, 18], yet it, too, has largely overlooked the broad impact of funding. As a result, there is a notable lack of quantitative analyses and accessible analytical approaches that science decision-makers can utilize effectively, as well as a substantial gap in our understanding of the full landscape of the influence of funding on science and society.

In recent years, extensive new datasets have emerged that capture not only the inner workings of the scientific enterprise in detail and at scale, including funding [19] and scientific outputs, but also trace the numerous ways in which science interfaces with the broader human society, ranging from healthcare [20] to government policies [21] to marketplace applications [22] to public perceptions [23], representing science as an interconnected ecosystem (Fig. 1). To truly unleash the potential of these datasets and to maximize the impact of science on society, however, we need new toolkits to analyze and understand the information that is now available to us. Here we present *Funding the Frontier* (*FtF*), a new visualization tool that combines the latest advances across two fields—the science of science (SciSci) [15, 24] and visual analytics [25]. Indeed, SciSci methods enable us to link various data sources to create a holistic and systematic approach to assessing and predicting the broad impact of funding, and visual analytics solutions empower us to deliver accessible insights to policymakers and institutional leaders directly. Overall, the system builds on a massive data collection that connects 7M research grants to 140M scientific publications, 160M patents, 10.9M policy documents, 800K clinical trials, and 5.8M newsfeeds, with 1.8B citation linkages among these entities, representing to our knowledge the largest and most comprehensive data aggregation of science funding and its downstream impacts. By assessing the broad impact of science and science funding, the system allows for more targeted investments, an improved ability to relieve roadblocks in the scientific

process, and, overall, a greater impact of science on society.

There are several challenges involved in designing such a system. First, impact assessment requires rational definitions and accessible metrics for evaluating the multidimensional impacts of funding on society [12]. These measures need to capture the wide range of potential outcomes of scientific research, while being straightforward and comprehensible to decision-makers who often lack technical backgrounds. Second, the diversity and volume of the data create a need for especially innovative visualization solutions to allow for effective exploration. The data not only cover a wide range, from grant information to scientific outcomes and their downstream impacts, but also introduce a high level of complexity, as they are characterized by heterogeneous citation networks, multidimensional variables, hierarchical topics, and temporal dynamics. Third, a system that can offer meaningful investment recommendations and guide future investment strategies requires a predictive engine that can analyze the massive amount of historical data on funding and its outcomes and suggest high-impact opportunities.

We began to tackle these challenges by interviewing multiple experts, including both decision-makers and active researchers, to characterize the domain problem. We then defined a set of measures to quantify the impact of funding from multiple dimensions and developed a prediction model to identify grant topics and research investigators with high potential to produce substantial impacts in the future. Using these historical data and prediction results, we designed and developed *Funding the Frontier* (*FtF*), a comprehensive visual analysis system that presents, analyzes, and predicts the multi-faceted impacts of science and science funding for a wide range of decision-makers in science, including funders, policymakers, university leaders, and more. The contributions of this work are summarized as follows:

- We characterize the problem domain of visual analytics for funding impact and propose a list of impact measures.
- We develop *FtF*, a comprehensive visual analysis system with novel designs and effective interactions to explore multi-dimensional funding impact. The system is available through the website: https://fundingthefrontier.com.
- We conduct an exploratory evaluation of the effectiveness and usability of the new system, including a quantitative study, two case studies, and expert interviews.

2 Related Work

Related literature spans three domains, including the science of science, funding data visualization, and network visualization.

2.1 The Science of Science

The SciSci community has long been interested in studying science funding and its outcomes. The early literature mainly focuses on the funding data itself, without considering its outcomes, revealing for example the racial and gender bias in grant applications [26, 27]. To the extent that funding outcomes are examined, the bulk of research has focused on the impact of funding within science itself, studying for instance how funding correlates with the amount and impact of papers that result [9, 28–31], or the impact of grants on subsequent career outcomes in science [9, 32, 33].

Separately, there is another line of inquiry that aims at quantifying the broad impacts of science. While these studies help us better understand how science interfaces with various aspects of human society, including marketplace applications [4], public perception [16, 23], clinical trials [1–3], and policy-making [6, 7, 10] (Fig. 1), they do not trace the science back to its upstream funding, highlighting the empirical gap linking funding to its broad impacts.

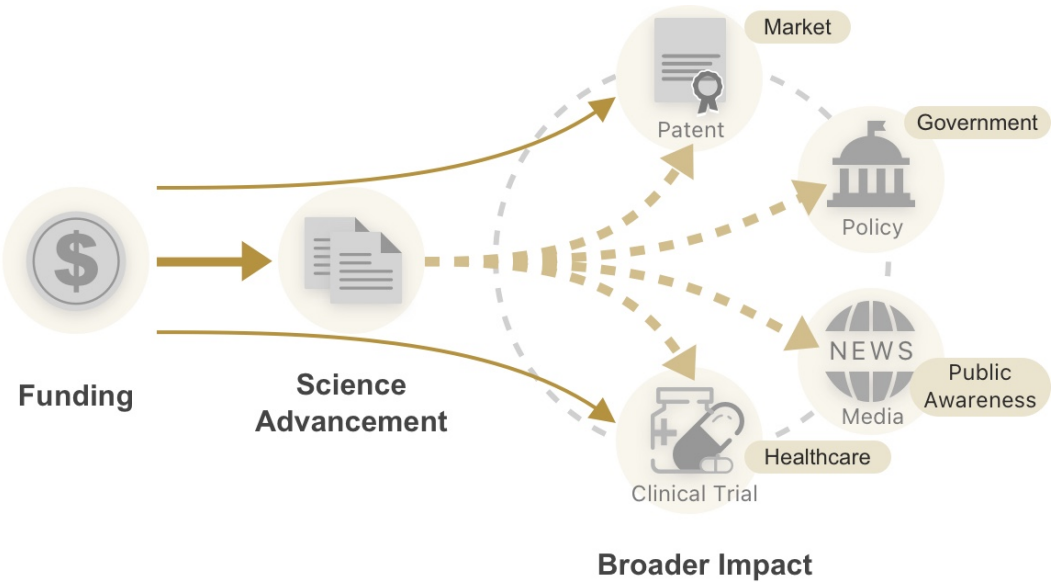


Figure 1: The science ecosystem, from the upstream funding to the science to the broader downstream impact. The direct outcomes of funding include scientific publications, patents, and clinical trials (solid dark yellow arrows), while the broader impacts emerge through the influence of funded papers on the broader society, such as public policy (dashed shallow yellow arrows).

More recently, there have been some initial attempts at linking funding to its broad impacts, which substantially inform our work. For example, Yin et al. consider the broad impacts of public funding by examining three types of public uses of science (i.e., policy, news media, and patents) [10]. Other studies link government funding to resulting patents, highlighting the impact of funding on marketplace applications [5, 34]. While these papers rely on datasets that are more limited in scope, scale, and detail, compared with our datasets, they are similar to our work in that they address the relationship between funding and its broad impact. More importantly, while these studies are descriptive by nature, the policy relevance of the insights derived from these studies highlights the need for a visual analytical approach to understanding the broad impact of the funding, which is the focus of the present paper. Indeed, the key differentiation is that these existing studies are descriptive by nature, and do not allow decision-makers, such as program officers at funding agencies, to explore the data or draw insights. To the best of our knowledge, the visual analytics system we propose here is the first user-oriented, interactive system for science decision-makers to analyze and visualize the multidimensional impacts of funding.

2.2 Scientific and Funding Data Visualization

The visualization community has widely used and analyzed scientific data on papers, patents, and grants. Previous literature mainly falls into two categories: (1) technique-centered work and (2) insight-centered work. Using the complex data structure of scientific data, which includes network [35,36], time-series [37,38], and multivariate [39] datasets, technique-centered studies often use these data to propose general visualization techniques. Insight-based studies, by contrast, either use comprehensive visualization systems [40–43] or are based on datasets that are developed to distill insights about the development of the visualization and HCI community itself [44–47].

More specifically, a number of studies have focused on science funding data. For example, DIA2 [48] visualizes the NSF funding portfolio. Liu et al. [49, 50] proposed a 2.5D treemap to visualize funding distributions. Scholar Plot [51] visualizes researchers' publication and funding data. The Public Innovations

Explorer [52] visualizes funding from federal agencies and small businesses across the US to explore innovative geolocations. The EnArgus system [53] provides semantic search functions for grant abstract data to focus on energy-related projects. And there are online platforms, including NIH RePORTER [54,55], Altmetric [23], PlumX Metrics [56], and Dimensions [19], which serve primarily as search engines that show grants-related meta-information and their outputs. Moreover, grant data have also been used to inform general visualization challenges, such as graph visualization using multidimensional grant metadata [57,58] and topic modeling and evolution using grant abstract [18,59]. Related, there is a body of visualization work that analyzes funding outcomes in science [17,60,61].

On one hand, these studies and systems highlight the importance of studying the impact of funding, but they have mainly focused on the impact within science, largely ignoring the impacts of funding on the broader human society. On the other hand, systems such as MedChemLens [62] and InnovationInsights [63] explore the interaction between science and its broader impact, including patents and medicine, but they do not consider upstream funding. By contrast, our work characterizes the domain problem of analyzing funding impact and aggregates multiple data sources to develop a first-of-its-kind visualization system that evaluates the broad outcomes of funding portfolios, especially those that go beyond science.

2.3 Network Visualization

Network visualization represents a broad area of study, covering a range of different network types, such as multi-variate [64], group [65], dynamic [66], and multi-layer networks [67]. Here we introduce studies that are most relevant to ours.

The data in the scientific ecosystem can be depicted as a large-scale heterogeneous network comprising various node types (e.g., grants, papers, and researchers) and links (e.g., paper-authorship and grant-paper citations). For large-scale networks, strategies such as aggregation and graph sampling are commonly employed to manage visual complexity, often supplemented with interactions like filtering and zooming. Compound graphs [66] serve to group nodes into categories based on user-defined attributes or clustering methods. For example, OnionGraph [68] aggregates paper nodes into groups based on their citation range. *FtF* also uses node aggregation to present the funding impact at the field level.

