

A Survey of Tasks and Visualizations in Multiverse Analysis Reports

Brian D Hall, Yang Liu, Yvonne Jansen, Pierre Dragicevic, Fanny Chevalier, Matthew Kay

▶ **To cite this version:**

Brian D Hall, Yang Liu, Yvonne Jansen, Pierre Dragicevic, Fanny Chevalier, et al.. A Survey of Tasks and Visualizations in Multiverse Analysis Reports. Computer Graphics Forum, 2022, 41 (1), pp.402-426. 10.1111/cgf.14443 . hal-03558950

HAL Id: hal-03558950
https://inria.hal.science/hal-03558950

Submitted on 5 Feb 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

**This is the authors' submitted version which has been accepted for publication on Dec. 14, 2021.
The final version of record will be accessible under DOI 10.1111/cgf.14443**

Volume xx (202x), Number x pp. x–x COMPUTER GRAPHICS *forum*

A Survey of Tasks and Visualizations in Multiverse Analysis Reports

Brian D. Hall[1], Yang Liu[2], Yvonne Jansen[3], Pierre Dragicevic[4], Fanny Chevalier[5], Matthew Kay[6]

[1]University of Michigan, MI, USA; [2]University of Washington, WA, USA; [3]Sorbonne Université, France;
[4]Inria, Université Paris-Saclay, France; [5]University of Toronto, ON, Canada; [6]Northwestern University, IL, USA

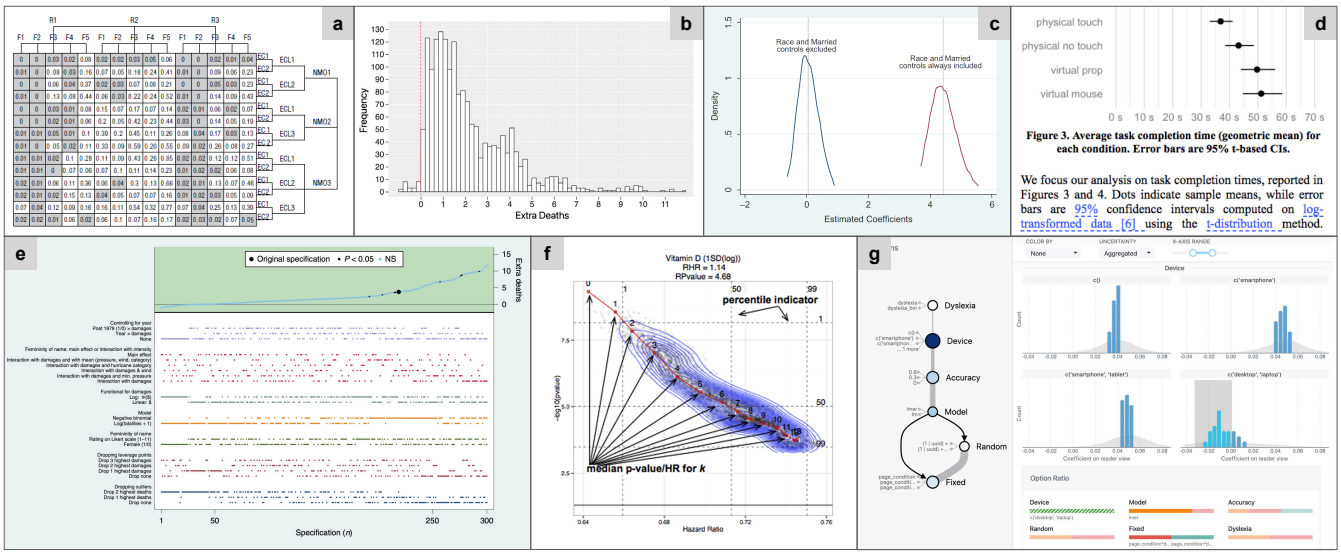


Figure 1: *Examples of multiverse analysis visualizations discussed in this survey: (a) outcome matrix [STGV16], (b) outcome histogram [STGV16], (c) outcome density plot [YH17], (d) explorable multiverse analysis reports [DJS*19], (e) specification curve [SSN20], (f) vibration of effects plot [PBI15], (g) Boba [LKAH20].*

Abstract

Analyzing data from experiments is a complex, multi-step process, often with multiple defensible choices available at each step. While analysts often report a single analysis without documenting how it was chosen, this can cause serious transparency and methodological issues. To make the sensitivity of analysis results to analytical choices transparent, some statisticians and methodologists advocate the use of "multiverse analysis": reporting the full range of outcomes that result from all combinations of defensible analytic choices. Summarizing this combinatorial explosion of statistical results presents unique challenges; several approaches to visualizing the output of multiverse analyses have been proposed across a variety of fields (e.g., psychology, statistics, economics, neuroscience). In this article, we (1) introduce a consistent conceptual framework and terminology for multiverse analyses that can be applied across fields; (2) identify the tasks researchers try to accomplish when visualizing multiverse analyses; and (3) classify multiverse visualizations into "archetypes", assessing how well each archetype supports each task. Our work sets a foundation for subsequent research on developing visualization tools and techniques to support multiverse analysis and its reporting.

Keywords: multiverse analysis, sensibility analysis, transparent reporting, statistical graphics.

1 Introduction: Multiverse Analyses and Visualizations

Analyzing data from experiments is a complex, multi-step process, with multiple choices available at each step, e.g., whether and how to exclude outliers, what approach to use to operationalize a variable, or what model and parameters to apply [LAH20]. While it is often possible to exclude some choices as invalid, often many alternatives remain that are equally valid. Faced with this complexity, analysts often try multiple analyses and report a single one without

© 2022 The Author(s)
Computer Graphics Forum © 2022 The Eurographics Association and John Wiley & Sons Ltd. Published by John Wiley & Sons Ltd.

documenting how it was chosen. This practice, sometimes termed *undisclosed flexibility*, can cause serious transparency and methodological issues [SNS11, WVA*16a, GL13] and has been identified as a major cause of the replicability crisis in psychology and other disciplines [NAB*15, MNB*17, Cum14]. Combined with a desire to report positive findings, undisclosed flexibility can be damaging because it substantially increases the chances of reporting erroneous findings, while being invisible to the reader.

One increasingly advocated solution to the issue of undisclosed flexibility is *pre-registration* [Ber12, CGD18], whereby all analytical choices are made before the data is collected and submitted to a verifiable registry. Pre-registration eliminates undisclosed flexibility, but still hides *analytic uncertainty*: the extent to which results are dependent on the particular analytic choices made. Different researchers who analyze the same data will often make different choices and get slightly—and sometimes widely—different results [SUM*18a]. If one of these researchers were to pre-register their analysis, their report would still convey an incomplete picture.

Some statisticians and methodologists have promoted the use of *multiverse analysis* to convey a much fuller picture of analytic uncertainty [STGV16, SSN19]. This approach consists of identifying a set of defensible analytical choices, performing all analyses corresponding to the possible combinations of such choices (possibly hundreds, thousands, or even millions) and reporting all outcomes, typically using summary visualizations. This idea is increasingly popular, with more and more academic papers reporting multiverse analyses; for example, a Google Scholar search for the term "specification curve"—a type of multiverse analysis visualization [SSN20]— returns 217 papers for the years 2019–2020.

However, multiverse analyses still raise many challenges, three of which serve as primary motivators for this work: (1) explaining and reporting the outcomes of hundreds or thousands of statistical analyses is difficult, especially when some of those analyses do not all point towards the same general conclusions [SSN20, DJS*19]; (2) literature specifically discussing the methodology of multiverse analyses is scattered across several fields and uses inconsistent terminology; and (3) visualization methods that have been proposed for helping to conduct and report multiverse analyses are similarly scattered across several fields and use inconsistent terminology. For example, considering the methodological literature (challenge 2), although the term *multiverse analysis* [STGV16] is recent, the core concept is found in older techniques under different names (e.g., *sensitivity analysis*, *robustness analysis*). These approaches have developed independently in different fields, leading to different terminology that often conflicts, which can make it difficult to communicate or reason about multiverse concepts. Similarly, static visual summaries of multiverse analyses (e.g., [STGV16, SSN19, PBI15]; Figure 1aef) or interactive visualizations of multiverse analyses (e.g., [DJS*19, LKAH20]; Figure 1dg) have been developed in different fields and under different names (challenge 3). Some more general visualization methods, like specification curve [SSN20], have been adopted in research papers, often with modifications, adaptations and improvements (e.g. [OP19a, BRRYD20]). Meanwhile, many papers use custom visualization methods for reporting multiverse analyses (e.g., [BNHC*20, BKB*20]). Some visualizations are domain-specific (e.g.: neuroimaging [Car12, BNHC*20]), or published in venues that may not be widely read outside of their field (e.g.: hydrology [Bie15]).

For a researcher who wants to report a multiverse analysis, these challenges make it hard to make informed choices about which visualizations to use; for a researcher who wants to study new multiverse visualization techniques, or teach the topic, it is hard to get a good overview of the state of the art. This article addresses the above challenges through a survey of academic articles that visualize multiverse analyses and related analyses. Importantly, our survey only covers ways visualization has been used to *report* multiple statistical analyses in an academic communication context. It does not discuss ways visualization has been used to help analysts *explore* multiple analyses, for example in the context of model steering and selection [DCCE19, MLMP17, CPCS19], ensemble data analysis [WHLS18], and visual parameter space analysis [SHB*14]. Our scope is further clarified in section 3.

In this survey, we (1) propose a conceptual framework and terminology for multiverse analyses that can be applied across fields, to support clarity when discussing this nascent family of concepts (section 3); (2) identify the *tasks* researchers try to accomplish with multiverse analysis visualizations, the questions one can seek to answer, and the central goal related to each category (section 5); and (3) classify multiverse visualizations into *archetypes*, assessing how well each *archetype* supports each *task*, their comparative limitations, key features, and what role they can play in an analysis (section 6). We close by discussing important design considerations surfaced by our survey—such as illusions of probability created by visualizing frequencies (subsection 7.1) and the largely unmet need to support validation and interpretation of multiverses (subsection 7.2)—as well as limitations and implications for future work (subsection 7.5). For visualization researchers looking to develop multiverse analysis visualizations, our work provides a foundational set of tasks for subsequent tools and techniques to support; for practitioners of multiverse analysis, our work provides a mapping between tasks they wish to accomplish and archetypes they can use to accomplish them.

2 An Example of Multiverse Analysis: are "Female" Hurricanes More Deadly?

To make our discussion throughout the rest of the paper more concrete, we will be using the multiverse analysis by Simonsohn et al. [SSN19] as a running example. We introduce this example here. The terms **in bold** are from our proposed multiverse analysis terminology and will be defined more precisely in section 3.

A 2014 study claimed that hurricanes whose names are female-gendered lead to more deaths, presumably because people do not take them as seriously as those with a male-gendered name [JSVH14]. However, later analyses of the same data called this finding into question [Mal14b, CC14, Mal14a]. It turns out that depending on how the analysis is carried out, it can be claimed that the data support the initial hypothesis, or the exact opposite. Simonsohn et al. [SSN19] conducted a multiverse analysis to investigate the space of possible analysis choices in more detail, and introduced the *specification curve* visualization (which we discuss

© 2022 The Author(s)
Computer Graphics Forum © 2022 The Eurographics Association and John Wiley & Sons Ltd.

in detail in subsection 6.2) to better understand the influence of analytical decisions on outcomes.

The subject of the original study and its re-analyses is a dataset of hurricanes, with their name and information such as number of victims. The multiverse as set up by Simonsohn et al. [SSN19] focuses on two **outcomes**: (1) *extra deaths*, the number of extra deaths occurring for hurricanes with female names compared to those with male names, and (2) a p-value reflecting the degree to which those extra deaths are surprising. Ultimately, their question is whether or not there is a statistically significant effect of hurricane name gender on extra deaths (i.e., if $p < .05$). Any single analysis (a **universe**) gives rise to a specific number of extra deaths, a specific p-value, and a single answer to that question. The whole **multiverse analysis report** makes it possible to assess whether those results are **sensitive** to different ways of conducting the analysis. For example, one might handle outliers from the dataset in different ways: *(i)* do not exclude any hurricane; *(ii)* exclude the most deadly hurricane, or *(iii)* exclude the two most deadly hurricanes. To reflect this, the multiverse has an *outliers* **parameter** that can take any of these **parameter values**. It also has other parameters, such as a *model* parameter for different model types that could be applied to the data, and a *femininity* parameter for different ways of operationalizing the gender of a hurricane name. Each universe is defined by a single combination of parameter values, which represents one unique way of analyzing the dataset. Simonsohn et al. report 1,728 universes, all produced by options they deemed to be reasonable.

If the outcomes of every universe were deemed to be practically equivalent the multiverse analysis need proceed no further, and one could infer simply that any of the examined choices can be selected without impacting final conclusions. In contrast, Simonsohn et al. found the estimated number of extra deaths attributable to the gender of hurricane names to range from about -1 to +12 (mean of 1.63), while only 37 out of the 1,728 universes (about 2%) yield $p < .05$. From those results they concluded that the proposed relationship between the gender of hurricane names and their deadliness is not robust to defensible analytical choices, and thus should not be accepted as correct on the basis of this evidence alone.

3 Definitions of Key Concepts

In this section, we introduce definitions that will serve to outline the scope of our survey. These are stipulative [Pap64] and are not meant to be authoritative.

Central to our survey is our definition of a multiverse analysis report:

> A **multiverse analysis report** is any statistical report that presents multiple analyses of the same raw dataset which answer the same question, are reported with a similar level of detail, and whose purpose is to learn from — or communicate insights about — that dataset.

Our definition is consistent with the way the term *multiverse analysis* (without the word *report*) is used by Steegen and Gelman [STGV16], who first introduced it and defined it as *"performing the analysis of interest across the whole set of data sets that arise from different reasonable choices for data processing."* The only previous usage we know of this full term is in Dragicevic et al. [DJS*19], though they do not explicitly define it. Our definition can be seen as a sharper version that more clearly distinguishes between multiverse analyses and related concepts.

Our definition has five key elements:

(1) *any statistical report*: this includes any narrative describing the result of a data analysis, in any format, even though in this survey we restrict ourselves to academic papers (see section 4). Thus, the focus is on what is reported, not what is analyzed. If multiple analyses are conducted but a single one is reported, as is commonly the case in empirical research [WVA*16b], then this cannot be considered to be a multiverse analysis report. Similarly, the process of building, selecting and tuning statistical models [DCCE19, MLMP17, CPCS19] is not within the scope of our definition, unless a report is written that uses multiple models to offer different perspectives on the same data.