When it comes to visualizing networks that include multiple node types, there are broadly four main approaches to visually differentiating node types within a network: embedding, superimposition, juxtaposition, and the use of visual node attributes [65]. To further convey hierarchy within a node type, contour overlays (e.g., circular treemaps [69]) have been shown to be intuitive and effective. However, these approaches often impose uniform data structures on each node type, which may neglect the diversity and the relationships across different node types. In *FtF*, we adopt a hybrid approach that combines node-link diagrams with hierarchical structures and glyph techniques, which is tailored for assessing the multi-dimensional impacts of funding.

3 System Overview

This section presents the analysis requirements and a system overview.

3.1 Analysis Tasks

Conversations with diverse stakeholders in the past few years underscored the urgent need to evaluate the impact of funding to enhance decision-making processes, such as funders' investment strategies. These stakeholders included directors and program officers in private and public funding agencies, leaders at top universities, policymakers, and SciSci researchers. We identified key challenges these decision-makers face, including the absence of comprehensive data on funded programs and their wide-ranging impacts, the lack of systematic approaches to measuring the impact of funding, and the scarcity of user-friendly tools for

those who lack specialized data analysis skills. These discussions highlighted the need for a visual analytics system capable of providing an intuitive, systematic, and efficient approach to evaluating the impacts of funding in order to enhance decision-making.

We began to address this challenge over the past few years, engaging with experts from various sectors. Our experts included a U.S. federal funding agency program officer with over eight years of experience (E_A), a director-level executive in a top private investment firm that specializes in investing in science (E_B), as well as experts in the SciSci field (E_C-E_E). E_C is an established SciSci professor and E_D is a professor who works on statistical analysis of funding data. E_E is a researcher focusing on science and innovation. We also gathered requirements and feedback from data analysis teams at the National Science Foundation (NSF) and the National Institutes of Health (NIH). Based on conversations with decision-makers, discussions with domain experts, and a comprehensive literature review, we leverage the expert-focused design study methodology [70] to identify the six analysis tasks listed below.

Tasks for Project Summarization (S). Profiles of funded projects offer a strategic entry point for analysis, enabling users to examine the investments and adjust strategies accordingly. Therefore, the system should provide both overview and detailed information about the funded projects and support a flexible information exploration approach.

S1: Dynamic Data Selection. The system should enable users to dynamically filter grants by various attributes, such as funding agencies and grant years, to tailor their analysis.

S2: Funding Overview Illustration. The system should provide a summary of funded fields to assist analysts in selecting fields of interest for detailed analysis.

S3: Researcher and Institution Characterization. The system should also characterize (1) the funded researchers (i.e., principal investigator (PI)), in terms of demographics and SciSci measures of research ability, and (2) recipient research institutions. This information supports diversity, equity, and inclusion (DEI) by highlighting the funding distribution among researchers and institutions and identifying PIs and institutions with potential for future investment.

Tasks for Impact Evaluation (E). Another key analysis requirement is the ability to assess the impact of funded projects. As funding agencies have a wide range of missions and investment strategies, the criteria for "success" in different funding portfolios can vary significantly. The analysis should, therefore, present the impact of funding from a multidimensional perspective and provide context to aid users in understanding these outcomes.

E1: Analysis of funding impacts. The system should offer a comprehensive set of easy-to-understand metrics to help analysts evaluate and select the type of impact they value most—ranging from direct scientific impacts to broader societal benefits. Intuitive visualizations and interactions should be integrated to present more detail regarding the impact, which is based on the citation relationships between grants and their outcomes (e.g., papers).

E2: Contextualization of funding impacts. To deepen users' understanding of funding outcomes and enrich the insights from the analysis, the system should go beyond citation relationships to incorporate relevant contextual data, including detailed information about the attributes of specific outcomes, such as the drugs that result from clinical trials or the assignees of the patents.

Tasks for Investment Recommendation (R). The ultimate aim of these evaluations of the impact of funding is to identify promising directions for future investments. Predictive analytics can play a pivotal role in guiding these forward-looking decisions.

R1: Identifying Investment Opportunities. Our system should support a forward-looking capability, allowing users to uncover potential investment opportunities based on their desired impact outcomes for the future. This involves pinpointing emerging research topics and, more importantly, identifying

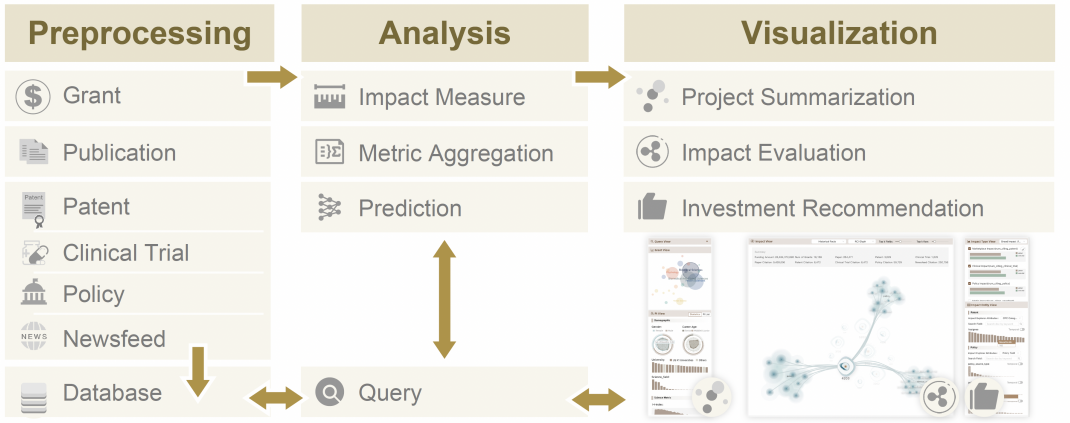


Figure 2: *FtF* system overview. The system consists of a preprocessing module, an analysis module, and a visualization module.

the right people (PIs) for collaboration or funding.

3.2 System Overview

Guided by these analysis tasks, we designed *FtF*, a web-based application that consists of three modules (Fig. 2): (1) the preprocessing module, (2) the analysis module, and (3) the visualization module. The preprocessing module cleans multiple data sources, stores them in the database, and supports dynamic data queries (**S1**). The analysis module provides a list of metrics (**S3**, **E1**) and prediction models to indicate grant topics, principal investigators (PI), and research institutions for future investment (**R1**). These two modules form the backend of the system, which is implemented by Google BigQuery, Python, and Flask. The visualization module uses coordinated views with intuitive interactions to present both historical data on grant outcomes (**S2**, **S3**, **E1**, **E2**) and prediction results (**R1**) to guide future grant-making, implemented by React.js, Redux.js, D3.js, and TypeScript. We introduce each module with more details in the following sections.

4 Data Analysis

The analysis module is responsible for the calculations that serve as the foundation for the visual analytics systems. First, this module assesses and aggregates the multidimensional impacts of science and funding at different levels, from individual papers and grants to broader categories at the PI and field levels. Second, the module predicts future research investment opportunities by estimating the potential impact of recent grants. We present more details regarding the system's data sources, preprocessing methods, funding impact assessments, and future investment recommendations below.

4.1 Data Sources and Preprocessing

This analysis of the broad impact of science funding requires the integration of the following four datasets containing information about the science ecosystem:

- **Dimensions** [19]. This dataset comprises data on global science funding (7M grants), scientific publications (140M), patents (160M), clinical trial records (800K), and their citation linkages (1.7B). It also contains information about researchers (51.7M), including their names and affiliations and their connections to grants (9M, mostly PIs) and paper-author pairs (461M).

- **Overton** [21]. This dataset includes global policy documents (10.9M) and citation linkages between policies and papers (13M).
- **Altmetric** [23]. This dataset provides insights into public engagement with science through newsfeeds (5.8M) and citations between news sources and papers (6M).
- **SciSciNet** [71]. This dataset includes a large-scale open data lake for the science of science research, covering over 134M scientific publications and millions of external linkages to funding and public uses. Specifically, we use a set of SciSci metrics (e.g., h-index and productivity) at the paper and researcher level, which is a great addition to the Dimensions dataset.

We preprocessed the data sources to limit our analysis to all types of documents published between 2000 and 2021 to ensure robust data coverage (see Supplementary Notes 1 and 2).

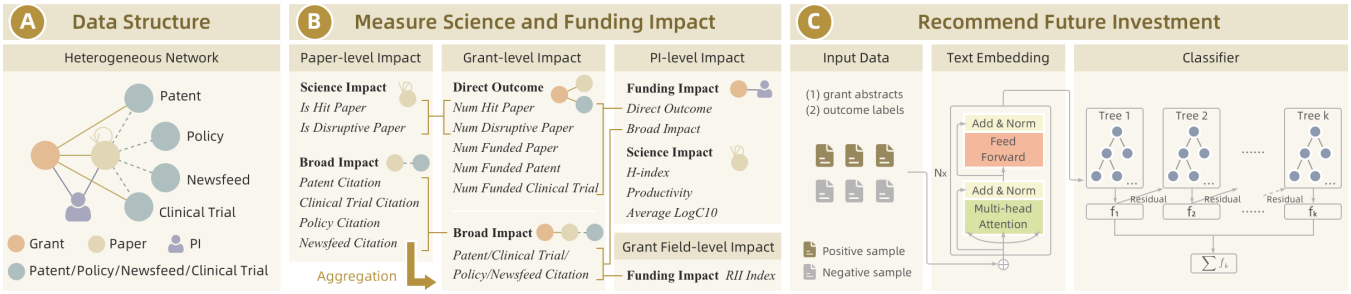


Figure 3: The data analysis procedure. (A) We preprocess the data into a heterogeneous network with multidimensional attributes. (B) We then define and calculate a list of metrics to measure the impact of science and funding at multiple levels. (C) Finally, we use a prediction model to recommend the grant topics and PIs that represent the best opportunities to achieve future impact in a specific impact topic.