(2) *of the same raw dataset*: the multiple analyses must be carried out on the same raw dataset. Carrying out the same analysis on different raw datasets does not qualify as a multiverse analysis. Examples are *(i)* ensemble data analysis, where multiple simulations are computed with different parameter settings, and the results are summarized and analyzed visually [WHLS18, SHB*14]; *(ii)* crowdsourced hypothesis testing, where multiple research teams conduct independent studies to answer the same research question [LJD*20]; and *(iii)* meta-analysis [GHF14], except when multiple meta-analyses are performed on the same set of studies [DBH19]. If different raw datasets (e.g., different experiments in a study) are subjected to the same set of analyses, there are as many multiverse analyses as there are raw datasets. A multiverse analysis can however involve the analysis of different *processed* datasets, as long as they all arise from the same *raw* dataset (e.g., when collapsing the levels of a variable in different ways; see the DATAVERSE example in [DJS*19]). Resampling techniques (e.g., bootstrapping; see the DANCE example in [DJS*19]) also generate multiple datasets from the same raw dataset, but in this survey we do not consider them as multiverse analyses, because their goal is only to assess statistical uncertainty in the original raw dataset.

(3) *answer the same question*: the multiple analyses need to answer the same question about the dataset. Statistical reports that use multiple analyses to answer different questions about a particular dataset (e.g., multiple subgroup analyses) do not qualify as multiverse analyses.

(4) *similar level of detail*: the multiple analyses need to be reported with a similar level of detail. A detailed data analysis followed by a cursory mention of additional analyses (e.g., "we redid the same analysis without outliers and obtained similar results") is not a multiverse analysis report. The outcomes from all analyses need to be reported with a similar level of detail. Similarly, a report that compares the goodness of fit of multiple statistical models but selects a single model to carry out the full data analysis does not qualify . However, we impose no lower limit on the number of analyses—a report with only two analyses can qualify as a multiverse analysis if the outcomes of both analyses are reported with a similar level of detail (e.g., [ESR17]). In addition, even if the analyses are heterogeneous in how they are conducted and reported (e.g., as in crowd-

© 2022 The Author(s)
Computer Graphics Forum © 2022 The Eurographics Association and John Wiley & Sons Ltd.

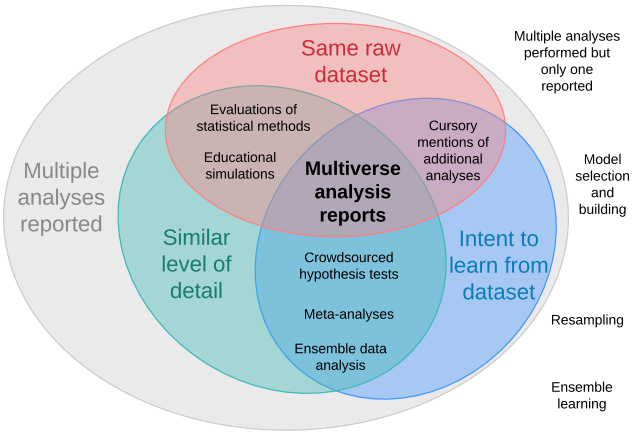


Figure 2: *Overview of the four major criteria making up our definition of multiverse analysis report (each criterion is an ellipse), and examples of cases that fulfill some but not all criteria.*

sourced analyses [BKB*20, BNHC*20]), they still qualify as long as all outcomes are reported in a similar fashion.

(5) *with the intent to learn from [...] that dataset*: several analysis types do not qualify multiverse analyses, as they do not have the goal to learn from the raw dataset itself: such examples that are not multiverse analyses include an evaluation of the coverage of different confidence interval procedures [Wah83], a sensitivity analysis carried out for model evaluation purposes [MLMP17], or an educational simulation illustrating how different analytical choices yield different outcomes [Fiv15]. Similarly, reporting multiple analyses with the intent to learn from — or communicate insights about — the analyses (not the datasets) would also not qualify. Furthermore, the entity that is expected to learn from the data must be a human. Thus, systems that learn from data by analyzing it in many different ways (e.g., ensemble learning algorithms [SR18]) are excluded, unless they explicitly convey the multiverse to human users. As stated initially, the multiple analyses need to be *reported*.

Figure 2 shows a Venn diagram where each ellipse stands for one of the criteria from our definition of multiverse analysis report. One criterion is not shown (i.e., that all analyses must answer the same question). The diagram regroups the edge cases we previously mentioned, and which fulfill most — but not all — of the criteria. We emphasize such edge cases because they help clarify the boundaries of our definition, and can speed up the classification of reports into multiverse or non-multiverse.

For the purpose of this survey, we additionally introduce the notion of trivial multiverse analysis report:

> A **trivial multiverse analysis report** is a multiverse analysis report with very few analyses and very little detail about each analysis, and which can be fully reported in the text without the need for tables or figures.

An example of a trivial multiverse analysis report is a paper that reports a p-value after excluding outliers, and a p-value without excluding outliers. Such analyses formally meet our definition of multiverse analysis report but will be excluded from our survey nonetheless, because little can be gained from visualizing them.

We draw from previous work [DJS*19] to define five basic elements that make up multiverse analysis reports, and which we will often refer to in this survey. In a multiverse analysis report:

> A **universe** or **analysis** is one of the multiple analyses that are conducted and reported in the multiverse analysis report.

> A **parameter** is a characteristic of the reported statistical analyses that varies across the multiverse.

> A **parameter value** is a possible value taken by a parameter. A synonym is *option* [DJS*19], but we use here the term *parameter value* for consistency with the rest of the terminology.

For example, suppose a paper uses three outlier exclusion methods to analyze data: (i) no exclusion; (ii) removing 3 standard deviations (SD) from the mean; and (iii) removing 2 SD from the mean. Thus, *outlier exclusion procedure* is a parameter of the multiverse analysis, and this parameter has three possible parameter values, each defining a different analysis or universe.

Similarly, in a multiverse analysis report:

> An **outcome** is a statistical result that is reported for all analyses in the multiverse.

> An **outcome value** is a possible value taken by an outcome.

In the previous example, suppose the paper reports a point estimate and a p-value for the main effect size of interest, computed for each of the three outlier exclusion methods. In this case, the multiverse analysis reports two outcomes (a point estimate and a p-value), and a total of six outcome values (two per universe).

A primary goal of multiverse analysis is to assess outcome sensitivity and robustness:

> **Outcome sensitivity** is the extent to which the values of an outcome vary across the multiverse.

> **Outcome robustness** is the opposite of outcome sensitivity, i.e., it is the extent to which the values of an outcome are stable across the multiverse.

Now we can define our main focus of investigation, which is the multiverse analysis visualization:

> A **multiverse analysis visualization** is any visual representation of the parameters, parameter values, outcomes and/or outcome values of multiple analyses in a multiverse analysis.

Visual representation means that at least some of the information is visually encoded [Mun14]. Thus, information conveyed exclusively via text and numerals (e.g., numerical tables) does not qualify, but hybrid representations that combine text or numerals with visual encodings (e.g., tabular visualization [PDF14]) qualify.

Finally, a last key concept central to this survey is the notion of *visualization archetype*:

> A **visualization archetype** (or simply archetype) is a class of multiverse analysis visualization designs that convey information about specific multiverse entities (i.e., parameters, parameter values, outcomes, and/or outcome values) using a specific combination of visualization idioms [Mun14].

© 2022 The Author(s)
Computer Graphics Forum © 2022 The Eurographics Association and John Wiley & Sons Ltd.

A visualization archetype thus defines a family of visualization designs that encode the same type of information in (more or less) the same manner. For example, a histogram of p-values and a histogram of effect sizes belong to the same archetype because they are both histograms of outcome values. However, a histogram of outcome values and a histogram of parameter values belong to different archetypes because they do not encode the same type of information, despite using the same visualization idiom.

4 Methodology

Our goal was to understand:

1. What tasks or analytical questions do researchers aim to perform or answer when reporting a multiverse analysis visualization?
2. What multiverse analysis visualizations do researchers use, and how do these visualizations support those tasks?

To answer these questions, we curated a corpus of research articles. To be considered for inclusion into our corpus, each article had to contain at least one multiverse analysis report, as well as at least one multiverse analysis visualization. We performed a systematic analysis of our corpus, to (i) derive a task taxonomy for multiverse analysis visualization, (ii) identify a set of visualizations archetypes, and (iii) analyze how well each archetype supports the tasks in our taxonomy.

4.1 Curating the Corpus

Multiverse analysis reports are being used across a wide body of literature in many different areas of science. We addressed the challenge of reviewing such a heterogeneous body of literature using a two-step approach (Figure 3). We first collected articles in a serendipitous fashion during the conduct of other research or reading activity, through social networks, or suggested by recommendation systems like Mendeley. This resulted in 52 *seed articles*. Since there was no agreed-upon or widely used term to refer to the concept of multiverse analysis, we extracted the terms used by the article authors, resulting in 8 terms (Table 1). We then used this list in a systematic literature search using the Google Scholar API through the *Publish or Perish* software [Har07] to find any documents with the terms appearing in the title, abstract or body text. We restricted the search to results published in 2015 or later, which for some keywords led to more results than the maximum of 1,000 returned by the API. To keep the number of articles we would need to analyze in detail manageable, we sorted each source list by the number of citations as counted by Google Scholar, and selected the first 20 items from each list. This led to 213 corpus candidates.

A second step consisted of checking whether each of the 213 corpus candidates was a research article. We replaced any item not passing this check with the next item from the respective source list. Using the definitions introduced in section 3, we then checked for each of the 213 corpus candidates that it (1) included at least one multiverse analysis report, (2) was not of a trivial nature, and (3) that the reported multiverse was visualized in some way. 36 of the seed articles and 19 articles discovered through the systematic literature search passed these checks for a total of 43 articles which form our final corpus (12 came up through multiple sources as detailed in Table 1).

More details on the corpus as well as the source lists from the systematic search are in the supplemental material.

Search type – keyword	Search results	In final corpus	In both I. and II.
I. Serendipitous (seed articles)	53	36	12
II. Systematic:	>4,893	19	12
– multiverse analysis	198	7	6
– specification curve	298	8	6
– vibration of effects	144	4	2
– crowdsourced analysis	264	3	2
– robustness analysis	>1,000	0	0
– multimodel analysis	989	0	0
– perturbation analysis	>1,000	1	0
– sensitivity analysis	>1,000	0	0

Table 1: *Quantitative background on the corpus curation, including the number of search results per search type (serendipitous vs. systematic) and per search term, the number of papers that met our inclusion criteria, and the number of papers from the systematic search that were already in our initial corpus of seed articles.*

4.2 Extracting Tasks on Multiverse Analysis Visualizations

To derive a taxonomy of the tasks researchers can perform with a multiverse analysis visualization, we performed a detailed analysis of parts of a subset of five articles in our corpus. These five were selected because their goal was to introduce a form of multiverse analysis as a *general method* rather than to use multiverse analysis to report specific findings [SSN19, STGV16, PBI15, YH17, SUM*18b]. Each paper analyzed one or more datasets as a demonstration of the technique being introduced, as well as detailing reasoning and broader implications of their methodology.

For each of these articles, we extracted all figures that contained a multiverse analysis visualization, as well as any relevant text either present directly on the figure or in the figure caption. We also searched for all references to that figure in the article's main text and extracted all statements about the figure from the corresponding paragraphs, as well as the ones preceding and following it. Each captured passage was split into individual quotes, then copied onto the digital equivalent of sticky-notes in a collaborative whiteboard platform (Miro board, `www.miro.com`). Three authors conducted an affinity diagramming exercise to cluster the quotes into themes, which facilitated the identification of common tasks that could be performed using multiverse analysis visualizations. A selection of quotes relevant to each task are presented in section 5.

Once all quotes from the initial articles were processed and a draft task taxonomy formed, we expanded and continued the analysis with additional articles from our corpus to ensure saturation was reached. Articles were chosen from reviewing the visualizations and discussion notes from our entire corpus, with a focus on selecting papers that were most likely to challenge our existing conceptions, judged from the distinctiveness of their associated visualizations and the topic of the articles themselves. The analysis of additional articles presenting interactive visualizations [DJS*19, LKAH20] and theoretical considerations of multiverse analysis [DGGS20] inspired the definition of the last category added to our taxonomy (Validate, subsection 5.5). Analysis of an

© 2022 The Author(s)
Computer Graphics Forum © 2022 The Eurographics Association and John Wiley & Sons Ltd.

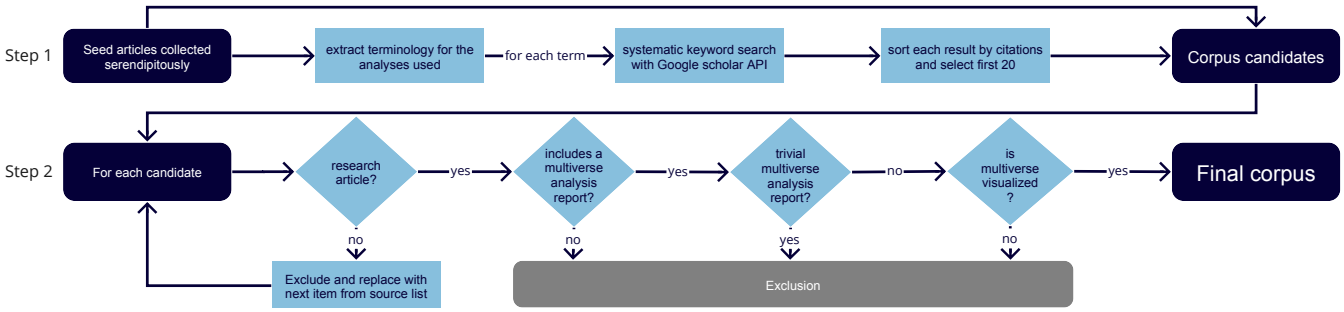


Figure 3: *Overview of our curation process. In step 1, we curated a corpus of candidates by combining serendipitously discovered articles with a systematic keyword search. In step 2, we analyzed each candidate to identify all research articles that contain a non-trivial multiverse analysis report and illustrate that report with some form of visualization.*

additional set of 7 articles [LKJA*19, ASGP18, PVB19, BKB*20, BNHC*20, Car12, ODT12], which featured distinctively different visualizations compared to the already included ones—and thus could likely challenge our task taxonomy—did not generate new tasks, categories, or change our taxonomy structure.

We present the outcome of our task analysis in section 5. The source material, including a PDF export of the Miro boards, can be found in the supplemental material.

4.3 Identifying Visualization Archetypes

To accomplish our second goal—identify multiverse analysis visualization archetypes and assess their capacity to support the tasks in our task taxonomy—we reviewed our full corpus of 43 articles, and extracted any figures and tables that initially appeared to satisfy our multiverse analysis visualization criteria.This resulted in a collection of 126 visualizations, which we trimmed so as to keep at least one representative figure for every distinct visual style present, as judged by all authors. The resulting set of 85 prospective archetypes was further reduced through closer review, with 16 being excluded as they were not actually multiverse analysis visualizations (e.g., visualizations of simulation studies, Sankey diagrams of a literature review), leaving 69 visualizations for further analysis.

To further distinguish between visualizations that supported different multiverse analysis tasks to some extent, from ones that only varied aesthetically, we conducted an in-depth iterative coding process. In each coding cycle we picked one of the prospective archetypes, then reviewed the source paper. We then graded the visualization's support for each of the tasks in our taxonomy on a scale of 0-3 (as detailed below), assuming a multiverse of similar proportions than that featured in the visualization. In each cycle, if a visualization was found to be equivalent to a previously scored visualization, it was labeled to be a variant of the same archetype and excluded from re-scoring.