4.2 Measure the Funding Impact

The impact of science and science funding is multidimensional and extends through a range of levels, from papers and grants at the lowest level to grant fields and funding agencies at the highest level. In collaboration with experts in the SciSci field, we developed a set of metrics to assess the diverse funding outcomes. These metrics are multidimensional, measuring not only the direct scientific impact of the funding based on publications and citations but also measuring a wide range of broader impacts, encompassing the marketplace, clinical trials, policy, and media. These impacts can be categorized into the direct outcomes of the funding and its broader impacts, which are distinguished by different citation relationships:

Direct Outcomes. The direct impact of funding can primarily be seen in scientific publications, patents, and clinical trials that cite grants directly. These outcomes are captured by the number of citations in papers, patents, and clinical trials to a grant. For a given grant $\mathscr{G}$, its direct impact can be measured by:

- *Number of Funded Papers*: the total number of published papers that acknowledge the funding $\mathscr{G}$.
- *Number of Funded Hit Papers*: the number of hit papers that acknowledge $\mathscr{G}$. A hit paper is a paper that has a citation count, C_p, that is above a particular percentile threshold T_y^f (e.g., top 5%) in a specific field f and year y.
- *Number of Funded Disruptive Papers*: the number of funded papers within the top 5% based on their disruption index, a measure derived from citation networks [72]. The disruption index (D) for a paper p quantifies its divergence from the existing literature, potentially indicating the opening of new research areas. It is calculated as: $D_p = \frac{n_i - n_j}{n_i + n_j + n_k}$, where n_i is the number of future papers that cite the focal paper

p but none of its references, n_j is the number of future papers that cite the paper p and one or more of its references, and n_k is the number of references in p that are not cited by papers citing p.

- *Number of Funded Patents*: the total number of patents that acknowledge $\mathscr{G}$.
- *Number of Funded Clinical Trials*: the number of clinical trials that acknowledge $\mathscr{G}$.

Broader Impact. Recent studies suggest that the true impact of funding goes far beyond publications, affecting industry and society more broadly [10]. To capture these broad outcomes from funding, we examine the citation linkages between funded papers and various downstream applications. By doing so, we not only expand the pools of patents and clinical trials stemming indirectly from funding [34], but also track the broad impact of science on other public sectors such as policymaking and media news. Given a funded paper $\mathscr{P}$, we define the following metrics to estimate its broader impact:

- *Patent Citation*: the total number of patents citing $\mathscr{P}$. This metric represents the marketplace impact of a paper.
- *Clinical Trial Citation*: the total number of clinical trials citing $\mathscr{P}$, measuring the clinical impact.
- *Policy Citation*: the total number of policy documents citing $\mathscr{P}$, measuring the policy impact.
- *Newsfeed Citation*: the total number of news articles citing $\mathscr{P}$, measuring the media impact.

Using the grant-to-paper linkages, we aggregate the impact metrics for all individual papers that cited a grant. We define the broader impact of a grant by summing up all such downstream citation counts received by papers funded by the grant.

Aggregation. Based on grant-level impact, we further aggregate and calculate the funding impact at the PI level and grant field level by summing the direct and broader impact measures of individual grants. Additionally, we calculate the following three metrics for a PI to measure the PI's scientific impact:

- *H-index*: the PI's h-index. The index of a researcher is h if she has h papers with at least h citations and all her remaining papers have fewer than h citations.
- *Productivity*: the PI's total number of publications.
- *Average LogC10*: the log value of the average number of citations within 10 years of publication for all of the PI's papers [73].

We also calculate a relative impact index for a grant field or a funding agency $\mathscr{X}$. We introduce the Relative Impact Index (RII) to benchmark their performance across various impact types, compared to a global baseline. For a given impact type i and a grant field or a funding agency $\mathscr{X}$, RII measures the fraction of grants in $\mathscr{X}$ with an impact in an impact type i, normalized by the same fraction obtained on all grants for that impact type:

$$RII_{i,\mathscr{X}} = \frac{\text{num_grant } in\ \mathscr{X}\ with\ impact\ type\ i / \text{num_grant } in\ \mathscr{X}}{Total\ \text{num_grant } with\ impact\ type\ i / Total\ \text{num_grant}} \tag{1}$$

4.3 Scientific Investment Recommendation

One powerful component of *FtF* is its capability to inform funding agencies about future investment opportunities. This aspect of the system allows funders who seek a particular outcome to identify optimal topics and principal investigators (PIs) for future grants. As a substantial amount of time must pass after a grant is given to observe evidence of the impact of the funding, we employ a machine learning model that forecasts the future impact of recent grants across various impact domains and topics, thereby allowing real-time insights into a recent grant's likely outcomes and allowing for the identification of emerging high-impact grant topics and PIs within a chosen impact domain.

In particular, we introduce a prediction model that utilizes grant abstracts as input to predict the intrinsic potential for various types of impacts, encompassing both direct outcomes and broader impacts, as detailed in Section 4.2. For instance, focusing on patents as direct outcomes, our prediction task is to assess the potential of a grant to directly result in at least one patent in different patent categories, with each patent topic requiring a specific machine learning model. To this end, we leverage SciBERT [74], which is a large language model based on the BERT architecture and has been pre-trained on a large corpus of scientific text data, enabling it to understand text within the scientific domain. SciBERT transforms words in a grant abstract into tokens, each of which is mapped to a 768-dimensional vector. These vectors are then averaged to generate a single vector representing the entire abstract. This resultant vector can be employed as input for machine learning classifiers. One such classifier is XGBoost [75], an efficient and scalable implementation of gradient boosting algorithms. XGBoost is known for its speed and performance in handling large-scale datasets and has been widely used in various machine-learning tasks due to its high accuracy and interpretability. Here, we utilize XGBoost as the classifier, which takes the SciBERT embeddings as the input and determines whether a grant generated the impact regarding the specific prediction task. Additionally, we employ patents as direct outcomes to compare the performance of XGBoost with other common classifiers (e.g., random forest). Our findings indicate that XGBoost demonstrates comparable performance with other classifiers.

Implementation Details. A number of impact prediction models have been trained in our system for different impact factors. To facilitate the understanding, we take the "direct patent outcome" prediction as an example. Other models are trained by following a similar procedure.

In particular, we employ the pre-trained SciBERT model provided by the Allen Institute for AI [76]. This model takes a text input with no more than 512 words. It is an uncased model that does not distinguish between uppercase and lowercase letters during training or inference. We use it to convert the grant text into latent vectors. In the next step, we utilize the official Python implementation of XGBoost [77] in our system as the prediction model. The model has been trained based on a binary classification task with the goal of estimating the probability of a given grant having patent outcomes in topic $\mathscr{A}$ in the future.

To train the model, we collected all grants either from a funding agency or from a funding field. Based on these data, we calculate the averaged time cost (denoted as Y) for filing a patent with the funding support from the grant. After that, Y is rounded as an integer to facilitate calculation. All the grants from 2000 to $2021-Y$ have been used as the training samples for our prediction model. Here, $2021-Y$ gives a reasonable time span for filing the patent, mitigating potential biases from using grants with insufficient time windows. After that, we prepared both the positive and negative samples. In particular, given a topic $\mathscr{A}$, we collect all the related grants based on which one or more patents have been filed from 2000 to $2021-Y$ as the positive samples, whereas an equal number of grants without any patent outcomes are randomly sampled as negative samples. 80% of the positive and negative samples were used as training samples, leaving the rest 20% for testing.

5 Visualization Design

In this section, we introduce the detailed visualization design of our system. We first briefly introduce the user interface as an overview and then describe the specific details of each visualization view.

5.1 User Interface

Our system integrates six coordinated views to support the analysis tasks outlined in Section 3.1 (Fig. 4). Given the large amount of grant data, our system always lets the users start by filtering the data via the *Query View* (**S1**, Fig. 4(A)), in which they can choose a funding agency or select a time range. The selected grants are shown in the *Grant View* (**S2**, Fig. 4(B)), in which the grants are clustered based on different

This is a working paper. Please do not distribute or share publicly.



Figure 4: The *FtF* system UI. The *Query View* (A) is for data filtering. The *Grant View* (B) and *PI View* (C) present grant project summaries. The *Impact Landscape View* (D) allows for a detailed exploration of the multidimensional impact of funding. The *Impact Type View* (E) enables users to select the type of impact they want to explore, and the *Impact Entity View* (F) displays contextual information about different types of impacts.

research fields as bubbles with the bubble size indicating the funding amount. Once a field is selected, the details such as demographics and profiles of the corresponding PIs are visualized in the *PI View* (**S3**, Fig. 4(C)) for users' reference. At the same time, the corresponding grants, as well as their outcomes (i.e., policies, patents, clinical trials, and news reports that cite these grants, indicating the direct or broader impact of these grants), are visualized in the *Impact Landscape View* (**E1**, Fig. 4(D)) in the form of a multi-partite graph. In particular, all the grants are aggregated by research topics and shown in the center as glyphs. Their outcomes are also aggregated by topics and visualized as different clusters of graph nodes surrounding and connecting to the center glyphs. The sizes of the clusters directly estimate the grants' impact in terms of the corresponding outcomes. Users can hover on the grant or impact topic nodes and use a word cloud to learn fine-grained text information about this topic (Fig. 7(A, B)). Users can switch between direct and broader impacts from the *Impact Type View* (**E1**, Fig. 4(E)). The outcome clusters in the *Impact Landscape View* will be changed accordingly. The system also incorporates a context view showing the outcome distribution over other dimensions, such as policy source organizations (**E2**, Fig. 4(F)).