We defined a 0-3 grading system as: 0 = no support for this task; 1 = information required for task is present, but requires a large amount of effort or mental calculations, or supports the task minimally; 2 = tasks are sometimes well supported and sometimes not, depending on factors that naturally vary between multiverses; 3 = supports the task in a way that makes it reasonably fast and easy to complete, usually through clear visual features or explicit encoding of relevant information into distinct visual channels. All scores disregard the learning-curve that may be required to use a visualization, and so adopt the perspective of a reader already familiar with that type of visualization. All scores were reviewed by at least two authors after all visualizations were coded, with any disagreements resolved by discussion until consensus was reached.

The primary results of this analysis are reported in section 6. The full set of visualizations reviewed and scored are available as supplemental material.

5 Taxonomy of Analysis Tasks

We identified twelve tasks that can be performed using a multiverse analysis visualization, summarized in Table 2. We organize these tasks into five analytical categories, with each category encompassing a general class of questions and goals that are common to most multiverse analyses. We denote each task definition as follows:

> **Category Name ▷ Task Name**: definition of this task.

In-text we use the notation Category Name ▷ Task Name to refer to specific tasks. We have given the categories and tasks a logical ordering primarily to make them easier to describe and understand; this order does not necessarily reflect the order in which these tasks are carried out or reported. For each category we provide an *Example Question* based on our running example from Simonsohn et al. [SSN19] as well as sample quotes taken from the corpus that were used to identify and synthesize these tasks.

5.1 Composition: Understand Composition of the Multiverse

Example Question: What are the different methods used to exclude outliers in this multiverse?

Goal: Understand the components and processes that define and makeup this multiverse.

Tasks in this category can involve descriptions of the dataset source, how the data was processed, the included variables in the data, and what analytical choices are being considered (parameters and their parameter values). These tasks lay the groundwork necessary for the later sense-making process of drawing conclusions from the multiverse analysis. This category is unique in that it does

© 2022 The Author(s)
Computer Graphics Forum © 2022 The Eurographics Association and John Wiley & Sons Ltd.

Category	Task
Composition	**Composition** ▷ **Process**: understand the process that defines and creates the universes being considered. **Composition** ▷ **Parameters**: understand the definition and composition of universe parameters and parameter values.
Outcome	**Outcome** ▷ **Range**: assess range or spread of outcome values across all universes. **Outcome** ▷ **Frequency**: assess overall frequency of outcome values across all universes.
Connect	**Connect** ▷ **OutcomeRange**: connect parameters to outcomes by comparing similarity or range of outcome values across a subset of universes defined by a specific parameter value. **Connect** ▷ **OutcomeFrequency**: connect parameters to outcomes by comparing frequency of outcome values across a subset of universes defined by a specific parameter value. **Connect** ▷ **SpecificOutcomes**: connect parameters to outcomes by examining specific outcome values of interest and identifying parameter values that lead to those outcomes.
Connect Combinations	**ConnectCombo** ▷ **OutcomeRange**: connect combinations of parameters to outcomes by comparing range of outcome values across subsets of universes defined by parameter values. **ConnectCombo** ▷ **OutcomeFrequency**: connect combinations of parameters to outcomes by comparing frequency of outcome values across subsets of universes defined by parameter values. **ConnectCombo** ▷ **Idiosyncratic**: connect combinations of parameters to outcomes according to idiosyncratic patterns particular to a given visualization or analysis.
Validate	**Validate** ▷ **Metrics**: assess validity metrics of universes or com-pare metrics across parameter values. **Validate** ▷ **Details**: assess validity of universes by examining the underlying details of analyses in each universe to interrogate their validity.

Table 2: *Overview of the taxonomy for multiverse analysis tasks derived from the multiverse analysis visualizations in our corpus.*

not consider the outcomes of any analyses. We refer to this category as **Composition**.

In most published reports this category of tasks is addressed solely through narrative descriptions, often in the form of lists in the text itself or as a table (see Figure 7). But as the composition of a multiverse grows in complexity, some authors choose to use visualizations to facilitate navigation and understanding of that complex structure. Two notable examples are the computation schematic of Patel et al. [PBI15] (Figure 13), and elements of the Boba interactive interface [LKAH20] (Figure 15).

> **Composition** ▷ **Process**: understand the process that defines and creates the universes being considered.

This task concerns the details and processes involved in creating individual universes, and thus the multiverse altogether. This can generally include data sources and data collection procedures, any processing of the data that is common to all universes, criteria for selecting outcomes of interest, and any other contextually relevant and important information of this kind.

For example, Patel et al. [PBI15] used the following narrative description to explain a few key steps in their process: *"First, we downloaded 417 self-reported, clinical, and molecular measures with linked all-cause mortality information in participants from NHANES 1999-2004. [. . .] We chose variables of interest that had data on at least 1,000 participants and at least 100 death events during follow-up."* In that work, the authors both described the process in the text and illustrated the steps in a diagram — the computation schematic visualization (Figure 13).

> **Composition** ▷ **Parameters**: understand the definition and composition of universe parameters and parameter values.

This task involves understanding how parameters and parameter values included in the multiverse are defined, as well as how they can combine to form universe specifications. In the hurricane multiverse (section 2), one parameter is *model*, with two parameter values: *negative binomial* and *log-normal*. In that multiverse, every combination of parameter values is considered valid, so there are no complex relationships between parameters and parameter values that need to be communicated. However, some multiverse analyses include more complex parameter contingencies, e.g., selecting one value for parameter A could render some available values for parameter B invalid. Communicating such relationships falls within the scope of this task as well.

5.2 Outcome: Assess Outcome Sensitivity

Example Question: Is the relationship between hurricane name genders and model-predicted fatalities stable across combinations of defensible analytical choices?

Goal: Assess the extent to which important outcomes vary among alternative analytical choices (sensitivity or robustness - see (see definitions in section 3).

The topic of this category is the fundamental concern of multiverse analysis: if all considered analytical choices lead to effectively the same conclusions, then there is no need to proceed any further in the multiverse analysis. If outcomes are not sensitive, one can conclude that which of the considered choices one prefers does not matter, as the ultimate conclusions one would reach are

© 2022 The Author(s)
Computer Graphics Forum © 2022 The Eurographics Association and John Wiley & Sons Ltd.

the same regardless. For example, in the hurricane study, only 37 of the 1,728 universes result in a p-value below .05, which indicates that some universes produce outcome values that differ substantially from the majority.

Importantly, how sensitive an outcome is depends upon context and expert judgment in the domain of the analysis. Assessing to what extent outcome values vary across a multiverse typically requires judgments of practical magnitude that are domain-dependent and subject to the analyst's interpretation. For example, if an analyst considers a certain range of effect sizes to be practically equivalent, then the effect size outcome is robust if it remains within that range. Similarly, if an analyst hinges their interpretation of p-values on a statistical significance threshold, then the p-value outcome is sensitive if, across universes, outcome values fall on both sides of that threshold.

> **Outcome ▷ Range**: assess range or spread of outcome values across all universes.

One way to assess outcome sensitivity is to examine the similarity (or spread, or range) of outcome values that occur within the multiverse, which is the goal of this task.

Simonsohn et al. [SSN19] describe the results of completing this task: *"The point estimates range from -1 to +12 additional deaths."* Similarly, Steegen et al. [STGV16] write: *"For fiscal political attitudes ... the remaining choice combinations lead to p values across the entire range from .05 to 1.0."*

> **Outcome ▷ Frequency**: assess overall frequency of outcome values across all universes.

Another way to assess outcome sensitivity is by examining the frequency or proportion of specific outcome values that occur within the multiverse. However, there is more than one way to interpret outcome frequencies, which necessitates a nuanced consideration of this task.

The first interpretation of outcome frequency is *probabilistic*; i.e. treating frequencies as estimates of relative likelihood, with outcomes that occur in more universes deemed more plausible than ones that occur in fewer universes. For example, Simonsohn et al. [SSN19] state: *"A researcher motivated to show a negative point estimate would be able to report twenty different specifications that do so, but the specification curve shows that a negative point estimate is atypical."* Simonsohn et al. [SSN19] even introduce a technique for calculating a p-value of statistical significance for the multiverse as a whole, which treats the selection of analytical choices as a probabilistic sampling process.

Alternatively, a *possibilistic* interpretation of outcome frequency is illustrated in Steegen et al. [STGV16]: *"If no strong arguments can be made for certain choices, we are left with many branches of the multiverse that have large p values. In these cases, the only reasonable conclusion on the effect of fertility is that there is considerable scientific uncertainty. [...] When only one choice is clearly and unambiguously the most appropriate one, variation [in outcomes] across this choice is uninformative."* In other words, frequency information can indicate the possibility that something could be true, but cannot be used to determine what outcomes are more or less likely. The second part of this quote goes even further, implying that relative frequency of outcomes for some options should not be interpreted as encoding any relevant meaning.

Consideration for how a reader could, or should, interpret outcome frequencies is important for visualization design, as we suspect different visualizations may invite incorrect probabilistic interpretations. We discuss this issue further in subsection 7.1. Note that this task is closely matched to what Amar et al. [AES05] refer to as a "Characterize Distribution" task.

5.3 Connect: Connect Parameters to Outcome Values to Identify Sources of Sensitivity

Example Question: Do some values within the "dropping outliers" parameter lead to consistently larger outcome values of model-predicted fatalities?

Goal: Identify which analytical choices cause outcomes to differ across universes.

This category explores potential relationships between individual parameters, parameter values, and outcome values. When outcomes have been determined to be sensitive to analytical choices (subsection 5.2), one can seek to determine which choices produce this sensitivity. For instance, it could be that only some small subset of parameter values produces a divergent outcome, in which case one might wish to focus on critically analyzing these few choices in greater detail. Further attention could either involve additional tasks described in this framework, or deeper theoretical considerations.

> **Connect ▷ OutcomeRange**: connect parameters to outcomes by comparing similarity or range of outcome values across a subset of universes defined by a specific parameter value.

As with the previously described Outcome ▷ Range task, this task examines the similarity or overall range of outcome values within a multiverse, but with the added detail of conditioning (subsetting) on a parameter or parameter value. It is this additional point that allows for sources of sensitivity to be identified, and for the impact of different parameter values to be compared.

An example from Steegen et al. [STGV16] describes two parameters identified as not being the primary drivers of outcome sensitivity: *"The different exclusion criteria and cycle day estimation options do not seem to have a large impact on fluctuation in the statistical conclusion."* In contrast, Silberzahn et al. [SUM*18b] describe the identification of two parameters that are sources of outcome sensitivity: *"The teams also varied in their approaches to handling the nonindependence of players and referees, and this variability also influenced both median estimates of the effect size and the rates of significant results."*

> **Connect ▷ OutcomeFrequency**: connect parameters to outcomes by comparing frequency of outcome values across a subset of universes defined by a specific parameter value.

© 2022 The Author(s)
Computer Graphics Forum © 2022 The Eurographics Association and John Wiley & Sons Ltd.

As with the previously described Outcome ▷ Frequency task, this task examines the frequency of outcome values, but now conditioned (subsetted) on a parameter or parameter value.

Silberzahn et al. [SUM*18b] compare the frequency of outcomes across the parameter *model form*: *"Fifteen teams used logistic models, and 11 of these teams found a significant effect [. . .] Six teams used Poisson models, and 4 of these teams found a significant effect."* Steegan et al. [STGV16] use a more roughly estimated proportion: *"For religiosity [. . .] most data sets constructed under the second option for relationship assessment (R2) yield a nonsignificant interaction effect."*

> **Connect ▷ SpecificOutcomes**: connect parameters to outcomes by examining specific outcome values of interest and identifying parameter values that lead to those outcomes.

Another approach to identifying sources of sensitivity is to instead focus on specific outcome values, and find what parameter values produce them. This can be particularly important when some outcome values are more consequential than others, such as when some outcome values imply a therapeutic intervention is harmful.

Simonsohn et al. [SSN19] considered negative effect sizes in this way: *"[. . .] we can see that obtaining a negative point estimate requires a fairly idiosyncratic combination of operationalizations."*

5.4 Connect Combinations: Connect Combinations of Parameters to Outcome Values to Identify Complex Relationships that Lead to Sensitivity

Example Question: Do the outcomes associated with the choice of model form strongly depend upon the choice of dropping outliers? In other words, do the parameters interact?

Goal: Identify which combinations of analytical choices cause outcomes to differ across universes.

In this category, the relationship between outcomes and analytical choices is further explored and characterized in ways that go beyond what was considered in category Connect (see subsection 5.3).

The primary additional factor is considering combinations of parameters and parameter values. As a simplified example, if some model forms are more sensitive to outliers, then any parameter value related to excluding outliers could theoretically have a combined effect that would not be noticeable when examining the parameter values individually.

> **ConnectCombo ▷ OutcomeRange**: connect combinations of parameters to outcomes by comparing range of outcome values across subsets of universes defined by parameter values.

This task extends task Connect ▷ OutcomeRange by considering combinations of parameter values, rather than treating parameters as effectively independent from one another. While we primarily consider the combination of only two parameter values at a time, conceptually there is no reason that more complex relationships might exist with even more parameter values, just as in a traditional multivariate analysis. However, just as in traditional multivariate analysis, it is extremely difficult to cognitively and intuitively consider higher-order interaction effects, and a three-way interaction is the most complex relationship we have an example for in our corpus.

Steegen et al. [STGV16] describe a two-way interaction effect between parameters thusly: *"Using the third option for relationship status assessment (R3) leads to more fluctuation, depending on the choices for the other processing steps."* In the report from Young et al. [YH17], the combined effect of two choices is a centrally important finding: *"Why do these estimates vary so much? Why is the distribution so non-normal? What combinations of control variables are critical to finding a positive and significant result? [. . .] In order to draw robust conclusions from these data, one must make a substantive judgment about two key modeling assumptions: the inclusion of race and marital status."*

> **ConnectCombo ▷ OutcomeFrequency**: connect combinations of parameters to outcomes by comparing frequency of outcome values across subsets of universes defined by parameter values.

This task similarly extends task Connect ▷ OutcomeFrequency by adding the consideration of a combination of multiple analytical choices, with a focus on the relative frequency of outcomes.

Steegen et al. [STGV16] provide an example of this task where proportion is considered with rough approximations: *"The first and third options (R1 and R3) consistently lead to a significant interaction effect in combination with the first and second option for fertility assessment (F1 and F2) and to a nonsignificant interaction effect in combination with F5, whereas data sets constructed under R1 or R3 in combination with F3 or F4 lead to more fluctuating conclusions, depending on the other choices for data processing."*

> **ConnectCombo ▷ Idiosyncratic**: connect combinations of parameters to outcomes according to idiosyncratic patterns particular to a given visualization or analysis.

This task encompasses a variety of special relationships and patterns that are described throughout the corpus. These patterns are generally specific to certain visualizations, and we discuss these cases in greater detail in section 6. However, as a brief example we consider here the most commonly described concept of modality/multi-modality of the outcome value distribution.

In a univariate analysis, distributions can have one or more modes, which are the value(s) that occur most often in that distribution. When all outcome values from a multiverse are analyzed as a distribution, there can be a single mode representing the value that the largest number of universes produce, or the distribution can be multi-modal. In Young et al.'s report [YH17], multi-modality is considered to possibly indicate that some parameter value, or combination of parameter values, are responsible for disparity of the outcome values. Having identified such parameter values, the authors state: *"In essence, there are two distinct modeling distributions to consider"*. This concept of modality is also described by Patel et al. [PBI15], referred to as "modality in the Vibration of Effects", and is given an equivalent interpretation: *"We observed three modes in the association between triglyceride levels and mor-*

© 2022 The Author(s)
Computer Graphics Forum © 2022 The Eurographics Association and John Wiley & Sons Ltd.

tality [...] The multimodal plots indicated that total cholesterol and diabetes were driving these modes."

5.5 Validate: Validate the Multiverse

Example Question: Are all combinations of parameter values equally reasonable or defensible? For instance, does model fit, or other statistical diagnostic metrics, suggest one model type may provide more reasonable estimates?

Goal: Determine the validity, reasonableness, plausibility, or defensibility of the multiverse overall.

This category is concerned with critically evaluating the validity of the constructed multiverse. Analytical choices and associated universes can be re-examined in light of additional insights gained from the multiverse analysis process itself. This can include examining model fits, statistical/predictive diagnostic criteria, re-evaluation of the handling of the underlying dataset, or other investigation of individual universes or sets of universes.

Conducting an analysis can lead one to reconsider some of the decisions that were included in the multiverse, or to realize other parameters and parameter values should be considered as well. Early work in multiverse analysis, such as that of Simonsohn et al. [SSN19] and Steegen et al. [STGV16], primarily considered analytical choices that could be considered defensible prior to examining the data, or at least without using the data to evaluate the appropriateness of the analytical choices themselves. However, an analyst could reasonably come to question whether some outcomes should be given greater weight than others, which would mean that some universes are not considered equally defensible, even if they cannot be definitely excluded as inappropriate.

This category ultimately represents a stage of reflection that would ideally come before final interpretation of the multiverse analysis results. While conducting an analysis, this might lead one to reconsider some of the decisions that were included, or to realize other parameters and parameter values that should be considered. It could also suggest that some outcomes should be given greater weight than others, which would mean that some universes are not considered equally defensible, even if they cannot be definitely excluded as inappropriate.

This category and its associated tasks are described in a broader and less exacting way, as there were fewer examples of these tasks and visualizations to support them.

> **Validate ▷ Metrics**: assess validity metrics of universes or compare metrics across parameter values.

This task considers the validity of universes that make up the multiverse using some form of metric, such as model fit metrics. For example, some model types may produce better model fits overall, or the model fits may vary across parameter values. Model fit is the only specific example of this task we identified in the corpus, but other metrics could certainly be used for a similar purpose.

In Boba [LKAH20], support for this task is described: *"Do we have evidence that certain outcomes are less trustworthy? We toggle the color-by drop-down menu so that each universe is colored by its model quality metric [...]. The large estimates are almost exclusively coming from models with a poor fit. We further verify the model fit quality by picking example universes and examining the model fit view [...]. The visual predictive checks confirm issues in model fit, for example the models fail to generate predictions smaller than 3 deaths, while the observed data contains plenty such cases. [. . .] we have reasons to be skeptical of the large estimates."*

> **Validate ▷ Details**: assess validity of universes by examining the underlying details of analyses in each universe to interrogate their validity.

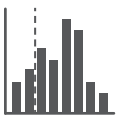
This general task is about investigating universes in a level of depth that may be more typical with traditional analyses, but which is difficult to do with an entire multiverse. This task is instead concerned with diving into either single universes or small sets of universes in greater detail, to allow for the richness and detail of a traditional analysis to be able to inform the construction and assessment of validity of the multiverse overall.

This task has some degree of limited support in a few different visualizations, but the primary source for the identification of this task is from the Explorable Multiverse Analysis Reports (EMAR) [DJS*19], an interactive media where balancing depth and richness with the comprehensiveness of a multiverse analysis is a primary design goal; for example: *"Four aspects of the analysis can be changed by the reader, which has the effect of immediately updating the two plots and some text elements such as explanations and figure captions."* While the technique was not designed or explicitly described with the goal of examining the validity of a multiverse, it is one of few multiverse visualizations that demonstrate how this task might be supported.

6 Multiverse Visualization Archetypes and Systems

We describe the set of multiverse visualization archetypes—families of similar visualizations designs— identified in our analysis, along with the tasks they support. We also discussed two interactive visualization systems designed to support multiverse analysis. A visual summary is shown in Figure 4 (see definition in section 3, and process in section 4).

6.1 Outcome Histogram



The *outcome histogram* conveys the frequency of the different outcome values that occur within a multiverse for a particular outcome, so that each individual universe outcome value is counted once. In Figure 5, the x-axis encodes the outcome values (here, point estimates of extra deaths for female hurricanes in the example of section 2), while the y-axis encodes the number of times binned outcome values occur within the multiverse. The dotted line serves as a visual aid to highlight the effect size of zero, which can be interpreted in the context of our running example as implying that there is no net effect of hurricane name femininity on predicted fatalities.

The outcome histogram allows a viewer to easily and simultaneously complete both Outcome (subsection 5.2) tasks: Outcome ▷ Range and Outcome ▷ Frequency. This is made possible

© 2022 The Author(s)
Computer Graphics Forum © 2022 The Eurographics Association and John Wiley & Sons Ltd.

	Name	Section	Icon	Composition ▸ Process	Composition ▸ Parameters	Outcome ▸ Range	Outcome ▸ Frequency	Connect ▸ Outcome Range	Connect ▸ Outcome Frequency	Connect ▸ Specific Outcomes	Connect Combo ▸ Outcome Range	Connect Combo ▸ Outcome Frequency	Connect Combo ▸ Idiosyncratic	Validate ▸ Metrics	Validate ▸ Details	References
Archetypes	Outcome Histogram	6.1		0	0	3	3	0	0	0	0	0	0	0	0	[STGV16] [DGGS20] [PVB19] [BI16] [DMH*18] [VKT19] [BI16] [CT16]
	Outcome Curve	6.2.1		0	0	3	2	0	0	0	0	0	0	0	0	[SSN20] [Coo18] [JKN18] [VKT19] [BYO19] [OP19a] [DS18] [SSN19]
	Universe Specification Panel	6.2.2		0	2	0	0	0	0	0	0	0	0	0	0	[SSN20] [HCM13] [GHF14] [SUM*17]
	Descriptive Specification Curve	6.2.3		0	2	3	2	3	2	3	2	1	3	0	0	[SSN20] [JKN18] [OP19a] [OP19b] [ODP19] [DGGS20] [BRRYD20] [VKT19] [RES17]
	Outcome Density Plot	6.3		0	1	3	3	2	2	1	2	2	3	0	0	[YH17] [LKJA*19] [HS06] [ODT12] [You18] [MY18]
	Vibration of Effects Plot	6.4		0	0	3	2	2	2	1	1	1	3	0	0	[PBI15] [DGGS20]
	Outcome Matrix	6.5		0	1	3	2	2	2	3	2	2	3	0	0	[STGV16] [CJT19] [DGH*18] [DKBK19] [DS18]
	Multiverse Computation Schematic	6.6		3	3	0	0	0	0	0	0	0	0	0	0	[PBI15]
Systems	Explorable Multiverse Analysis Reports	6.7.1		0	2	1	1	1	1	1	1	1	0	0	3	[DJS*19]
	Boba	6.7.2		3	3	3	3	3	3	3	1	1	3	3	0	[LKAH20]

Figure 4: *Overview of the archetypes and interactive systems described in section 6. Shaded cells indicate how well an archetype or system supports an analysis task in our taxonomy, on a scale of 0 (not supported) to 3 (fully supported).*

because both the full range of outcome values, as well as their proportions, are explicitly encoded in the plot. For instance, Figure 5 allows to identify that the most common outcome values are near zero, and that there are also many more results above zero than below it. One can also see that the the positive effect sizes go to greater magnitudes than the negative ones (+12 versus -1). However, with no mapping of parameters and options to outcomes, the viewer cannot explore which analytical choices are responsible for this variation.

Though frequency is a fundamental feature of the outcome histogram, and Steegen et al. themselves were clearly aware of the dangers of a probabilistic interpretation (as described in subsection 5.2), under a strictly possibilistic interpretation the existence of even one seemingly valid universe with a given outcome value is evidence that outcome cannot be ruled out. This suggests a potential issue with this (and other) frequency-based encodings: they may invite unintended or incorrect interpretations of multiverse outcomes. We discuss this further in subsection 7.1.

The outcome histogram is a general approach that we encountered frequently in our corpus, e.g. [STGV16, DGGS20, PVB19, BI16, DMH*18, VKT19]. We also note one variation where the outcome is a *p*-curve [BI16], while Cirillo et al. reported multiple varieties of this type [CT16].

6.2 Descriptive Specification Curve

The *descriptive specification curve* is an example of a *composite visualization*—a visualization that is made up of two or more linked components, each of which could individually function as stand-alone visualizations on their own. Some composites feature

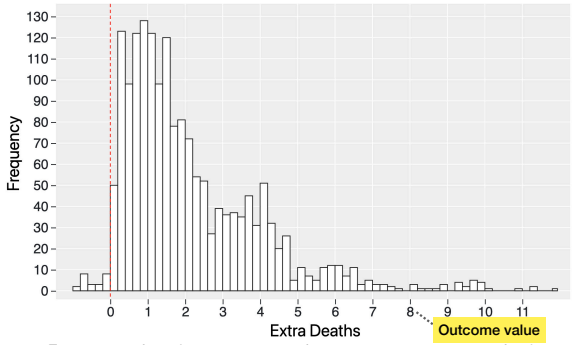


Figure 5: *Example of an outcome histogram. Recreated after Steegen et al. [STGV16], but using the hurricane dataset (section 2). The x-axis encodes outcome values (effect size estimates), while the y-axis shows the count across the multiverse.*

super-additive functionality, which is when a composite visualization supports more tasks than all of the individual components considered separately, and this archetype is the primary example of this concept. Note that the term *specification curve* has been ambiguously used in the literature to refer to a multiverse analysis, the full composite (Figure 6), or just the top panel (Figure 6a). Following Simonsohn et al. [SSN20], we use *descriptive specification curve* to refer to the full composite. We first review each component individually before discussing the composite.

6.2.1 Outcome Curve (Component)



The core component of the descriptive specification curve is the *outcome curve* (Figure 6a). The y-axis

© 2022 The Author(s)
Computer Graphics Forum © 2022 The Eurographics Association and John Wiley & Sons Ltd.

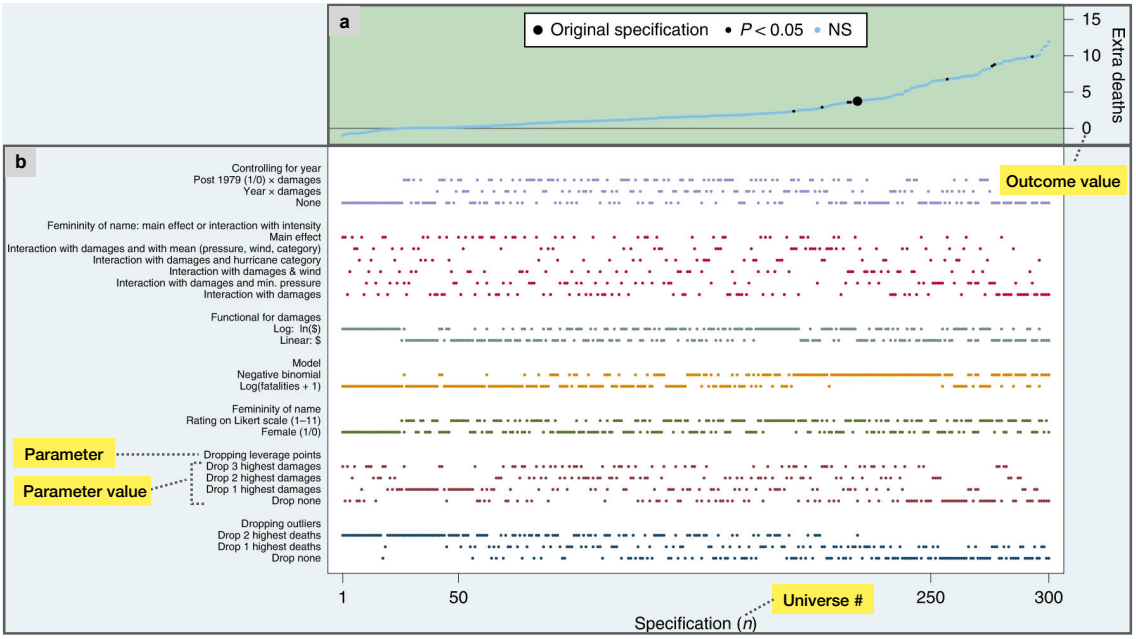


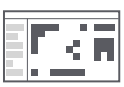
Figure 6: *Example of a specification curve [SSN19]. We treat the full figure as a composite, made up of two components: (a) an outcome curve, (b) an universe specification panel. The composite visualization has super-additive functionality, enabling tasks that neither component supports by itself. 300 universes are shown here, out of the full multiverse of 1,728. The 50 universes with the smallest and largest outcome values are shown, along with a random sample of 200 other universes.*

encodes the outcome values (here, extra deaths), and universes are sorted along the x-axis according to outcome value, giving this visualization its distinctive shape. In the design of Simonsohn et al. [SSN20] shown in Figure 6, dot color encodes a second outcome (black for statistically significant and blue for non-significant). In addition, due to limited horizontal space and the large size of their multiverse, the authors chose to only display a subset of the 1,728 universes: only those with the top and bottom 50 outcome values are shown, along with 200 other randomly sampled universes.

This visualization supports the same two tasks as the histogram of outcomes, i.e., Outcome ▷ Range and Outcome ▷ Frequency. The outcome curve resembles a cumulative distribution function (CDF) with the axes swapped. Because frequency is not explicitly encoded, the task Outcome ▷ Frequency is more difficult and less precise, especially when values being compared are not adjacent.

The outcome curve is commonly presented as a stand-alone visualization, e.g., [Coo18, JKN18, VKT19, BYO19, OP19a, DS18], especially in papers explicitly reporting a *specification curve analysis*. Simonsohn et al. [SSN19] include 3 examples of the curve presented alone, with only one example of the full descriptive specification curve.

6.2.2 Universe Specification Panel (Component)



The second component of the descriptive specification curve is the *universe specification panel* (Figure 6b). It consists of a tabular visualization [PDF14] where columns are individual universes, and rows are parameter values clustered by parameter. Columns may be sorted by outcome value, although outcome values themselves are not shown in this component. A cell in this table indicates when a universe (column) includes a given parameter value (row) in its specification. As this visualization shows no outcome values, it only supports task Composition ▷ Parameters.