Despite the above design that supports impact exploration, our system also predicts the potential impacts of each recent grant. The topics (i.e., the nodes in the impact graph) of the grants that may have a higher impact in the future are highlighted by purple rings (**R1**, Fig. 7(C)). The corresponding PIs are also listed in Fig. 4(C2) for users (funding agencies) to make a future investment.

5.2 Impact Landscape View

We designed the *Impact Landscape View* to illustrate the multidimensional impact of grants, using citation links between the grants and their outcomes in the form of a multipartite impact graph (**E1**, Fig. 4(D)). A grant may have four different types of outcomes (i.e., direct or broader impact), including policies, patents, clinical trials, and news reports that acknowledge or mention the grant or cite the funded papers of that grant. These outcomes reveal the grant's impact from different aspects. To ensure scalability, both the grants and their outcomes are aggregated by their topic labels provided in our data. As a result, in the impact graph, each node either indicates a collection of grants under the same research topic (denoted as the ***grant topic node***) or a collection of outcomes indicating the same type of impact and sharing the same topic (denoted as the ***impact node***). The node size represents the number of grants or outcomes in the cluster.

Layout. To layout the graph, we introduce a hybrid layout method. In particular, we first use a force-directed layout method to layout the entity-type level graph (e.g., grant, paper, and patent). Second, within each impact type (i.e., paper, patent, clinical trial, and newsfeed), we employ a bubble treemap layout [78] to pack the same type of impact nodes together as a cluster with a hierarchical topic structure. The centers of these bubble treemaps are then positioned at the locations of their corresponding entity-type nodes within the entity-type level graph. Third, the cluster of grant topic nodes is initially placed in the middle of the view, surrounded by clusters of different impact nodes that are placed at the edge of the view. Another force-directed layout is employed to fine-tune the positions of the grant topic nodes to better reveal their relationships with a visually coherent and balanced layout. Specifically, the following forces are introduced to fine-tune the initial layout of the grant topic nodes:

- *Impact Force (F_{impact}):* an attraction force between a grant topic node in the middle and the impact nodes in the surrounding clusters. It is formally defined as:

$$F_{\text{impact}} = \sum_{i=1}^{n} \begin{cases} \left(RII_i - 1\right) \cdot d_i \cdot \hat{u}_i, & RII_i \geq 1, \\ \left(RII_i - 1\right) \cdot (1/d_i) \cdot \hat{u}_i, & RII_i < 1, \end{cases} \tag{2}$$

  where F_{impact} is the impact force, RII_i is the RII value for the i-th impact factor, $d_i = \|X_{\text{impact}_i} - X_{\text{grant}}\|$ is the Euclidean distance between the i-th impact type node and the grant node, $\hat{u}_i = \left(X_{\text{impact}_i} - X_{\text{grant}}\right)/d_i$ is the unit direction vector, X_{impact_i} is the position vector of the i-th impact node, and X_{grant} is the position vector of the grant glyph.

- *Containment Force ($F_{contain}$):* the force to constrain the positions of the grant topic nodes within the boundary of its containing cluster. It pulls a node towards the center if it strays beyond a predefined maximum distance.

$$F_{\text{contain}} = - \cdot \max(0, d - d_{\max}) \cdot \hat{r}, \tag{3}$$

  where d is the distance from the node to the center, $d_{\max}$ is the maximum allowed distance, and $\hat{r}$ is the unit vector pointing from the center to the node.

- *Collision Force ($F_{collide}$):* a force that prevents nodes from overlapping with a padding parameter p. It is proportional to the amount of overlap between nodes:

$$F_{\text{collide}} = \cdot \max(0, r_1 + r_2 + p - d) \cdot \hat{r}, \tag{4}$$

  where F_{collide} is the collision force, r_1 and r_2 are the radii of the two nodes, d is the distance between the centers of the two nodes, p is a constant padding term, and $\hat{r}$ is the unit vector pointing from the center of one node to the center of the other node.

This is a working paper. Please do not distribute or share publicly.

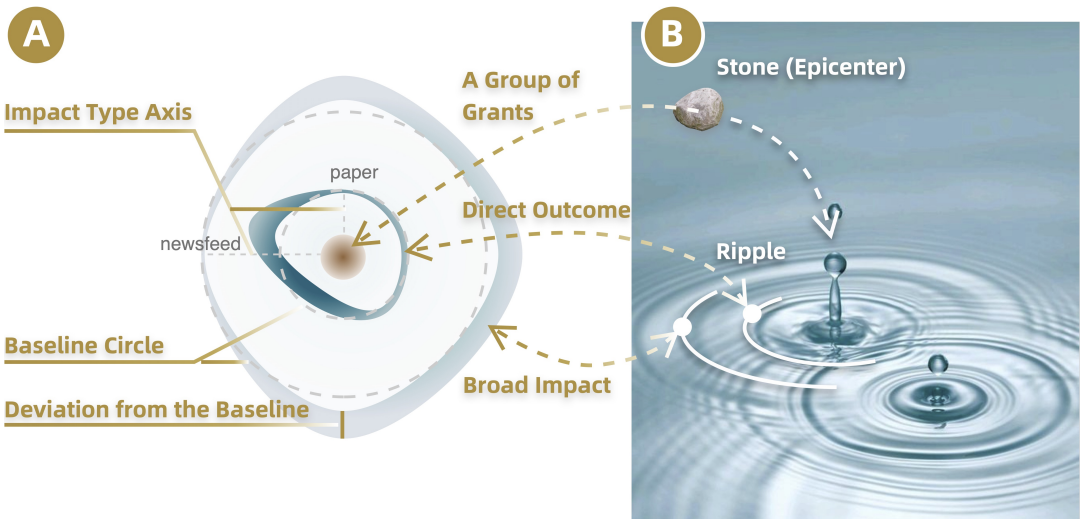


Figure 5: The visual design of the *ImpactGlyph* is inspired by the metaphor of a ripple. (A) In the historical mode, the glyph consists of a central grant circle surrounded by two concentric impact ripples. (B) An illustrative ripple metaphor, with a stone as the "epicenter."

The three forces are balanced by three parameters to adjust the primary effect of the final layout:

$$\vec{F}_{\text{total}} = \alpha\vec{F}_{\text{impact}} + \beta\vec{F}_{\text{contain}} + \gamma\vec{F}_{\text{collide}}, \tag{5}$$

where $\vec{F}_{\text{total}}$ is the total force acting on the node, α, β, and γ are parameters that control the strength of the impact force, containment force, and collision force, respectively, and $\vec{F}_{\text{impact}}$, $\vec{F}_{\text{contain}}$, and $\vec{F}_{\text{collide}}$ are the impact force, containment force, and collision force, respectively.

We simulate the force system with alpha cooling and velocity decay, allowing it to gradually stabilize into a visually coherent and well-balanced layout. Finally, the citation linkages between the grant topic nodes and the impact nodes are bundled by impact types to reduce visual clutter and facilitate reading.

ImpactGlyph. To summarize the multidimensional impact of each collection of grants represented by the grant topic node and facilitate the impact comparisons across different grant topics, the *ImpactGlyph* is introduced. The design is inspired by a visual metaphor of the ripple caused by throwing a stone into the water, as shown in Fig. 5.

In particular, the glyph consists of a center node (i.e., the collection of grants in the same topic) and two levels of outer rings respectively, indicating the corresponding direct and broader impacts of these grants. Two grey circular dashed lines set the global baseline for the RII index for the direct and broad impacts, respectively. Two radar charts symbolize the direct and broader impacts as concentric ripples from multiple dimensions, including patents, clinical trials, policies, and news mentions. In our design, we use a belt with thinness to replace the polyline that indicates multidimensional values in the traditional radar chart. The thickness of the belt encodes the grants' impact deviations across different impact dimensions from the baseline, which magnifies the difference. The color of the belt transitioning from darker to lighter green illustrates the grants' spillover effect from immediate outcomes to wider societal impacts. Altogether, this visualization enables users to swiftly grasp the impact of a group of grants across various dimensions relative to the global baseline using the shape and color of the ripple.

In our system, the above glyph design has two modes: (1) the historical mode and (2) the prediction mode. In the historical mode (Fig. 5), the historical funding impact is compared with global baselines using

the RII index. The size of the center node reflects the number of historical grants within the timeframe selected by the user, serving as the "epicenter" of the ripple. In the prediction mode (Fig. 7(C)), after specifying an impact topic of interest, the number of recent grants with high-prediction scores (above a threshold) will be highlighted as a purple ring, wrapping the previous historical center node, with the size encoding the number of high-prediction score grants. The two impact ripples are hidden for simplicity.

TopicLens. Despite the above views, a *TopicLens* (Fig. 7(A, B)) is designed to illustrate the detailed topic keywords in the form of a wordle when the mouse hovers over a grant topic node in the impact graph. Word size encodes the overall keyword frequency with an overlaid sparkline showing its temporal trend, which provides insights into evolving grant or impact topic foci over time.

5.3 Contextual Views

Despite the above primary views, *FtF* provides a list of side views with contextual information, enabling users to understand the funding impacts and make informed decisions for future investments.

Query View. In this view, users can filter grants based on criteria such as grant year, funding amount, and agency, enabling the flexible search for grants of interest (**S1**, Fig. 4(A)).

Grant View. This view summarizes the funded projects for the selected grants in terms of grant fields or sub-institutions (**S2**, Fig. 4(B)). The size of the bubble encodes the total funding amount within each grant field or sub-institution. The bubble's position indicates the similarity among grant groups (i.e., the semantic similarity of grant fields and funded grant fields of sub-institutions), as determined by dimension reduction techniques (i.e., t-SNE).

PI View. This view shows the statistics of funded PIs, including demographics, universities, fields of study, and impact metrics (**S3**, Fig. 4(C)). In addition, we also provide a query and ranking using universities and impact metrics to enable locating detailed information in the PI list.