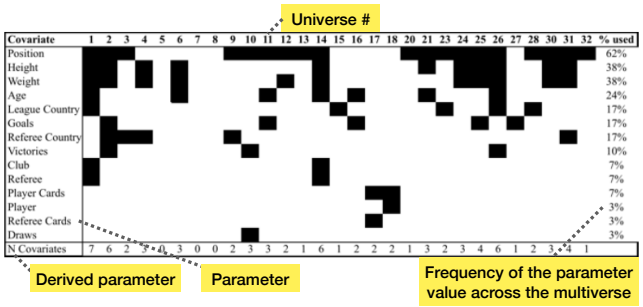


Figure 7: *A variant of a universe specification panel [SUM* 18a]. Each column is a team of analysts (i.e. a universe) having analyzed the same dataset using different analytical choices, as defined by black cells indicating the selection of parameters values. The bottom row indicated the number of parameter values in each universe, and the rightmost column indicates the frequency of a given parameter value across the sparse multiverse.*

Figure 7 shows a variant of this archetype, designed for sparse multiverses, i.e., where not every combination of parameter values is used. The plot indicates on the far-right column how many of the universes has each parameter value enabled. In this example, columns are also sorted by the number of covariates included in the analysis performed by a team.

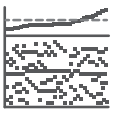
Note that the number of covariates is not a free parameter, but is

© 2022 The Author(s)
Computer Graphics Forum © 2022 The Eurographics Association and John Wiley & Sons Ltd.

instead a function of other parameter values (which would be, for example, one parameter per covariate that indicates if it was used in the analysis). We refer to this as a **derived parameter**. Derived parameters can be visualized the same as any other parameter.

There are a number of other examples of this archetype in our corpus, e.g. [HCM13, GHF14], but all have equivalent task support.

6.2.3 Descriptive Specification Curve (Composite)



Combined together on a common x-axis, the components above form the full composite *descriptive specification curve* (Figure 6), which allows the viewer to connect outcome values to analytical choices.

The composite supports all the tasks that its individual components support, but also supports all tasks in the Outcome and Connect categories (subsection 5.2 and subsection 5.3). Consider, for instance the *dropping outliers* parameter: at the bottom of the specification panel (Figure 6b), the eye is drawn to a continuous pattern of dark blue dots indicating that the *drop 2 highest deaths* parameter value (i.e. exclude the two deadliest hurricanes Katrina and Audrey) leads to the all of the lowest outcome values. The viewer can read up to see that all of the outcome values below zero are associated with this parameter value (Connect ▷ OutcomeRange). Alternatively, if the viewer were interested in outcome values below zero, they could have started in the Outcome Curve (Figure 6a) and read down (Connect ▷ SpecificOutcomes), leading to the same observation, with other similar patterns observed for the *controlling for the year* parameter value *none* (purple), or *model* parameter value *log(fatalities + 1)* (yellow).

The tasks ConnectCombo ▷ OutcomeRange and ConnectCombo ▷ OutcomeFrequency can be completed in the same manner, but with less ease because columns that satisfy a combination of more than one parameter values (e.g., *controlling for year = year × damages* and *feminity of name = rating on Likert scale (1–11)*) are not clustered together, making it difficult to identify whether the corresponding outcome values exhibit any particular pattern.

This visualization also enables identification of Simonsohn et al. termed *idiosyncratic specifications* (ConnectCombo ▷ Idiosyncratic), e.g., pointing out that only a particular, small subset of the available parameter values lead to negative effect sizes. We discuss such interpretations of outcome frequencies in more depth in subsection 7.1.

6.2.4 Variants of the Descriptive Specification Curve

Figure 8 shows notable variants featuring interesting adaptations and improvements. In Figure 8a, statistical significance is color-coded on both the outcome curve and the universe specification panel (red is significant), and standard error is shown using an error band around the outcome values. This places more visual emphasis on statistical significance and confidence within each universe. Figure 8b maps significance to color but uses a three-color scheme that also indicates the sign of the effect.

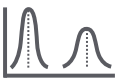
Figure 8c also uses an error band and a different three-color scheme for statistical significance (blue: $\alpha = 0.05$, red: $\alpha = 0.10$, and black: non-significant). Note also that the columns in this variant are the result of a depth-first sorting across the parameter values. This makes some tasks in Connect Combinations (subsection 5.4) easier compared to Figure 6 (so long as the desired combinations of parameter values are clustered together) while making tasks in Connect (subsection 5.3) more difficult (by disrupting the sorting within single parameter values).

Figure 8d presents a multiverse of meta-analyses, where each universe is one meta-analysis. The number of studies within each universe is color-coded (red = 2, blue = 18), and plotted as a frequency plot as an additional middle panel. More generally, this is mapping an additional outcome variable onto color in the specification curve. This allows a task specific to multiverse meta-analysis (ConnectCombo ▷ Idiosyncratic): reasoning about the validity of individual universes based on the number of studies included in their meta-analyses.

Figure 8e is a variant of the outcome curve component (stand alone). It uses confidence intervals around a bootstrapped null distribution instead of around the outcome value, but is otherwise similar to other variants that use error bands.

While not strictly variants of this archetypal family, the standard forest plot, e.g., Arslan's Figure 4 [ASGP18], and dot-interval plot, e.g., Silberzahn's Figure 2 [SUM*18b] could be considered as ancestors of the outcome curve, and have some similar visual features and functionality, though to show only a very small number of universes. See the supplemental material for more detail, including a number of other examples of this archetype [OP19a, OP19b, ODP19, BRRYD20, DGGS20, RES17, VKT19].

6.3 Outcome Density Plot



The *outcome density plot* shows the distribution of outcome values as a density plot. In Figure 9, the outcomes of a multiverse analysis examining potential racial and gender bias in a mortgage-lending dataset are shown. The parameters in this universe indicate whether a specific variable (such as a mortgage applicant's race, marital status) was included as a covariate in a statistical model. The x-axis encodes outcome values of estimated effect size, while the y-axis encodes the relative proportion of universes with the associated effect size.

While similar in function to the outcome histogram, this archetype splits the multiverse into two distribution lines (blue and red) corresponding to two different subsets of the multiverse defined by chosen parameter values. This allows it to support additional task categories, Connect (subsection 5.3) and Connect Combination (subsection 5.4), by isolating subsets of the parameter space of interest. This also means that these tasks are only supported for the particular parameter value(s) or subsets that are directly encoded. While one can easily imagine plotting more than two curves in one plot, it can quickly become cluttered. See the supplemental material for more examples of this archetype [LKJA*19, HS06, ODT12, You18, MY18].

The limited scalability of this archetype in terms of the number of parameters that can be supported is emblematic of an important tradeoff in multiverse visualization design: some visualizations are

© 2022 The Author(s)
Computer Graphics Forum © 2022 The Eurographics Association and John Wiley & Sons Ltd.

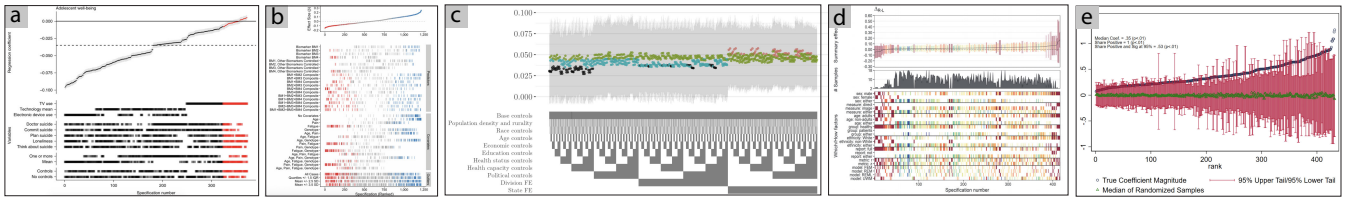


Figure 8: *Example variants of the specification curve archetype, notable for their alternative mappings of the color channel and integration of uncertainty quantification metrics. (a) Figure 1 from Orben et al. [OP19a], (b) Figure 5 from Del Giudice et al. [DGGS20], (c) Figure 7 from Burstyn et al. [BRRYD20], (d) Figure 2 from Voracek et al. [VKT19], (e) Figure 5 from Jelveh et al. [JKN18]*

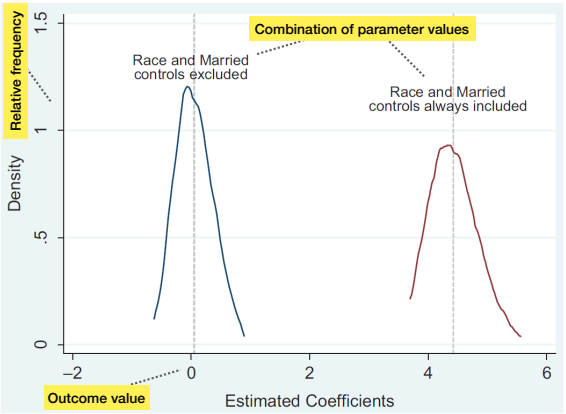


Figure 9: *Example of an outcome density plot, from Young et al. [YH17]. Here, each density curve represents the relative frequency of outcome values across a subset of universes, defined by combinations of parameter values.*

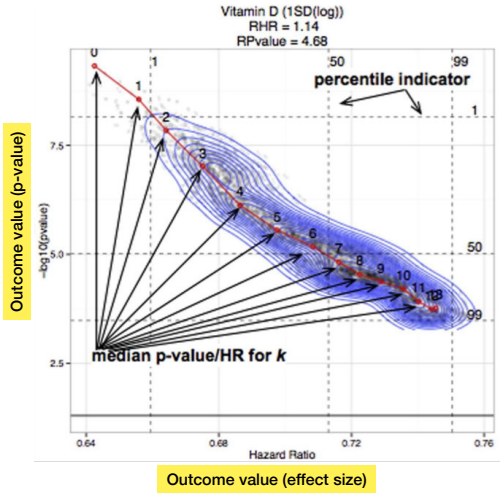


Figure 10: *Example of a vibration of effects plot [PBI15]. The x-axis encodes outcome values (effect size estimates), and the y-axis encodes the statistical significance (negative log transform of p-value). Blue contour lines are used to show the relative frequency of outcomes within the multiverse.*

better for identifying the source of sensitivity in a multiverse overall, while visualizations like the outcome density plot can effectively show the sensitivity of a small selection of parameters after having identified them by other means.

Multi-modality in a density curve of outcome values may indicate that a small subset of parameter values, or combination of parameters, are especially important as they are uniquely responsible for widely different outcome values. As an example of task ConnectCombo ▷ Idiosyncratic, Young et al. [YH17] identify variables for race and marital status as being especially important in their study, and use Figure 9 to illustrate the effect of these decisions on outcome sensitivity. The distribution is multi-modal: all outcome values are close to zero (left curve) or they span large positive effect sizes (right curve). The importance of modality of the outcome distribution is also emphasized in vibration of effects plots (subsection 6.4).

6.4 Vibration of Effects Plot



Figure 10 depicts a multiverse analysis concerned with the reliability of hazard ratios (an effect size) associated with various health factors, like blood levels of vitamin D. The parameters in this analysis are thirteen covariates that can individually be included or excluded, resulting in 8,192 universes. In this *vibration of effects plot* (also called a *volcano plot* by Patel et al. [PBI15]), the effect size is plot-

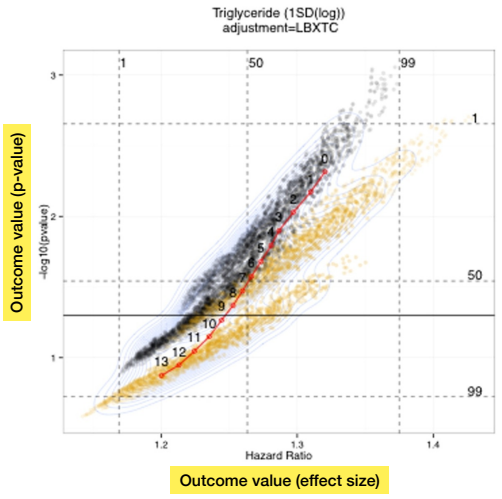


Figure 11: *Variant of the vibration of effects plot [PBI15] where parameter values for a given parameter (inclusion vs. exclusion) are color-coded.*

© 2022 The Author(s)
Computer Graphics Forum © 2022 The Eurographics Association and John Wiley & Sons Ltd.

ted on the x-axis and statistical significance is plotted on the y-axis of a scatter plot with density contour lines. Some other variants in Patel et al. [PBI15] use 2D binned heatmaps instead of scatterplots.

All tasks in Outcome (subsection 5.2) are well-supported by this plot to the extent that density contours and overplotted scatterplots support frequency estimation. All tasks in Connect (subsection 5.3) are supported with comparable ease, and in much the same way, as the Outcome Density plot (subsection 6.3). Similar caveats apply: generally only a small set of combinations of parameter values can be compared at once, e.g., by mapping parameters to colors (Figure 11). However, the 2D density of statistical significance and effect size may allow additional clusters of outcomes to be visible that would not be visible in a 1D density chart, potentially aiding identification of interesting clusters of parameter values.

The identification of potentially important clusters in outcome values is an example of the task ConnectCombo ▷ Idiosyncratic. Patel et al. [PBI15] dedicate extensive discussion of visual patterns exhibited by vibration of effects plots and their interpretation. For example, while the color coding of parameter values in Figure 11 shows this parameter is part of the cause of multimodality in outcomes, there are still at least two visually distinct regions within the outcomes associated with this parameter. This suggests this parameter is not the only cause of multimodality, and that there may be an interaction with another parameter. This ability to identify interaction effects is a unique feature of this archetype, though identifying what specific parameters are responsible (ConnectCombo ▷ OutcomeFrequency or ConnectCombo ▷ OutcomeRange) requires creating additional charts—Patel et al. [PBI15] describe how hundreds of such figures are to be generated to this end.

Patel et al. [PBI15] also describe many idiosyncratic visual patterns and corresponding relationships that can be identified with vibration of effects plots. As an example, outcomes may form a U-shape around 0, which indicates that there are universes that show opposite effect sizes, which Patel et al. call the *Janus effect* (after the Roman god with two faces). Other patterns feature when all universes had the same direction of effect, but disagreed only on magnitude or statistical significance of the effect.

Another common feature of vibration of effects plots is the red line with numerically labeled points, where each points is the median outcome value of all universes with the corresponding number of covariates included (a *derived parameter* as defined in subsubsection 6.2.2). This allows identification of patterns concerned with the joint combination of effect size, statistical significance, and the number of covariates used (ConnectCombo ▷ Idiosyncratic). For example, Patel et al. reported finding cases where more adjusting variables were associated with smaller effect sizes, larger effect sizes, and cases where the effect size appeared to have no dependence on this parameter.

Overall, this archetype represents an effective way of getting an overview of the outcome of a multiverse where two outcome metrics are jointly important. The only other example we found of this archetype was in del Guidace et al. [DGGS20], but this was a near-exact reproduction of the style of this archetype that differed primarily in color choice.

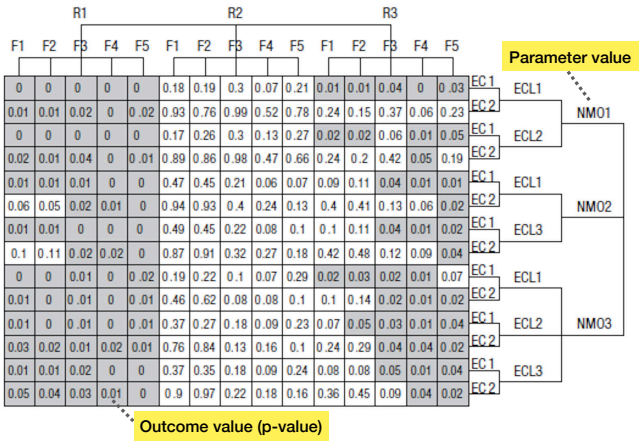
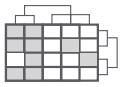


Figure 12: *Example of an outcome matrix [STGV16]. The double-dendrogram structure encodes parameter specification: each level is a parameter, and each node at a given level is a parameter value. Each cell in the matrix thus corresponds to a universe, and indicates the outcome value for this universe (also color-coded).*

6.5 Outcome Matrix



An outcome matrix is a tabular visualization [PDF14] where both rows and columns are parameter values, and each cell reports an outcome value. In Figure 12 each cell reports a *p*-value, both using numerals and a color (statistically significant in gray). In this figure, the axes are dendrograms where each level of the tree is a parameter and each branch a parameter value, thus a path through the tree shows the combinations of parameter values defining each universe. Insofar as the size of the tree is able to scale to the size of multiverse, there is good support for Composition ▷ Parameters in that the structure and relationships within and between parameter values can be derived easily.

Figure 12 is an example of the *outcome matrix* from Steegen et al. [STGV16], a work of the authors who coined the term *multiverse analysis* itself. They chose to visualize their analyses both with this archetype (the *outcome matrix*) and the previously described *outcome histogram* (subsection 6.1). They examined data that explored the relationship between human fertility and religious and political attitudes, across a multiverse defined by data exclusion and operationalization parameters. The outcome of interest is a *p*-value.

The color coding of the outcome values supports Outcome ▷ Frequency (here, the more gray cells the more occurrences of a significant outcome). Outcome ▷ Range for other types of outcomes (e.g., effect size) could be supported given a more granular color coding, although known issues with heatmaps may make certain tasks difficult [KAB*20, GW12].

Tasks in Connect (subsection 5.3) are generally well supported, with a few qualifiers. Tasks Outcome ▷ Range and Outcome ▷ Frequency are relatively easily accomplished when the specified parameter is at the top of the hierarchical axis (e.g., *R1* Figure 12 spans 5 adjacent columns), but require more mental effort otherwise as all the relevant universes are not found within

© 2022 The Author(s)
Computer Graphics Forum © 2022 The Eurographics Association and John Wiley & Sons Ltd.

adjacent columns or row (e.g., *F1* spans 3 non-adjacent columns). The ease of connecting specific outcomes to parameters (Connect ▷ SpecificOutcomes) depends on the hierarchical structure of the parameters as it impacts how outcome values cluster with parameter values: in Figure 12 one can easily observe that all significant *p*-values are in R1 and R3, but if the axes were ordered differently (e.g. swap the order of the *R* and *F* parameters), or if the viewer were interested in a more specific outcome value, the task can become difficult. Similarly, ConnectCombo ▷ OutcomeRange and ConnectCombo ▷ OutcomeFrequency may be well-supported for some combinations of outcome values and axis orderings, making this one of the few visualizations that can support these tasks (at least in some cases). However, the difficulty of all of these tasks depends heavily on row and column ordering and the resulting clusters, as with matrix visualization in general [BBHR*16].

6.5.1 Variants of the Outcome Matrix

Variants of the outcome matrix in our corpus were generally less structurally complex than the example shown in Figure 12, as they omitted the use of a hierarchical axis on either columns or rows. Multiple examples used only one axis to represent parameters, while the other axis was used to show outcomes of interest [CJT19, DKBK19, DS18]. Multiple variants used continuous outcomes and applied different color maps (e.g. diverging palette for positive-negative effect and magnitude), illustrating how this archetype is not fundamentally limited to binary outcomes types [DGH*18, DS18].

6.6 Multiverse Computation Schematic

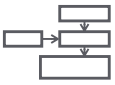


Figure 13, also from Patel et al. [PBI15], is an example of the *multiverse computation schematic* archetype. This is one of the few archetypes whose focus is on Composition (subsection 5.1)—as opposed to reporting outcome values—providing the most support for the tasks Composition ▷ Process and Composition ▷ Parameters in our corpus.

Each panel of Figure 13 denotes a single major stage of the analysis pipeline for creating this multiverse analysis. Panel A describes the data source and Panel B describes the dependent variable in the analysis. Supporting Composition ▷ Parameters, Panel C lists parameters (here, parameter values are either include or exclude) and Panel D describes the statistical model used to produce outcome values for each universe (in some multiverses this would be a parameter if there were more than one model type). Panel E is a miniature vibration of effects plot (subsection 6.4). Panel F contains two metrics the authors use to quantify the spread of outcome values of a multiverse (Outcome ▷ Range), though this is not an essential part of this archetype and the vast majority of multiverse analyses in our corpus do not use such metrics. The illustrated pipeline helps a viewer gain a high level understanding of the multiverse structure (Composition ▷ Process) and the process of analysis.

6.7 Interactive Visualization Systems

While most of the visualizations in our corpus are static, we identified two interactive visualization systems designed to support multiverse analysis. These systems are the primary inspiration for category Validate (subsection 5.5), as these tasks are largely unsupported by the other visualizations in our corpus.

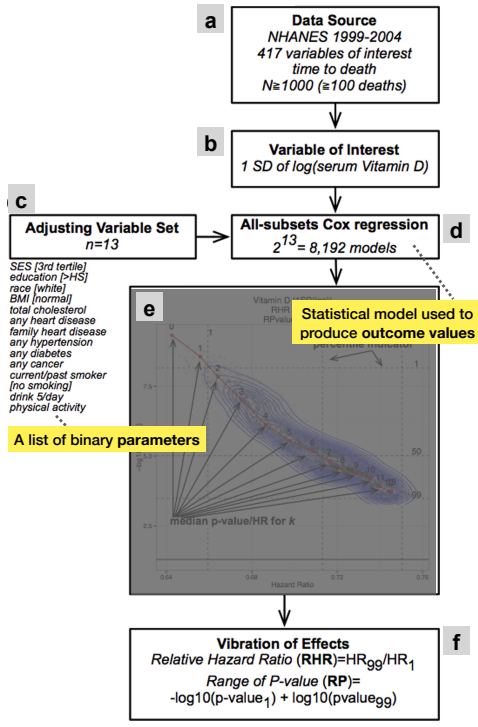


Figure 13: *Example of a multiverse computation schematic [PBI15], describing data source (a), variable of interest (b), and parameters (c,d) composing the multiverse; and elements of the multiverse analysis report: a vibration of effects plot (e); and measures of outcome value spread (f).*

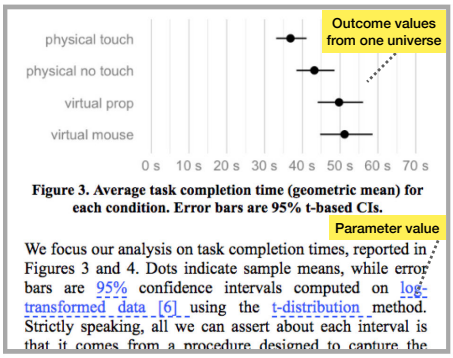
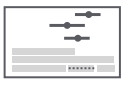


Figure 14: *Excerpt from an explorable multiverse analysis report [DJS*19], where parameter values can be selected dynamically through interactive text widgets, resulting in figures, numerals and text updating accordingly in the report. See* `https://explorablemultiverse.github.io/`.

6.7.1 Explorable Multiverse Analysis Reports (EMAR)



Explorable Multiverse Analysis Reports (EMARs) (Figure 14) are interactive variants of academic ar-

© 2022 The Author(s)
Computer Graphics Forum © 2022 The Eurographics Association and John Wiley & Sons Ltd.

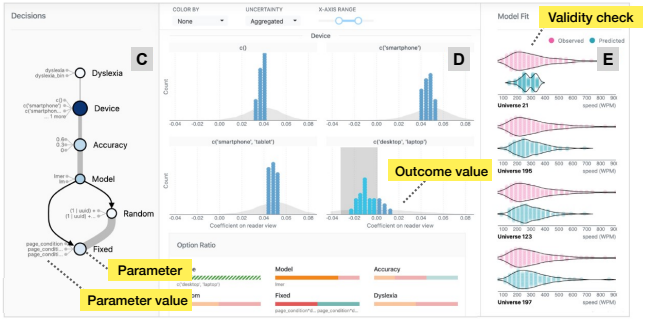


Figure 15: *Screenshot of the Boba system [LKAH20]. Panel C shows the design space of parameters and their relationships; parameters that are source of sensitivity are in a darker color. Panel D is a trellis of dotplots of outcome values, subsetted by parameter values. Panel D shows predictive distributions from each universe compared to the observed data.*

ticles inspired by *explorable explanations* [Vic11]. EMARs allow readers to interactively explore individual universes by selecting combinations of parameter values directly in the report, and see the full analysis report resulting from the corresponding universe update accordingly. For example, the dot-interval plot in Figure 14 is not itself a multiverse visualization; instead, each parameter value in the text is an interactive widget that allows the reader to select different values for that parameter, which updates the body text and all visualizations in the report to describe the analysis resulting from the selected universe. For example, clicking on the *t-distribution* widget allows the reader to switch to bootstrapped confidence intervals.

Unlike the summary visualizations in our corpus, EMARs allow the reader to inspect the full statistical report for a single universe. This allows a reader to make more informed judgments about the validity of each universe (Validate ▷ Details). However, this can make it more difficult to gain a higher-level understanding of outcome sensitivity (subsection 5.2). EMARs address this by allowing the reader to animate over all of the universes to see how much individual visualizations of outcomes change depending on the active universe (Outcome ▷ Frequency).

6.7.2 Boba



Boba (Figure 15) is an interactive system designed to support multiverse analysis. As a full system it supports many tasks in our taxonomy, but the support for some tasks are limited.. It supports tasks in Connect (subsection 5.3) by allowing viewers to interactively select parameters of interest (Figure 15c), which it uses to show dotplots of outcome values faceted by parameter values (Figure 15d) . It has some support for Connect Combination (subsection 5.4) tasks by allowing the viewer to select multiple parameters, though the scalability of these tasks is limited by the fact that faceting is itself limited to two axes. It does not support Validate ▷ Details as it mainly relies on summary visualizations.

A unique contribution of this system is that it explicitly considers model fit (Figure 15e) as a component of assessing multiverse validity (Validate ▷ Metrics). This is because a cross-product of *a priori* reasonable parameters may produce many universes with poor model quality, and some universes may not provide a sound basis for inference [DGGS20]. Support for this task is provided by allowing the viewer to examine model fit (Figure 15e) and exclude outcome values from poor-fitting models in the final interpretation.

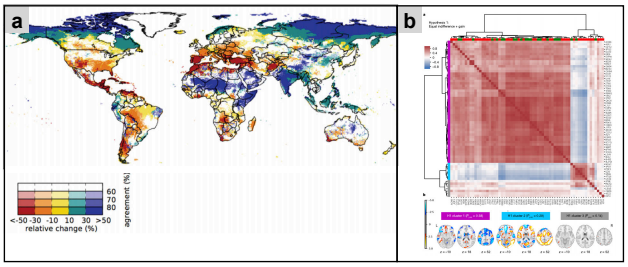


Figure 16: *Two examples of domain-specific visualizations of multiverse analyses. (a) outcome values are contextualized in a geographical map [Bie15], (b) correlation matrix of outcome values [BNHC*20].*

6.8 Domain-Specific Visualizations

We selected the archetypes above for full description as we believe they are likely to be widely applicable to multiverse analyses, regardless of domain. Some of the visualizations in our corpus are instead highly domain-specific [Car12, BZ08, PV17, BKB*20]. A common example is spatial data, such as encountered in geographic and medical research. We present two examples of this type of visualization that both employ heatmaps to encode multiverse outcome data together with domain-specific visualizations that would otherwise only show the result of a single analysis.

Figure 16a shows the output of water runoff (discharge) predictions from 55 climate models [Bie15]. Outcome values and sensitivity are encoded on a bivariate color scale: mean predicted change in water runoff (outcome value) is mapped to hue, and percentage agreement between universes (outcome sensitivity) is mapped to saturation, helping the viewer assess the range of outcomes in each region on the map (Outcome ▷ Range).

Figure 16b shows the correlation between outcomes across universes in a neuroimaging analysis multiverse. The top panel is a correlation matrix: rows and columns are universes, and each cell shows the correlation of outcome values between two universes. The dendrogram axes are similar to those of the outcome matrix (subsection 6.5), but are the result of a clustering algorithm rather than a direct representation of parameters. The color-coding on the rows links the results in the matrix to the models of human brains in the lower panel. Each brain model uses a heatmap to show the averaged relative activation of certain brain areas. This aids in assessing the sensitivity of outcomes to different analysis choices (Connect, subsection 5.3).

7 Discussion

In this section, we discuss some difficulties and limitations of current multiverse analysis visualizations, implications for design following from our survey, and directions for future work.

© 2022 The Author(s)
Computer Graphics Forum © 2022 The Eurographics Association and John Wiley & Sons Ltd.

7.1 The Illusion of Probability in Multiverse Visualizations

One issue with existing multiverse visualizations that show outcome values stems from the subtle yet important distinction between *probabilistic* and *possibilistic* interpretations of frequencies. Although this is a general difficulty when interpreting any multiverse analysis, it may be exacerbated by visualizations.

Under a *probabilistic* interpretation, all specified universes would be assumed to be equally likely to be correct, so outcome values that occur more frequently within the multiverse must be more likely to be correct. Yet the set of reported universes in a multiverse analysis is not a random sample of all reasonable specifications, and universes are themselves not statistically independent [SSN20]. Authors and readers may even disagree on the validity of some universes [SSN20]. It follows that when interpreting visualizations such as outcome histograms, the relative frequencies of outcomes should *not* be treated probabilistically.

Instead, variation in outcomes should be treated *possibilistically* [GK75]: the presence of an outcome in a multiverse indicates that it is a possible result of reasonable analytical choices. Under this interpretation, no outcome value can be considered any more or less likely to be correct solely based on how frequently it occurs; one must instead examine the validity of universes leading to particular outcomes.

However, the probabilistic interpretation is very tempting: we suspect many readers may interpret visualized outcome frequencies as probabilities or likelihoods, and we have encountered such interpretations while reading multiverse analysis reports. This relates to classic notions of visualization *expressiveness*: density plots, histograms, dotplots, and so forth. all invite a probabilistic interpretation even though that interpretation is not intended for multiverse data. One might consider this misinterpretation a kind of **illusion of probability**. This illusion puts designers of visualizations for multiverse analysis in a bind, as the frequency information that creates the illusion is still useful for many tasks (e.g., Connect ▷ OutcomeFrequency and ConnectCombo ▷ OutcomeFrequency). How can we visualize this frequency information while preventing erroneous probabilistic interpretations? One potential direction may be to use visualization types explicitly designed for possibilistic uncertainty, such as probability boxes [FS11]; see Bonneau et al. [BHJ*14] for further discussion of possibilistic versus probabilistic uncertainty visualization. We have not seen examples of possibilistic uncertainty visualizations applied to multiverse analysis as yet.