Impact Type View. This view provides a list of impact types for users to select, which will be further explored in the *Impact Landscape View* (**E1**, Fig. 4(E)). For each impact type, a bar chart shows the average number of impact documents produced by a grant against a global baseline. Users can switch between direct and broader impacts via a dropdown menu.

Impact Entity View. In addition to topics, this view provides context through bar charts summarizing the distribution of impact documents across various dimensions (**E2**, Fig. 4(F)), such as patent assignees, policy sources, clinical trial phases, and newsfeed origins, aiding in a deeper understanding of funding impacts.

6 Evaluation

We evaluated *FtF* through a quantitative study of the prediction model, two case studies, and a series of interviews with experts.

6.1 Quantitative Evaluation of the Prediction Model

We evaluated our model's performance using the Area Under the Curve (AUC) metric and assessed its prediction accuracy and scalability across the two datasets utilized in the case studies (Section 6.2). Employing the prediction pipeline described earlier, we focused on the top K topics in the predictions for each type of impact, ensuring these topics collectively cover more than 80% grants in each prediction task. We then filtered out topics with fewer than 100 positive samples.

- **AUC**. We present the AUC results on the test sets for the two case studies across different types of impacts in Fig. 6. The figure shows that overall prediction performance is good and remains robust for all types of impacts.

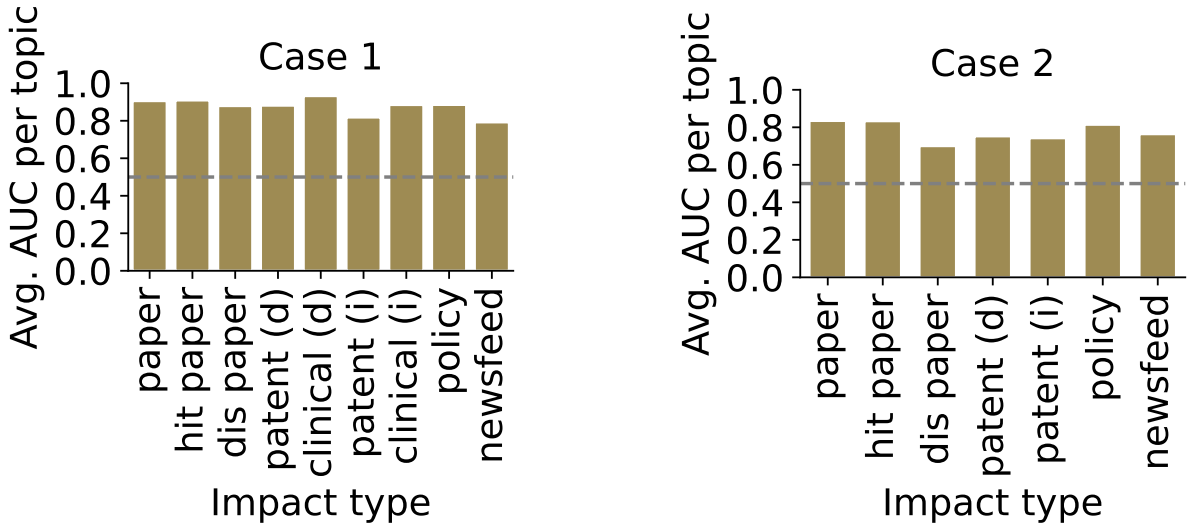


Figure 6: The prediction performance (AUC) on the test set for two case studies. The average AUC is reported for each type of impact, based on the performance across topics for that type of impact. The dashed lines at 0.5 indicate the baseline performance of random guesses. Overall, the results demonstrate good performance for our prediction tasks. (d) represents direct impact and (i) represents indirect impact.

- **Scalability**. XGBoost is known for its scalability and efficiency [75]. We also ran the codes in parallel for different topics to accelerate the prediction process. Moreover, the SciBERT embeddings and prediction model are pre-executed and do not impact the visualization system in real time.

6.2 Case Study

We invited our experts to explore the system and then summarized their observations and formed two case studies to demonstrate our system.

6.2.1 Funding Impact for a Funding Agency

We first demonstrated our system from a funder's perspective, engaging our experts E_A and E_B and showcasing the ability of *FtF* to help them access the funding impact and inform investment decisions.

Overview of the funding portfolio. The experts first query grants from a federal funding agency (**S1**, Fig. 4(A)) and start from the *Grant View* and *PI View* to get an overview of all funded projects. This agency has a focus on four main fields (**S2**, Fig. 4(B)). Zooming in on these grant fields, E_A noticed a gender disparity across some of the topics they fund, where female PIs appear underrepresented (**S3**, Fig. 4(C1)). Looking into the radar chart for details, he found that projects by female PIs were on par and, in some cases surpassed those by male PIs in certain impact measures (e.g., policy impact). This insight about gender disparity clearly caught the attention of the expert, *"We should definitely look more into this."*

Funding impact. The expert E_A was very excited to explore the societal impacts of their grants, as these impact measures represent entirely new information for him. Indeed, given the novelty of the data and the linkages, the various impact dimensions offered by *FtF* present information that has not been available to decision-makers like E_A, who described himself as "blind" to these outcomes. E_A turned to *Impact Type View* (**E1**, Fig. 4(E)), selecting the broad impact. He was thrilled to learn that the marketplace, clinical, and policy impacts of their grants were notably higher than the global baseline. For further details about these types of impacts, he went to the *Impact Landscape View* (**E1**, Fig. 4(D)). The *ImpactGlyph* allowed E_A to focus on the field with the largest proportion of funding expenditures which has high impacts of the three impact types (Fig. 4(D1)) and to zoom in on the detailed topic-level impacts. The system revealed that the

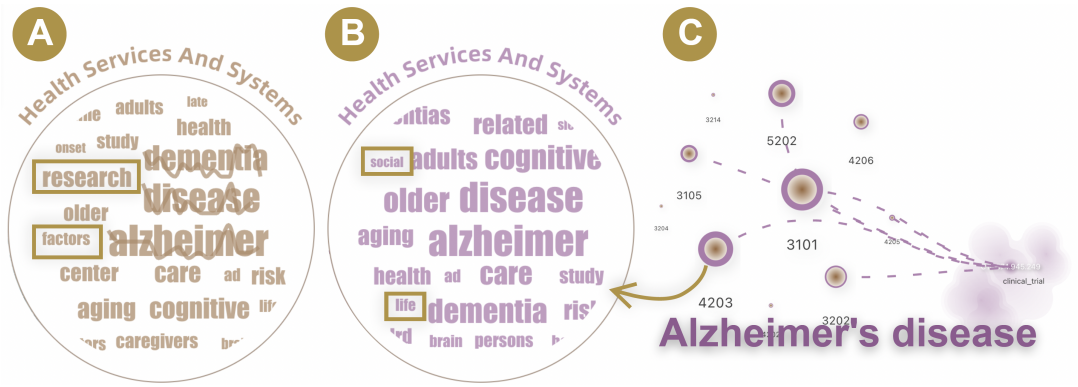


Figure 7: Predicted result in Case 1. (A) Historically, grants with impacts on Alzheimer's clinical trials have focused on disease understanding. (B) The recent grants with high prediction scores are oriented toward social and life dimensions of the disease. (C) The topic distribution of recent high-prediction score grants (purple rings) to have an impact on Alzheimer's disease clinical trials.

funded research had been cited in many policy topics (e.g., "Virus Disease", Fig. 4(D2)). Although the agency mainly funds research in the US, he was delighted to learn that the policy impact of their funded projects extends across the globe, with numerous policy citations from the Germany, United Kingdom, and other intergovernmental organizations (IGOs) (**E2**, Fig. 4(F1)), so as the patent impact (Fig. 4(F2)). He exclaimed, *"My boss should see this!"* Our interactions with E_A suggest that *FtF* provides him with hard data that can substantially inform their decision-making. We were especially delighted to see that some of the insights were derived from novel data linkages that offered relevant yet previously unknown information to key decision-makers in science, like E_A.

Future investment guidance. Whereas E_A is from a public funder, our expert E_B comes from a private investment firm that specializes in investing in biomedical research in academia to advance healthcare. E_B was particularly interested in Alzheimer's research, probably due to the increased investment in this area by his organization. Here he was drawn to our predictive analytics, which showcases the predicted clinical impacts of current projects related to Alzheimer's disease (**R1**, Fig. 7(C)). Insights from predictive analytics are, in general, new to decision makers and must be interpreted with care. Nevertheless, by looking at the detailed grant keywords, E_B noticed an important trend. Whereas the research on Alzheimer's disease that has high clinical impacts over the past twenty years focused primarily on disease understanding (Fig. 7(A)), the predictive analytics suggest that current grants with high predicted clinical impacts (i.e., high prediction scores) seem to be more concentrated on the social and life dimensions of the disease (Fig. 7(B)). Musing over this trend, E_B noted, *"This is very interesting. I wonder if we should look more into combining medical treatment with social support systems."* E_B also consulted the *PI View* (Fig. 4(C2)) and noticed several top PIs with relevant research who had not yet been funded by his organization. The system allowed him to easily navigate to these PIs' research profiles to assess potential fit and collaboration opportunities. Here, Alzheimer's disease is presented merely as an illustrative example to demonstrate the system's predictive capabilities, and similar analyses could be conducted for other diseases or research areas.

Overall, our system offers domain experts key information that they were previously unaware of. This knowledge broadens and sharpens the search space over which they evaluate investments (in ideas or people). While funding or investment decisions depend on a large variety of factors and considerations, *FtF* helps decision makers reduce key blind spots and ultimately offers a better and more comprehensive set of data to inform evidence-based decision making.

6.2.2 Funding Landscape for the Visualization Community

Given the growing interest in interdisciplinary research [79], our SciSci experts (E_C-E_E) examined visualization field grants to illustrate the diverse impact of global funding on interdisciplinary research.

Overview of the funding impact. The experts began their exploration of the system by looking at an overview of the funded projects (**S2**, Fig. 8(B1)). Reflecting the highly interdisciplinary nature of the vis community, the funded projects cover a wide range of topics beyond computer science, such as education, biomedicine, and design. Exploring the broad impact of these projects, they found that the projects have had a significant influence on transportation, environmental, and health policies (**E1**, Fig. 8(A1)). Digging deeper into the policies about oceans, they discovered that these policies had a strong emphasis on climate change and that this impact was seen in the policy documents of multiple countries, as well as IGOs (**E2**, Fig. 8(A1)). An expert noted that the system documents a contribution of visualization research that had not necessarily been recognized in the past: *"Interesting... This suggests that vis research plays a role in addressing key societal challenges such as climate change, which is currently underappreciated."* The experts then turned to the media impact (**E1**, Fig. 8(A2)). Both specialized and general-audience media organizations have covered vis research, touching on a wide range of topics, including AI and pandemics. The breadth of coverage reflects the far-reaching media interest in and the applicability of vis research. While experts may have been aware of this broad appeal at a high level, *FtF* provided them with concrete evidence of this range and specificity that they had not seen before. The system also allowed the experts to compare the direct and indirect outcomes of grants in terms of patents (**E1**, Fig. 8(A3)). They discovered a new patent category in the top frequent patent fields in indirect impact: transmission of digital information (H04L). *"It seems that researchers in vis don't typically file patents in this technology topic, but their scientific expertise is widely leveraged in related areas,"* noted an expert, seeing an *"opportunity in leveraging vis for digital communication technologies."*

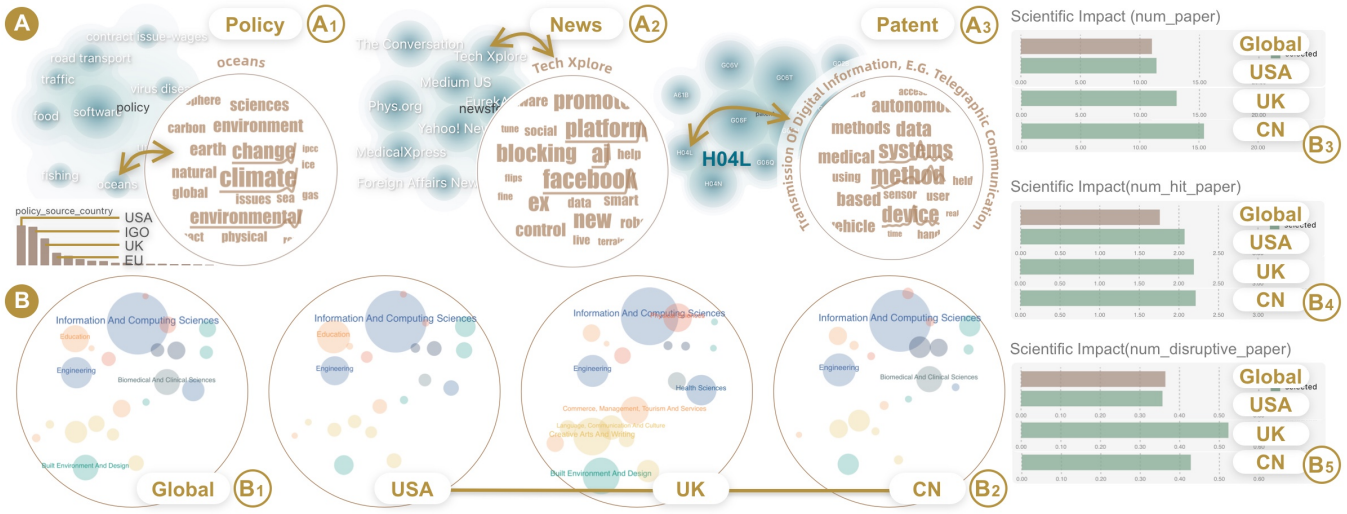


Figure 8: The impacts of funding for Case 2. (A) The broader impacts of funded visualization research include the policy impact for climate change (A1), media impact for AI (A2), and marketplace impact for digital communication (A3). (B) The comparison of funded visualization research in different countries shows interesting variation between the fields that are funded ((B1)-(B2)) and the science impact ((B3)-(B5)).

National disparities in visualization research. The experts were also interested in exploring national differences in funded vis projects and their impacts. For this, the experts compared three countries: the

United States, the United Kingdom, and China. Most funding for vis projects goes to the computer science topic, but there were notable differences between the countries in the distribution of other fields (**S2**, Fig. 8(B2)): the U.S. prioritizes education, the U.K. favored art and design, and China emphasized biomedical fields. Turning to the scientific impacts, though vis projects in the three countries all tend to be greater than the global average, the experts found that there were notable variations across countries that they had not seen before (**E1**, Fig. 8(B3)-(B5)). The U.K. showed the highest propensity for disruptive science among the three countries (Fig. 8(B5)). Interestingly, vis papers from China feature a higher fraction of disruptive work than the U.S.. When the experts combined this insight about disruptive science with other related information the system offered, including the hit papers statistics, they learned that vis funding from China tends to yield more papers per grant (Fig. 8(B3)), but not substantially more hit papers than most other countries (Fig. 8(B4)).

In this case, *FtF* offers experts a wide range of insights—from the interdisciplinarity of vis work to national differences in the impact.

6.3 Expert Interview

In addition to feedback from the experts described in Section 3.1, we presented the system to new experts, including a founding director of a leading university's Technology Transfer Office (P_A) and senior researchers who focus on the impact of science on public policy and health (P_B and P_C). We also interviewed three visualization experts to potentially identify other uses of the system beyond its applications for funding agency decision-making, including a senior Ph.D. experienced in grant writing (P_D) and two graduates with over three years of research experience (P_E and P_F). Each interview with these new experts lasted 90 minutes, covering project introductions, system and visual design illustrations (using case 1), free time to explore the system, and a semi-structured interview. Experts were encouraged to share their thoughts and impressions during the process. We summarize the feedback from the two groups of experts below.

Analysis Workflow. All experts reported that the system's workflow was straightforward and that they could complete the analysis tasks, from data query and funded project analysis to the evaluation of impact. E_B praised the system's sophistication and its advances beyond the capabilities offered by their daily used tools like WOS and Scopus. E_A was also excited about the system's ability to identify research topics and PIs using prediction results: *"This changes the game. Instead of passively waiting for proposals, I can use the tool as a radar to scan all suitable topics and PIs to invest in."*

Investment prediction and impact metrics. E_A appreciated the prediction feature that allowed him to effectively identify potential investment opportunities, and he commended the impact metrics for their clarity and usefulness. E_A also praised the relative impact index and the *ImpactGlyph*, which he saw as valuable tools for quickly identifying the strengths and weaknesses of a group of grants. However, he called for greater transparency in the prediction model. Voicing similar enthusiasm, P_C appreciated the PI and grant field level impact measure to compare high-level funding impacts such as funders and countries.

Visualization. The experts agreed that the visualization components satisfy all of the analysis tasks. E_B, P_A, and P_E were especially impressed with the citation flows in the *Impact Landscape View* that features an intuitive presentation of the impact of funding at the topic level. P_D commented that the word cloud design was both readable and easy to understand. Many of the experts appreciated the *ImpactGlyph*. One observed, *"the ripple metaphor shows the multi-dimensional impact in a vivid way."* Nevertheless, E_A and P_A noted that there would be a learning curve in their use of the glyph, acknowledging, *"it's a bit new to me."* All experts were able to understand and use it effectively, however, after being given some time to explore the system. One expert affirmed, *"I think it will be very useful to compare group level grants, as the difference among glyph shapes is quite apparent."*

Suggestions. The experts saw the current metrics as straightforward and useful for analyzing the impact

of funding, but they had some thoughtful suggestions adding more innovative measures such as "risk." E_A also recommended adding a raw data table that presents both grants and impact documents for deeper analysis. E_B pointed out potential limitations in the data, *"many patents do not directly indicate the support of grants. We need to do deep data cleaning to search the whole patent doc to see whether they have claimed grant support."*

7 Discussion

Our study illustrates how connecting funding to its downstream impacts opens new possibilities for both research and practice. By linking grants to publications, patents, clinical trials, policy, and media at scale, *Funding the Frontier* (*FtF*) provides not only a descriptive map but also an integrative infrastructure for future work in the science of science and science policy.

The availability of large-scale data linking upstream funding to diverse downstream impacts presents new research opportunities. With *FtF*, researchers can examine how different funding models shape the trajectory of science and innovation, compare funding portfolios across nations, and study the efficiency and time lags of investment returns across domains. The system can also support investigations into long-standing questions. For example, how public versus private funders differ in their societal footprint, or whether funding strategies that emphasize interdisciplinarity systematically yield broader impacts. In this sense, *FtF* represents not an end point but a foundation for new lines of work across the science of science, innovation studies, and policy research.

Equally important, *FtF* was built to be actionable by non-technical stakeholders. Policymakers, program officers, university leaders, and even philanthropic funders face growing pressure to justify investments, prioritize areas of opportunity, and demonstrate societal relevance. By making complex data accessible and interpretable, *FtF* equips these decision-makers to explore their portfolios, identify blind spots, and consider new directions. For example, program officers might use the system to surface emerging investigators with outsized potential; university leaders might assess how their institution's grants contribute to policy or market outcomes; and policymakers might gain perspective on whether federal portfolios align with national and regional priorities. The diversity of these use cases underscores the broad value of making SciSci insights transparent and navigable.

While our case studies focus on selected funders and domains, the underlying datasets and framework are global in scope. The system is not limited to U.S. agencies but can, in principle, extend to philanthropic organizations, multinational funding programs, and national science foundations worldwide. This breadth makes *FtF* a potential global public good: a resource for comparing funding impacts across borders and for enabling international comparisons at a moment of intensifying global competition in science and technology.