7.2 Visualizations to Better Support Multiverse Validation and Interpretation are Needed

We considered proposing a sixth task category, "Interpret the Multiverse" as the logical final step in a multiverse analysis: to make some inference about the original dataset (not about the sensitivity of that inference). We decided against doing so as we did not find examples of tasks in this category that were substantially supported by multiple sources in the corpus, generalizable, and explicitly a feature of a visualization. Overall, we found that interpretations of a multiverse vary widely between authors, are often domain-dependent, and are not strongly tied to specific features of any visualization.

Del Giudice et al. [DGGS20] stated that, *"Going forward, multiverse-style methods should not be narrowly thought of as a means to promote transparency in reporting, but rather as an analytic tool that can profitably aid the interpretation of data and inform the development of theoretical models."* This echoes similar suggestions made in earlier works [SSN20, STGV16], but most multiverse reports we reviewed did not go beyond tasks from the Outcome (subsection 5.2) and Connect (subsection 5.3) categories, or at least not in a way that explicitly referenced a visualization.

Only 2 visualizations provided support for tasks under the Validate category (subsection 5.5). Two recent threads of research have suggested the need to more carefully validate the universes in a multiverse, possibly pruning some universes. Liu et al. [LKAH20] suggest doing so by examining model fit and provide some support for this task in Boba (Figure 15). They suggest an analyst might wish to iteratively redefine the multiverse itself as a result of a previous round of multiverse analysis, given that some analytical choices may no longer be considered equally defensible after having run them on the data.

Relatedly, Del Giudice et al. [DGGS20] argue that analysts should explicitly consider whether analytical choices have *principled equivalence*, *principled non-equivalence*, or if there is *uncertainty* about their equivalence; each conclusion leads to different choices about whether to include a parameter value in the multiverse. They argue that if poor analysis choices were truly excluded, most multiverses would be much smaller than ones seen in practice. Simonsohn et al. [SSN20] note that, *"While all included specifications should be theoretically justified, statistically valid and non-redundant, researchers may nevertheless consider some specifications superior to others and that some should be given greater weight than others."* However, to date we are not aware of reported multiverse analyses that attempt such relative weightings.

7.3 Multiplexing and Interaction to Investigate Parameter Combinations

Few visualizations provided substantial support for tasks in the Connect Combinations category (subsection 5.4). For most visualizations , this support comes with the caveat that meaningful combinations of parameters have been selected ahead of time (e.g. Figure 9), which does not address how visualization might be used to discover these interesting relationships in the first place. Two strategies in the corpus were used to help analysts discover the impact of arbitrary parameter combinations on outcome sensitivity: multiplexing in space (e.g., faceting), and interactivity.

While the vibrations of effects plot (subsection 6.4) can only compare across a small number of parameter values at once, Patel et al. [PBI15] describe a full analysis workflow in which an analyst reviews potentially hundreds of vibration of effects plots representing combinations of parameter values. On a smaller scale, Poarch et al. [PVB19] faceted by both variables of interest and parameters, producing an 8-by-6 of outcome histograms (subsection 6.1) to report their multiverse analysis. In theory, faceting by parameter can be performed with any base-plot type, but in our corpus faceting

© 2022 The Author(s)
Computer Graphics Forum © 2022 The Eurographics Association and John Wiley & Sons Ltd.

was primarily used with archetypes that did not otherwise support connecting parameters to outcome values (subsection 5.3 and subsection 5.4).

Boba (Figure 15) combined faceting with interactivity, allowing viewers to facet according to interactively-selected parameters. Interactivity removes the need to present all faceted plots at once and could aid in more focused exploration. However, there is an untapped potential to enhance the value of other plot types in our corpus through interactivity, beyond just interactively selecting facets: the outcome matrix plot (Figure 12), for example, could benefit from interactive row and column reordering to aid in cluster identification [PDF14]; similar functionality could also help reduce tradeoffs in fixed column ordering on specification curve charts (e.g. Figure 8a versus Figure 8b). Such approaches could be used in interactive systems aimed at analysts, like Boba [LKAH20], or incorporated into interactive reports aimed at readers, like EMARs [DJS*19].

7.4 Importance of Multiverse Scale and Structure

Multiverses vary in their scale, in terms of both the number of parameters and the number of universes those parameters form in combination. Some multiverses are *dense*, if most or all combinations of parameter values are included, while some are *sparse*, if many theoretically possible combinations of parameters are not included. Sparse multiverses are typical in analyses constructed by using only the specifications found in previous work, or when specifications are crowdsourced (e.g., [SUM*18b]). Some archetypes explicitly visualize this structure (e.g. the dendrograms in outcome matrices; Figure 12) and may not scale well to large numbers of parameters or complex relationships between them, while others do not depict any particular structure and are thus usable regardless (e.g. the outcome histogram; Figure 5).

Part of the inherent difficulty of multiverse analysis is that the data is not easily reduced or summarized without losing information that is critical for supporting important tasks, such as Connect ▷ OutcomeRange or ConnectCombo ▷ OutcomeRange. Summarization of outcome values can appear trivial at first, such as when stating the proportion of universes with outcomes values that were statistically significant, or presenting outcomes with a histogram (subsection 6.1). As discussed previously (subsection 6.1 & subsection 7.1), under a possibilistic interpretation even this task is fraught with the danger of misinterpretation. While frequency can also serve as an indicator for how much of the examined choice space is connected to any given outcome, summarizing outcomes severs the threads that connect outcomes to parameter values, thus preventing one from performing any Connect-related tasks (subsection 5.3). It may be that supporting some tasks better will tend to reduce support for other tasks. This implies that designers and researchers may be best served by building up a toolbox of multiverse visualizations that support their desired tasks, rather than trying in vain to create an all-in-one solution.

Given this, the design of visualizations must take into account the scale of the multiverses they are to support. In the visualization table in the supplement we provide our estimation of the scale of multiverses that are supported by each archetype, both in terms of number of parameters and number of universes. As an example, the vibration of effects plot (subsection 6.4) scales to an unlimited number of universes, but is only able to show one (or very few) parameter values in a single plot. By contrast, an interactive system is not limited in the amount of parameters it can support overall, but the component visualizations are still limited to simultaneously displaying a number of parameters on the order of tens. Future work might investigate ways to scale multiverse visualizations that already have good support for some tasks to larger multiverses.

7.5 Limitations of this Survey and Future Work

There are several ways in which our survey is limited. We set out to survey tasks and visualizations for multiverse analysis reports, as detailed in section 3. Since adjacent concepts, such as model comparison, or parameter space exploration (also see Figure 2) likely entail different tasks, we curated our corpus by strictly applying the definitions presented in section 3. The eight relevant keywords identified from our list of 53 seed articles resulted in a total of 213 corpus candidates. In analyzing these candidates we only found a total of 43 articles fulfilling our criteria. Consequently, our survey may have missed some potentially relevant visualizations.

Our survey only covers multiverse visualizations reported in academic papers, most of which are static. We had to exclude many visualization designs and tools—some of which are interactive—that have been designed for related purposes (see Figure 2). Future work should examine how such tools can inspire the design of multiverse analysis reports, while remaining aware of differences in goals. For example, interactive visualization tools for model building [DCCE19, MLMP17, CPCS19] and for ensemble data analysis [WHLS18, SHB*14] focus on using data visualization to help analysts prune vast spaces of possibilities, often with the goal of identifying one optimal model or set of parameters. In contrast, in a typical multiverse analysis the entire multiverse is reported as it was decided before the data was analyzed, irrespective of the outcomes of those analyses. Nevertheless, pruning tools require effective data overview techniques, which can be re-purposed for multiverse analysis reporting. In addition, adding interactive pruning tools to multiverse analysis reports could help readers navigate them.

Our survey covers how multiverse visualizations have been used across disciplines, but few of the papers we examined are from within the field of information visualization. This is because such visualizations are not broadly used, and we know of very few examples in information visualization. Our focus is however less on helping information visualization researchers *use* such visualizations in their own papers, and more on helping them *study* them as a research subject. We however expect that many of the insights gained by looking at practices across disciplines can transfer to visualization papers, as methodologies for analyzing and reporting experiments and transparency criteria are very similar across research areas.

None of the tasks discussed in this work are unique to any single domain or discipline, and the vast majority of datasets being analyzed are well expressed in tabular data structures familiar to all quantitative analysts. Major challenges to be addressed by future researchers will involve finding ways to effectively communi-

© 2022 The Author(s)
Computer Graphics Forum © 2022 The Eurographics Association and John Wiley & Sons Ltd.

cate multiverse results of data and analyses with additional structural complexity. For example, hierarchical data and modeling techniques can require multiple visualizations to adequately communicate the results of a single analysis. Similarly, there is no reason why multiverse analysis techniques cannot be applied to analyses of other data structures, such as networks. While domain-specific techniques applied to spatial data may provide some inspiration (subsection 6.8), considerable innovation may be required.

8 Conclusion

This state of the art report has reviewed the development and advances made in the visual design and communication of multiverse analysis results, starting with related techniques that go back long before the term *multiverse analysis* was first coined, and carried through the year 2020. We surveyed literature across multiple fields and disciplines, considering visualizations from areas as diverse as psychology, statistics, economics, and visualization.

We contributed a coherent and operational terminology to provide researchers with a common vocabulary so they can better communicate and reason about multiverse analyses (section 3). We assembled a taxonomy of analysis inspection tasks that multiverse visualizations should support, grounded in an extensive analysis of the curated corpus (section 5). Finally, we discussed the design and functionality of major multiverse visualization archetypes and assessed how well each of them supports our tasks (section 6), in order to guide analysts in the selection of appropriate visualizations to use when conducting or reporting multiverse analyses.

Our work was motivated by the fact that visualization solutions to multiverse analysis and reporting have, to date, been largely explored in isolation. We contribute a conceptual framework and reflections that can help shed light on this rich design space. Ultimately, no single multiverse visualization has dominant support for all tasks, and there is ample opportunity for future work to investigate improvements to existing visualizations, new visualizations, or even combinations of visualizations to better support the range of tasks needed for a complete reporting of a multiverse analysis.

9 Acknowledgements

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE 1256260.

References

[AES05] AMAR R., EAGAN J., STASKO J.: Low-level components of analytic activity in information visualization. In *IEEE Symposium on Information Visualization, 2005. INFOVIS 2005.* (2005), IEEE, pp. 111–117. 8

[ASGP18] ARSLAN R. C., SCHILLING K. M., GERLACH T. M., PENKE L.: Using 26,000 diary entries to show ovulatory changes in sexual desire and behavior. *Journal of Personality and Social Psychology* (2018). 6, 13

[BBHR*16] BEHRISCH M., BACH B., HENRY RICHE N., SCHRECK T., FEKETE J.-D.: Matrix reordering methods for table and network visualization. In *Computer Graphics Forum* (2016), vol. 35, Wiley Online Library, pp. 693–716. 16

[Ber12] BERRY D.: Multiplicities in cancer research: ubiquitous and necessary evils. *Journal of the National Cancer Institute 104*, 15 (2012), 1125–1133. 2

[BHJ*14] BONNEAU G.-P., HEGE H.-C., JOHNSON C. R., OLIVEIRA M. M., POTTER K., RHEINGANS P., SCHULTZ T.: Overview and state-of-the-art of uncertainty visualization. In *Scientific Visualization*. Springer, 2014, pp. 3–27. 18

[BI16] BRUNS S. B., IOANNIDIS J. P.: P-curve and p-hacking in observational research. *PloS one 11*, 2 (2016), e0149144. 11

[Bie15] BIERKENS M. F. P.: Global hydrology 2015: State, trends, and directions. *Water Resources Research 51*, 7 (2015), 4923–4947. 2, 17

[BKB*20] BASTIAANSEN J. A., KUNKELS Y. K., BLAAUW F. J., BOKER S. M., CEULEMANS E., CHEN M., CHOW S.-M., DE JONGE P., EMERENCIA A. C., EPSKAMP S., ET AL.: Time to get personal? the impact of researchers choices on the selection of treatment targets using the experience sampling methodology. *Journal of psychosomatic research 137* (2020), 110211. 2, 4, 6, 17

[BNHC*20] BOTVINIK-NEZER R., HOLZMEISTER F., CAMERER C. F., DREBER A., HUBER J., JOHANNESSON M., KIRCHLER M., IWANIR R., MUMFORD J. A., ADCOCK R. A., ET AL.: Variability in the analysis of a single neuroimaging dataset by many teams. *Nature 582* (2020), 84–88. 2, 4, 6, 17

[BRRYD20] BURSZTYN L., RAO A., ROTH C., YANAGIZAWA-DROTT D.: Misinformation during a pandemic. *University of Chicago, Becker Friedman Institute for Economics Working Paper*, 2020-44 (2020), 1–114. 2, 13, 14

[BYO19] BRYAN C. J., YEAGER D. S., O'BRIEN J. M.: Replicator degrees of freedom allow publication of misleading failures to replicate. *Proceedings of the National Academy of Sciences 116*, 51 (2019), 25535–25545. 12

[BZ08] BARALDI P., ZIO E.: A combined monte carlo and possibilistic approach to uncertainty propagation in event tree analysis. *Risk Analysis: An International Journal 28*, 5 (2008), 1309–1326. 17

[Car12] CARP J.: On the plurality of (methodological) worlds: estimating the analytic flexibility of fmri experiments. *Frontiers in neuroscience 6* (2012), 149. 2, 6, 17

[CC14] CHRISTENSEN B., CHRISTENSEN S.: Are female hurricanes really deadlier than male hurricanes? *Proceedings of the National Academy of Sciences 111*, 34 (2014), E3497–E3498. 2

[CGD18] COCKBURN A., GUTWIN C., DIX A.: Hark no more: on the preregistration of CHI experiments. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (2018), ACM, p. 141. 2

[CJT19] CESARIO J., JOHNSON D. J., TERRILL W.: Is there evidence of racial disparity in police use of deadly force? analyses of officer-involved fatal shootings in 2015–2016. *Social psychological and personality science 10*, 5 (2019), 586–595. 16

[Coo18] COOKSON J. A.: When saving is gambling. *Journal of Financial Economics 129*, 1 (2018), 24–45. 12

[CPCS19] CASHMAN D., PERER A., CHANG R., STROBELT H.: Ablate, variate, and contemplate: Visual analytics for discovering neural architectures. *IEEE transactions on visualization and computer graphics 26*, 1 (2019), 863–873. 2, 3, 19

[CT16] CIRILLO P., TALEB N. N.: On the statistical properties and tail risk of violent conflicts. *Physica A: Statistical Mechanics and its Applications 452* (2016), 29–45. 11

[Cum14] CUMMING G.: The new statistics: Why and how. *Psychological science 25*, 1 (2014), 7–29. 2

[DBH19] DONNELLY S., BROOKS P. J., HOMER B. D.: Is there a bilingual advantage on interference-control tasks? a multiverse meta-analysis of global reaction time and interference cost. *Psychonomic bulletin & review 26*, 4 (2019), 1122–1147. 3

© 2022 The Author(s)
Computer Graphics Forum © 2022 The Eurographics Association and John Wiley & Sons Ltd.