As a prototype, *FtF* has several limitations. Our predictive models, while effective, can be further improved with more data and more advanced machine learning methods, and can benefit from further development to maximize transparency and interpretability, particularly for stakeholders unfamiliar with the technical aspects. In addition, while our datasets are among the most comprehensive, they are inevitably shaped by coverage biases, such as the predominance of English-language sources in policy documents and newsfeeds, which may underrepresent impacts in other linguistic and regional contexts. Further, while our evaluation demonstrates the system's usability, the evaluation should be considered as exploratory, and future work may further expand the datasets, integrate additional impact measures, and tailor customized visual workflows to distinct user groups, from administrators and leaders to data scientists and researchers. Addressing these challenges will improve the robustness of the system, further enhancing its value as an infrastructure for both decision-making and research.

Taken together, *FtF* demonstrates how the combination of SciSci insights and visualization design can

help address a long-standing gap between funding inputs and societal outcomes. At a time when science budgets face scrutiny, geopolitical competition intensifies, and public trust in science is contested, the ability to assess and communicate the multidimensional impacts of funding is critical. By providing a transparent, scalable, and extensible framework, *FtF* offers both a practical resource for decision-makers and a foundation for future research into the broad impacts of science and science investment.

Data Availability

Data necessary to reproduce all plots will be made freely available.

Code Availability

Code necessary to reproduce all plots will be made freely available.

References

[1] Williams, R. S., Lotia, S., Holloway, A. K. & Pico, A. R. From scientific discovery to cures: Bright stars within a galaxy. *Cell* **163**, 21–23 (2015). URL https://doi.org/10.1016/j.cell.2015.09.007. 2, 3

[2] Wang, X., Zhang, S., Liu, Y., Du, J. & Huang, H. How pharmaceutical innovation evolves: The path from science to technological development to marketable drugs. *Technological Forecasting and Social Change* **167**, 120698 (2021). URL https://doi.org/10.1016/j.techfore.2021.120698. 2, 3

[3] Thelwall, M. & Kousha, K. Are citations from clinical trials evidence of higher impact research? An analysis of ClinicalTrials.gov. *Scientometrics* **109**, 1341–1351 (2016). URL https://doi.org/10.1007/s11192-016-2112-1. 2, 3

[4] Ahmadpoor, M. & Jones, B. F. The dual frontier: Patented inventions and prior scientific advance. *Science* **357**, 583–587 (2017). URL https://doi.org/10.1126/science.aam9527. 2, 3

[5] Fleming, L., Greene, H., Li, G., Marx, M. & Yao, D. Government-funded research increasingly fuels innovation. *Science* **364**, 1139–1141 (2019). URL https://doi.org/10.1126/science.aaw2373. 2, 4

[6] Council, N. R. *et al. Using science as evidence in public policy* (National Academies Press, 2012). URL https://doi.org/10.17226/13460. 2, 3

[7] Furnas, A. C., LaPira, T. M. & Wang, D. Partisan disparities in the use of science in policy. *Science* **388**, 362–367 (2025). 2, 3

[8] Bush, V. *Science, the endless frontier* (Princeton University Press, 2020). URL https://doi.org/10.1515/9780691201658. 2

[9] Azoulay, P., Graff Zivin, J. S. & Manso, G. Incentives and creativity: evidence from the academic life sciences. *The RAND Journal of Economics* **42**, 527–554 (2011). URL https://doi.org/10.1111/j.1756-2171.2011.00140.x. 2, 3

[10] Yin, Y., Dong, Y., Wang, K., Wang, D. & Jones, B. F. Public use and public funding of science. *Nature Human Behaviour* **6**, 1344–1350 (2022). URL https://doi.org/10.1038/s41562-022-01397-5. 2, 3, 4, 9

[11] Jones, B. F. Science and innovation: The under-fueled engine of prosperity. *Rebuilding the Post-Pandemic Economy, ed. Melissa S. Kearney and Amy Ganz (Washington DC: Aspen Institute Press, 2021)* (2021). 2

[12] NSF's perspectives on broader impacts. https://new.nsf.gov/funding/learn/broader-impacts. 2, 3

[13] Price, D. J. D. S. *Little science, big science* (Columbia University Press, 1963). URL https://doi.org/10.7312/pric91844. 2

[14] Fortunato, S. *et al.* Science of science. *Science* **359**, eaao0185 (2018). URL https://doi.org/10.1126/science.aao0185. 2

[15] Wang, D. & Barabási, A.-L. *The science of science* (Cambridge University Press, 2021). URL https://doi.org/10.1017/9781108610834. 2

[16] Weller, K. Social media and altmetrics: An overview of current alternative approaches to measuring scholarly impact. *Incentives and performance: Governance of research organizations* 261–276 (2015). 2, 3

[17] Molnar, A., McKenna, A. F., Liu, Q., Vorvoreanu, M. & Madhavan, K. Using visualization to derive insights from research funding portfolios. *IEEE Computer Graphics and Applications* **35**, 91–c3 (2015). URL https://doi.org/10.1109/MCG.2015.68. 2, 5

[18] Dou, W., Yu, L., Wang, X., Ma, Z. & Ribarsky, W. HierarchicalTopics: Visually exploring large text collections using topic hierarchies. *IEEE Transactions on Visualization and Computer Graphics* **19**, 2002–2011 (2013). URL https://doi.org/10.1109/TVCG.2013.162. 2, 5

[19] Dimensions. https://www.dimensions.ai/. 2, 5, 7

[20] Aggregated analysis of ClinicalTrials.gov (AACT). https://aact.ctti-clinicaltrials.org/. 2

[21] Overton. https://www.overton.io/. 2, 8

[22] Patentsview. https://patentsview.org/. 2

[23] Altmetric. https://www.altmetric.com/. 2, 3, 5, 8

[24] Gates, A. J. & Barabási, A.-L. Reproducible science of science at scale: pySciSci. *Quantitative Science Studies* **4**, 700–710 (2023). 2

[25] Keim, D. *et al.* Visual analytics: Definition, process, and challenges. In *Information visualization: Human-centered issues and perspectives* (2008). 2

[26] Forscher, P. S., Cox, W. T., Brauer, M. & Devine, P. G. Little race or gender bias in an experiment of initial review of NIH R01 grant proposals. *Nature Human Behaviour* **3**, 257–264 (2019). URL https://doi.org/10.1038/s41562-018-0517-y. 3

[27] Ginther, D. K. *et al.* Race, ethnicity, and NIH research awards. *Science* **333**, 1015–1019 (2011). URL https://doi.org/10.1126/science.1196783. 3

[28] Li, D. & Agha, L. Big names or big ideas: Do peer-review panels select the best science proposals? *Science* **348**, 434–438 (2015). URL https://doi.org/10.1126/science.aaa0185. 3

[29] Wang, J. & Shapira, P. Funding acknowledgement analysis: an enhanced tool to investigate research sponsorship impacts: the case of nanotechnology. *Scientometrics* **87**, 563–586 (2011). URL https://doi.org/10.1007/s11192-011-0362-5. 3

[30] Miao, L., Larivière, V., Wang, F., Ahn, Y.-Y. & Sugimoto, C. R. Cooperation and interdependence in global science funding. *arXiv preprint arXiv:2308.08630* (2023). URL https://doi.org/10.48550/arXiv.2308.08630. 3

[31] Vasan, K. & West, J. D. The hidden influence of communities in collaborative funding of clinical science. *Royal Society Open Science* **8**, 210072 (2021). 3

[32] Wang, Y., Jones, B. F. & Wang, D. Early-career setback and future career impact. *Nature communications* **10**, 4331 (2019). URL https://doi.org/10.1038/s41467-019-12189-3. 3

[33] Bol, T., De Vaan, M. & van de Rijt, A. The Matthew effect in science funding. *Proceedings of the National Academy of Sciences* **115**, 4887–4890 (2018). URL https://doi.org/10.1073/pnas.1719557115. 3

[34] Li, D., Azoulay, P. & Sampat, B. N. The applied value of public investments in biomedical research. *Science* **356**, 78–81 (2017). URL https://doi.org/10.1126/science.aal0010. 4, 9

[35] Dörk, M., Riche, N. H., Ramos, G. & Dumais, S. PivotPaths: Strolling through faceted information spaces. *IEEE Transactions on Visualization and Computer Graphics* **18**, 2709–2718 (2012). URL https://doi.org/10.1109/TVCG.2012.252. 4

[36] Ye, Y., Huang, R. & Zeng, W. VISAtlas: An image-based exploration and query system for large visualization collections via neural image embedding. *IEEE Transactions on Visualization and Computer Graphics* 1–15 (2022). URL https://doi.org/10.1109/TVCG.2022.3229023. 4

[37] Wang, Y. *et al.* Seek for success: A visualization approach for understanding the dynamics of academic careers. *IEEE Transactions on Visualization and Computer Graphics* **28**, 475–485 (2022). URL https://doi.org/10.1109/TVCG.2021.3114790. 4

[38] Yan, X. & Ma, Y. Turing Galaxy: Visualizing turing award laureates. In *IEEE VIS 2021 Posters* (IEEE, 2021). 4

[39] Nobre, C., Streit, M. & Lex, A. Juniper: A Tree+ Table approach to multivariate graph visualization. *IEEE Transactions on Visualization and Computer Graphics* **25**, 544–554 (2019). URL https://doi.org/10.1109/TVCG.2018.2865149. 4

[40] Li, Z., Zhang, C., Jia, S. & Zhang, J. Galex: Exploring the evolution and intersection of disciplines. *IEEE Transactions on Visualization and Computer Graphics* **26**, 1182–1192 (2020). URL https://doi.org/10.1109/TVCG.2019.2934667. 4

[41] Dong, A., Zeng, W., Chen, X. & Cheng, Z. VIStory: Interactive storyboard for exploring visual information in scientific publications. In *Proceedings of the 12th International Symposium on Visual Information Communication and Interaction* (2019). URL https://doi.org/10.1145/3356422.3356430. 4

[42] Guo, Z., Tao, J., Chen, S., Chawla, N. & Wang, C. SD^2: Slicing and dicing scholarly data for interactive evaluation of academic performance. *IEEE Transactions on Visualization and Computer Graphics* **29**, 3569–3585 (2023). URL https://doi.org/10.1109/TVCG.2022.3163727. 4

[43] Narechania, A., Karduni, A., Wesslen, R. & Wall, E. VITALITY: Promoting serendipitous discovery of academic literature with transformers & visual analytics. *IEEE Transactions on Visualization and Computer Graphics* **28**, 486–496 (2022). URL https://doi.org/10.1109/TVCG.2021.3114820. 4

[44] Chen, J. *et al.* VIS30K: A collection of figures and tables from IEEE Visualization conference publications. *IEEE Transactions on Visualization and Computer Graphics* **27**, 3826–3833 (2021). URL https://doi.org/10.1109/TVCG.2021.3054916. 4

[45] Deng, D. *et al.* VisImages: A Fine-Grained Expert-Annotated visualization dataset. *IEEE Transactions on Visualization & Computer Graphics* **1**, 1–1 (2023). URL https://doi.org/10.1109/TVCG.2022.3155440. 4

[46] Isenberg, P. *et al.* Vispubdata.org: A metadata collection about IEEE Visualization (VIS) publications. *IEEE Transactions on Visualization and Computer Graphics* **23**, 2199–2206 (2017). URL https://doi.org/10.1109/TVCG.2016.2615308. 4

[47] Cao, H., Lu, Y., Deng, Y., McFarland, D. A. & Bernstein, M. S. Breaking out of the ivory tower: A large-scale analysis of patent citations to HCI research. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (2023). URL https://doi.org/10.1145/3544548.3581108. 4

[48] Madhavan, K. *et al.* DIA2: Web-based cyberinfrastructure for visual analysis of funding portfolios. *IEEE Transactions on Visualization and Computer Graphics* **20**, 1823–1832 (2014). URL https://doi.org/10.1109/TVCG.2014.2346747. 4

[49] Liu, S., Cao, N., Lv, H. & Su, H. The visual funding navigator: Analysis of the NSF funding information. In *Proceedings of the 15th ACM International Conference on Information and Knowledge Management*, 882–883 (2006). URL https://doi.org/10.1145/1183614.1183778. 4

[50] Liu, S., Cao, N. & Lv, H. Interactive visual analysis of the NSF funding information. In *IEEE Pacific Visualization Symposium*, 183–190 (IEEE, 2008). URL https://doi.org/10.1109/PACIFICVIS.2008.4475475. 4

[51] Kwon, K., Majeti, D., Brian, U. & Pavlidis, I. Scholar Plot: Visualizing scientific careers at a glance. In *Proceedings of the 2nd Annual International Conference on Computational Social Science* (2016). URL https://doi.org/. 4

[52] Schimmel, S. The Public Innovations Explorer: A geo-spatial & linked-data visualization platform for publicly funded innovation research in the united states. CUNY Academic Works (2021). URL https://academicworks.cuny.edu/gc_etds/4409/. 5

[53] Oppermann, L. *et al.* Finding and analysing energy research funding data: The EnArgus system. *Energy and AI* **5**, 100070 (2021). URL https://doi.org/10.1016/j.egyai.2021.100070. 5

[54] NIH RePORTER. https://reporter.nih.gov/. 5

[55] Hesselberth, J. *et al.* nihexporter: an R package for NIH funding data. *bioRxiv* 033456 (2015). URL https://doi.org/10.1101/033456. 5

[56] PlumX Metrics. https://plumanalytics.com/learn/about-metrics/. 5

[57] Ghani, S., Elmqvist, N. & Ebert, D. S. MultiNode-Explorer: A visual analytics framework for generating web-based multimodal graph visualizations. In *EuroVA: International Workshop on Visual Analytics* (2012). URL https://doi.org/10.2312/PE/EuroVAST/EuroVA12/067-071. 5

[58] Ghani, S., Kwon, B. C., Lee, S., Yi, J. S. & Elmqvist, N. Visual analytics for multimodal social network analysis: A design study with social scientists. *IEEE Transactions on Visualization and Computer Graphics* **19**, 2032–2041 (2013). URL https://doi.org/10.1109/TVCG.2013.223. 5

[59] Dou, W., Wang, X., Chang, R. & Ribarsky, W. ParallelTopics: A probabilistic approach to exploring document collections. In *2011 IEEE Conference on Visual Analytics Science and Technology (VAST)*, 231–240 (IEEE, 2011). URL https://doi.org/10.1109/VAST.2011.6102461. 5

[60] Gilbert, J. E., Burnett, M., Ladner, R. E., Rosson, M. B. & Davis, J. Applying the NSF broader impacts criteria to HCI research. In *CHI'11 Extended Abstracts on Human Factors in Computing Systems*, 459–462 (Association for Computing Machinery, 2011). URL https://doi.org/10.1145/1979742.1979534. 5

[61] Lazar, J. *et al.* Workshop on engaging the human-computer interaction community with public policymaking internationally. In *CHI'13 Extended Abstracts on Human Factors in Computing Systems*, 3279–3282 (Association for Computing Machinery, 2013). URL https://doi.org/10.1145/2468356.2479666. 5

[62] Shi, C. *et al.* MedChemLens: An interactive visual tool to support direction selection in interdisciplinary experimental research of medicinal chemistry. *IEEE Transactions on Visualization and Computer Graphics* **29**, 63–73 (2023). URL https://doi.org/. 5

[63] Wang, Y., Qian, Y., Qi, X., Cao, N. & Wang, D. InnovationInsights: A visual analytics approach for understanding the dual frontiers of science and technology. *IEEE Transactions on Visualization and Computer Graphics* **30**, 518–528 (2024). URL https://doi.org/10.1109/TVCG.2023.3327387. 5

[64] Nobre, C., Meyer, M., Streit, M. & Lex, A. The state of the art in visualizing multivariate networks. In *Computer Graphics Forum*, vol. 38, 807–832 (2019). URL https://doi.org/10.1111/cgf.13728. 5

[65] Vehlow, C., Beck, F. & Weiskopf, D. The state of the art in visualizing group structures in graphs. *EuroVis (STARs)* 21–40 (2015). URL https://doi.org/10.2312/eurovisstar.20151110. 5

[66] Beck, F., Burch, M., Diehl, S. & Weiskopf, D. A taxonomy and survey of dynamic graph visualization. In *Computer Graphics Forum*, vol. 36, 133–159 (2017). URL https://doi.org/10.1111/cgf.12791. 5

[67] McGee, F., Ghoniem, M., Melançon, G., Otjacques, B. & Pinaud, B. The state of the art in multilayer network visualization. In *Computer Graphics Forum*, vol. 38, 125–149 (2019). URL https://doi.org/10.1111/cgf.13610. 5

[68] Shi, L. *et al.* OnionGraph: Hierarchical topology+ attribute multivariate network visualization. *Visual Informatics* **4**, 43–57 (2020). URL https://doi.org/10.1016/j.visinf.2020.01.002. 5

[69] Holten, D. Hierarchical edge bundles: Visualization of adjacency relations in hierarchical data. *IEEE Transactions on Visualization and Computer Graphics* **12**, 741–748 (2006). URL https://doi.org/10.1109/TVCG.2006.147. 5

[70] Sedlmair, M., Meyer, M. & Munzner, T. Design study methodology: Reflections from the trenches and the stacks. *IEEE Transactions on Visualization and Computer Graphics* **18**, 2431–2440 (2012). URL https://doi.org/10.1109/TVCG.2012.213. 6

[71] Lin, Z., Yin, Y., Liu, L. & Wang, D. Sciscinet: A large-scale open data lake for the science of science research. *Scientific Data* **10**, 315 (2023). 8

[72] Wu, L., Wang, D. & Evans, J. A. Large teams develop and small teams disrupt science and technology. *Nature* **566**, 378–382 (2019). URL https://doi.org/10.1038/s41586-019-0941-9. 8

[73] Sinatra, R., Wang, D., Deville, P., Song, C. & Barabási, A.-L. Quantifying the evolution of individual scientific impact. *Science* **354**, aaf5239 (2016). URL https://doi.org/10.1126/science.aaf5239. 9

[74] Beltagy, I., Lo, K. & Cohan, A. SciBERT: Pretrained language model for scientific text. In *EMNLP* (2019). URL https://doi.org/10.48550/arXiv.1903.10676. arXiv:1903.10676. 10

[75] Chen, T. & Guestrin, C. XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–794 (2016). URL https://doi.org/10.1145/2939672.2939785. 10, 15

[76] SciBERT implementation. https://github.com/allenai/scibert. 10

[77] XGBoost implementation. https://xgboost.readthedocs.io/en/stable/python/. 10

[78] Görtler, J., Schulz, C., Weiskopf, D. & Deussen, O. Bubble treemaps for uncertainty visualization. *IEEE transactions on visualization and computer graphics* **24**, 719–728 (2017). 12

[79] Park, M., Maity, S. K., Wuchty, S. & Wang, D. Interdisciplinary papers supported by disciplinary grants garner deep and broad scientific impact. *arXiv preprint arXiv:2303.14732* (2023). URL https://doi.org/10.48550/arXiv.2303.14732. 17

Acknowledgments

We thank all members of the Center for Science of Science and Innovation (CSSI) at Northwestern University for helpful discussions, and Alyse Freilich for editing. This work is supported by the National Science Foundation (award number 2404035) and the Future Wanxiang Foundation. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation and the Future Wanxiang Foundation.