[DCCE19] DAS S., CASHMAN D., CHANG R., ENDERT A.: Beames: Interactive multimodel steering, selection, and inspection for regression tasks. *IEEE computer graphics and applications 39*, 5 (2019), 20–32. 2, 3, 19

[DGGS20] DEL GIUDICE M., GANGESTAD S. W., STEVEN W.: A traveler's guide to the multiverse: Promises, pitfalls, and a framework for the evaluation of analytic decisions, 2020. 5, 11, 13, 14, 15, 17, 18

[DGH*18] DUBOIS J., GALDI P., HAN Y., PAUL L. K., ADOLPHS R.: Resting-state functional brain connectivity best predicts the personality dimension of openness to experience. *Personality neuroscience 1* (2018). 16

[DJS*19] DRAGICEVIC P., JANSEN Y., SARMA A., KAY M., CHEVALIER F.: Increasing the transparency of research papers with explorable multiverse analyses. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (New York City, NY, USA, 2019), Association for Computing Machinery, pp. 1–15. 1, 2, 3, 4, 5, 10, 16, 19

[DKBK19] DEJONCKHEERE E., KALOKERINOS E. K., BASTIAN B., KUPPENS P.: Poor emotion regulation ability mediates the link between depressive symptoms and affective bipolarity. *Cognition and Emotion 33*, 5 (2019), 1076–1083. 16

[DMH*18] DEJONCKHEERE E., MESTDAGH M., HOUBEN M., ERBAS Y., PE M., KOVAL P., BROSE A., BASTIAN B., KUPPENS P.: The bipolarity of affect and depressive symptoms. *Journal of personality and social psychology 114*, 2 (2018), 323. 11

[DS18] DENNY M. J., SPIRLING A.: Text preprocessing for unsupervised learning: Why it matters, when it misleads, and what to do about it. *Political Analysis 26*, 2 (2018), 168–189. doi:10.1017/pan.2017.44. 12, 16

[ESR17] ELSHERIF M. M., SABAN M. I., ROTSHTEIN P.: The perceptual saliency of fearful eyes and smiles: A signal detection study. *PloS one 12*, 3 (2017), e0173199. 3

[Fiv15] FIVETHIRTYEIGHT: Hack your way to scientific glory, 2015. 4

[FS11] FERSON S., SIEGRIST J.: Verified computation with probabilities. In *IFIP Working Conference on Uncertainty Quantification* (New York City, NY, USA, 2011), Springer, pp. 95–122. 18

[GHF14] GILDERSLEEVE K., HASELTON M. G., FALES M. R.: Do women's mate preferences change across the ovulatory cycle? a meta-analytic review. *Psychological Bulletin 140*, 5 (2014), 1205. 3, 13

[GK75] GAINES B. R., KOHOUT T. L.: Possible automata. *Proceedings of the International Symposium on Multiple-Valued Logic* (1975). 18

[GL13] GELMAN A., LOKEN E.: The garden of forking paths: Why multiple comparisons can be a problem, even when there is no "fishing expedition" or "p-hacking" and the research hypothesis was posited ahead of time. *Department of Statistics, Columbia University* (2013). 2

[GW12] GEHLENBORG N., WONG B.: Points of view: Heat maps. *Nature Methods 9*, 3 (2012), 213. 15

[Har07] HARZING A.: Publish or perish, 2007. 5

[HCM13] HARRIS C. R., CHABOT A., MICKES L.: Shifts in methodology and theory in menstrual cycle research on attraction. *Sex Roles 69*, 9-10 (2013), 525–535. 13

[HS06] HEGRE H., SAMBANIS N.: Sensitivity analysis of empirical results on civil war onset. *Journal of conflict resolution 50*, 4 (2006), 508–535. 13

[JKN18] JELVEH Z., KOGUT B., NAIDU S.: Political language in economics. *Columbia Business School Research Paper*, 14-57 (2018), 1–67. 12, 14

[JSVH14] JUNG K., SHAVITT S., VISWANATHAN M., HILBE J. M.: Female hurricanes are deadlier than male hurricanes. *Proceedings of the National Academy of Sciences 111*, 24 (2014), 8782–8787. 2

[KAB*20] KRAUS M., ANGERBAUER K., BUCHMÜLLER J., SCHWEITZER D., KEIM D. A., SEDLMAIR M., FUCHS J.: Assessing 2d and 3d heatmaps for comparative analysis: An empirical study. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (2020), pp. 1–14. 15

[LAH20] LIU Y., ALTHOFF T., HEER J.: Paths explored, paths omitted, paths obscured: Decision points & selective reporting in end-to-end data analysis. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (2020), pp. 1–14. 1

[LJD*20] LANDY J. F., JIA M. L., DING I. L., VIGANOLA D., TIERNEY W., DREBER A., JOHANNESSON M., PFEIFFER T., EBERSOLE C. R., GRONAU Q. F., ET AL.: Crowdsourcing hypothesis tests: Making transparent how design choices shape research results. *Psychological Bulletin 146*, 5 (2020), 451–479. 3

[LKAH20] LIU Y., KALE A., ALTHOFF T., HEER J.: Boba: Authoring and visualizing multiverse analyses, 2020. arXiv:2007.05551. 1, 2, 5, 7, 10, 17, 18, 19

[LKJA*19] LONSDORF T. B., KLINGELHÖFER-JENS M., ANDREATTA M., BECKERS T., CHALKIA A., GERLICHER A., JENTSCH V. L., DREXLER S. M., MERTENS G., RICHTER J., ET AL.: Navigating the garden of forking paths for data exclusions in fear conditioning research. *Elife 8* (2019), e52465. 6, 13

[Mal14a] MALEY S.: Statistics show no evidence of gender bias in the public's hurricane preparedness. *Proceedings of the National Academy of Sciences 111*, 37 (2014), E3834–E3834. 2

[Mal14b] MALTER D.: Female hurricanes are not deadlier than male hurricanes. *Proceedings of the National Academy of Sciences 111*, 34 (2014), E3496–E3496. 2

[MLMP17] MÜHLBACHER T., LINHARDT L., MÖLLER T., PIRINGER H.: Treepod: Sensitivity-aware selection of pareto-optimal decision trees. *IEEE transactions on visualization and computer graphics 24*, 1 (2017), 174–183. 2, 3, 4, 19

[MNB*17] MUNAFÒ M. R., NOSEK B. A., BISHOP D. V., BUTTON K. S., CHAMBERS C. D., DU SERT N. P., SIMONSOHN U., WAGENMAKERS E.-J., WARE J. J., IOANNIDIS J. P.: A manifesto for reproducible science. *Nature Human Behaviour 1*, 1 (2017), 0021. 2

[Mun14] MUNZNER T.: *Visualization analysis and design*. CRC press, 2014. 4

[MY18] MUÑOZ J., YOUNG C.: We ran 9 billion regressions: Eliminating false positives through computational model robustness. *Sociological Methodology 48*, 1 (2018), 1–33. 13

[NAB*15] NOSEK B. A., ALTER G., BANKS G. C., BORSBOOM D., BOWMAN S. D., BRECKLER S. J., BUCK S., CHAMBERS C. D., CHIN G., CHRISTENSEN G., ET AL.: Promoting an open research culture. *Science 348*, 6242 (2015), 1422–1425. 2

[ODP19] ORBEN A., DIENLIN T., PRZYBYLSKI A. K.: Social media's enduring effect on adolescent life satisfaction. *Proceedings of the National Academy of Sciences 116*, 21 (2019), 10226–10228. 13

[ODT12] OLKIN I., DAHABREH I. J., TRIKALINOS T. A.: Gosh–a graphical display of study heterogeneity. *Research Synthesis Methods 3*, 3 (2012), 214–223. 6, 13

[OP19a] ORBEN A., PRZYBYLSKI A. K.: The association between adolescent well-being and digital technology use. *Nature Human Behaviour 3*, 2 (2019), 173–182. 2, 12, 13, 14

[OP19b] ORBEN A., PRZYBYLSKI A. K.: Screens, teens, and psychological well-being: evidence from three time-use-diary studies. *Psychological science 30*, 5 (2019), 682–696. 13

[Pap64] PAP A.: Theory of definition. *Philosophy of science 31*, 1 (1964), 49–54. 3

[PBI15] PATEL C. J., BURFORD B., IOANNIDIS J. P.: Assessment of vibration of effects due to model specification can demonstrate the instability of observational associations. *Journal of clinical epidemiology 68*, 9 (2015), 1046–1058. 1, 2, 5, 7, 9, 14, 15, 16, 18

[PDF14] PERIN C., DRAGICEVIC P., FEKETE J.-D.: Revisiting bertin matrices: New interactions for crafting tabular visualizations. *IEEE transactions on visualization and computer graphics 20*, 12 (2014), 2082–2091. 4, 12, 15, 19

© 2022 The Author(s)
Computer Graphics Forum © 2022 The Eurographics Association and John Wiley & Sons Ltd.

[PV17] PIIRONEN J., VEHTARI A.: Comparison of bayesian predictive methods for model selection. *Statistics and Computing 27*, 3 (2017), 711–735. 17

[PVB19] POARCH G. J., VANHOVE J., BERTHELE R.: The effect of bidialectalism on executive function. *International Journal of Bilingualism 23*, 2 (2019), 612–628. 6, 11, 18

[RES17] ROHRER J. M., EGLOFF B., SCHMUKLE S. C.: Probing birth-order effects on narrow traits using specification-curve analysis. *Psychological Science 28*, 12 (2017), 1821–1832. 13

[SHB*14] SEDLMAIR M., HEINZL C., BRUCKNER S., PIRINGER H., MÖLLER T.: Visual parameter space analysis: A conceptual framework. *IEEE Transactions on Visualization and Computer Graphics 20*, 12 (2014), 2161–2170. 2, 3, 19

[SNS11] SIMMONS J. P., NELSON L. D., SIMONSOHN U.: False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological science 22*, 11 (2011), 1359–1366. 2

[SR18] SAGI O., ROKACH L.: Ensemble learning: A survey. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery 8*, 4 (2018), e1249. 4

[SSN19] SIMONSOHN U., SIMMONS J. P., NELSON L. D.: Specification curve: Descriptive and inferential statistics on all reasonable specifications, 2019. 2, 3, 5, 6, 8, 9, 10, 12

[SSN20] SIMONSOHN U., SIMMONS J. P., NELSON L. D.: Specification curve analysis. *Nature Human Behaviour 4*, 11 (2020), 1208–1214. 1, 2, 11, 12, 18

[STGV16] STEEGEN S., TUERLINCKX F., GELMAN A., VANPAEMEL W.: Increasing transparency through a multiverse analysis. *Perspectives on Psychological Science 11*, 5 (2016), 702–712. 1, 2, 3, 5, 8, 9, 10, 11, 15, 18

[SUM*18a] SILBERZAHN R., UHLMANN E. L., MARTIN D. P., ANSELMI P., AUST F., AWTREY E., BAHNÍK Š., BAI F., BANNARD C., BONNIER E., CARLSSON R., CHEUNG F., CHRISTENSEN G., CLAY R., CRAIG M. A., DALLA ROSA A., DAM L., EVANS M. H., FLORES CERVANTES I., FONG N., GAMEZ-DJOKIC M., GLENZ A., GORDON-MCKEON S., HEATON T. J., HEDEROS K., HEENE M., HOFELICH MOHR A. J., HÖGDEN F., HUI K., JOHANNESSON M., KALODIMOS J., KASZUBOWSKI E., KENNEDY D. M., LEI R., LINDSAY T. A., LIVERANI S., MADAN C. R., MOLDEN D., MOLLEMAN E., MOREY R. D., MULDER L. B., NIJSTAD B. R., POPE N. G., POPE B., PRENOVEAU J. M., RINK F., ROBUSTO E., RODERIQUE H., SANDBERG A., SCHLÜTER E., SCHÖNBRODT F. D., SHERMAN M. F., SOMMER S. A., SOTAK K., SPAIN S., SPÖRLEIN C., STAFFORD T., STEFANUTTI L., TAUBER S., ULLRICH J., VIANELLO M., WAGENMAKERS E.-J., WITKOWIAK M., YOON S., NOSEK B. A.: Many Analysts, One Data Set: Making Transparent How Variations in Analytic Choices Affect Results. *Advances in Methods and Practices in Psychological Science 1*, 3 (sep 2018), 337–356. URL: http://journals.sagepub.com/doi/10.1177/2515245917747646, doi:10.1177/2515245917747646. 2, 12

[SUM*18b] SILBERZAHN R., UHLMANN E. L., MARTIN D. P., ANSELMI P., AUST F., AWTREY E., BAHNÍK Š., BAI F., BANNARD C., BONNIER E., ET AL.: Many analysts, one data set: Making transparent how variations in analytic choices affect results. *Advances in Methods and Practices in Psychological Science 1*, 3 (2018), 337–356. 5, 8, 9, 13, 19

[Vic11] VICTOR B.: Explorable explanations. Online. http://worrydream.com/ExplorableExplanations/, 2011. 17

[VKT19] VORACEK M., KOSSMEIER M., TRAN U. S.: Which data to meta-analyze, and how? *Zeitschrift für Psychologie 227* (2019), 64–82. 11, 12, 13, 14

[Wah83] WAHBA G.: Bayesian "confidence intervals" for the cross-validated smoothing spline. *Journal of the Royal Statistical Society: Series B (Methodological) 45*, 1 (1983), 133–150. 4

[WHLS18] WANG J., HAZARIKA S., LI C., SHEN H.-W.: Visualization and visual analysis of ensemble data: A survey. *IEEE transactions on visualization and computer graphics 25*, 9 (2018), 2853–2872. 2, 3, 19

[WVA*16a] WICHERTS J. M., VELDKAMP C. L., AUGUSTEIJN H. E., BAKKER M., VAN AERT R., VAN ASSEN M. A.: Degrees of freedom in planning, running, analyzing, and reporting psychological studies: A checklist to avoid p-hacking. *Frontiers in Psychology 7* (2016), 1832. 2

[WVA*16b] WICHERTS J. M., VELDKAMP C. L., AUGUSTEIJN H. E., BAKKER M., VAN AERT R., VAN ASSEN M. A.: Degrees of freedom in planning, running, analyzing, and reporting psychological studies: A checklist to avoid p-hacking. *Frontiers in psychology 7* (2016), 1832. 3

[YH17] YOUNG C., HOLSTEEN K.: Model uncertainty and robustness: A computational framework for multimodel analysis. *Sociological Methods & Research 46*, 1 (2017), 3–40. 1, 5, 9, 14

[You18] YOUNG C.: Model uncertainty and the crisis in science. *Socius 4* (2018). 13

© 2022 The Author(s)
Computer Graphics Forum © 2022 The Eurographics Association and John Wiley & Sons Ltd.