

Visualization Atlases: Explaining and Exploring Complex Topics through Data, Visualization, and Narration

Jinrui Wang , Xinhuan Shu , Benjamin Bach , and Uta Hinrichs 

Abstract—This paper defines, analyzes, and discusses the emerging genre of *visualization atlases*. We currently witness an increase in web-based, data-driven initiatives that call themselves “atlases” while explaining complex, contemporary issues through data and visualizations: climate change, sustainability, AI, or cultural discoveries. To understand this emerging genre and inform their design, study, and authoring support, we conducted a systematic analysis of 33 visualization atlases and semi-structured interviews with eight visualization atlas creators. Based on our results, we contribute (1) a definition of a visualization atlas as a compendium of (web) pages aimed at explaining and supporting exploration of data about a dedicated topic through data, visualizations and narration. (2) a set of design patterns of 8 design dimensions, (3) insights into the atlas creation from interviews and (4) the definition of 5 visualization atlas genres. We found that visualization atlases are unique in the way they combine i) exploratory visualization, ii) narrative elements from data-driven storytelling and iii) structured navigation mechanisms. They target a wide range of audiences with different levels of domain knowledge, acting as tools for study, communication, and discovery. We conclude with a discussion of current design practices and emerging questions around the ethics and potential real-world impact of visualization atlases, aimed to inform the design and study of visualization atlases.

Index Terms—Visualization Atlases, Information Visualization, Data-driven Storytelling

1 INTRODUCTION

Over the past years, we noticed an increase in highly polished web-based projects that use visualizations to make data about complicated contemporary problems accessible to a wide audience. Projects like *The Atlas of Economic Complexity* [50], *Our World in Data* [42], or *The Atlas of Sustainable Development Goals* [39] summarize and explain large and multidimensional datasets and explain key insights about global challenges, such as climate change, sustainable development, economics, artificial intelligence, or cultural discoveries. Many of these projects are created by large organizations such as the World Bank and the UN, or research labs at prestigious institutions, with a mission to promote open data, methodological transparency, and education.

Many of these projects call themselves “atlas”, “observatory”, or “index” and cannot easily be described through existing genres or forms of visualization alone, such as dashboards [4], narrative visualizations [44], collections of journalistic articles [22], or interactive applications. In fact, these projects seem to act as umbrellas that cover a palette of techniques, providing different perspectives and tools to investigate, explore, and understand a topic through its data: interactive visualizations, key performance indicators, overviews, details-on-demand, narration and storytelling—all of which are embedded into well-organized pages, offering bespoke navigation aids, overview pages with introductions, and mechanisms for onboarding. However, we lack a detailed understanding of the uniqueness and characteristics of these kind of projects, which we decided to term *visualization atlases*, motivated by the fact that *a) ‘atlas’* is the term most frequently used across projects and *b) its compelling analogy to geographic atlases with their purpose of mapping ‘unknown’ terrain for lookup, learning, and exploration*. Hence, we ask *What is a visualization atlas and how to inform its design?* We aim for a definition, design patterns, a typology of atlas types, and to understand the motivations that lead to their creation. We believe this

can inform future atlases, further study, and respective authoring tools.

To that end, this paper adopts two complementary methods. First, we collected 33 atlases from the internet, as no scientific writings on visualization atlases exist to date. We then analyzed those atlases’ design and identified 45 **design patterns** along 8 design dimensions (*Sec. 3*): *content page design, entry page design, atlas structure, visualization interaction, atlas navigation, visualization types, atlas onboarding, and data transparency*. Second, we conducted **semi-structured interviews** with eight atlas designers (*Sec. 4*) who specialized in visualization design. We asked about the motivation for the atlas project, the intended audiences, and what they see as the characteristics of this atlas format. These two approaches allowed us to identify 5 higher-level visualization atlas **genres** (*Sec. 5*) that describe usage scenarios and imply design patterns.

Finally, we formulate 9 key characteristics of atlases with respect to three main aspects: **TOPIC** (Visualization atlases present *complex* topics to a wide range of audiences in a *data-driven* and *comprehensive* way), **CURATION** (Visualization atlases are *scoped, structured, and visually curated*), and **VISUALIZATION** (Visualization atlases are *visualization-driven, explanatory* and *exploratory*) (*Sec. 6*). Based on these characteristics, we formulate a shorter definition of visualization atlases that is more memorable: it defines a visualization atlas as a compendium of (web) pages aimed at explaining and supporting exploration of data about a dedicated topic through data, visualizations and narration.

All materials including atlas cases and detailed descriptions of design patterns can be found online: <https://vis-atlas.github.io>.

2 VISUALIZATION ATLASES: EXAMPLES AND COLLECTION

We start our investigation by describing three representative examples of visualization atlases (*Fig. 1*) to give background and prepare our more detailed investigations in Sections 3 and 4.

The Atlas of Economic Complexity [50] (ECONCOMPLEXITY) by Harvard’s Growth Lab reports on economic dynamics and growth opportunities for 250 countries worldwide, centering on the Economic Complexity Index (ECI). The online visualization atlas is fed by research published in an academic book of the same name [23] but supports interactive exploration and storytelling to present the underpinning economic theories. The atlas is designed for long-term usage; the underlying data is updated on a yearly basis.

The atlas features four qualitatively very different complementary components, shown on its landing page in *Fig. 1-2*. The first component is an exploratory visualization tool where a visitor can freely

-
- Jinrui Wang and Uta Hinrichs are with The University of Edinburgh. E-mail: s2251950,uhinrichs@ed.ac.uk.
 - Xinhuan Shu is with Newcastle University. E-mail: xinhuan.shu@newcastle.ac.uk.
 - Benjamin Bach is with Inria and The University of Edinburgh. E-mail: benjamin.bach@inria.fr

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org.
Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxx



Example: Atlas of Economic Complexity

Example: ATLAS of Sustainable Development Goals 2023

Example: Our World in Data

Fig. 1: Three visualization atlas examples.

customize the data dimensions and countries they are interested in. A variety of visualizations are provided in full screen: treemaps, node-link diagrams, stream graphs, maps, etc. A step-by-step walkthrough familiarizes users with the customizing features. The second component is a set of dedicated country profiles; each profile starts with a key metrics overview and then features a slideshow [44] explaining how economic complexity theory applies to a country. Readers can move between slides linearly or choose topics from a menu including market dynamics, growth opportunities, export statistics, etc. On each slide, the visualization is almost full-screen and accompanied by factual and automatically-generated statements (“*The United States of America exported products worth USD \$2.55 trillion in 2021*”). The remaining two components in this atlas are an article that explains growth projection methods and a page featuring three interactive visualizations for reporting country & product complexity rankings. The atlas features extensive documentation on methodology and data sources; videos are provided for atlas usage, key concepts, research talks, and publications.

The atlas is used in Harvard classes and the Growth Lab’s researchers use it for field research with governments across the world. Policy makers in local and national governments use it as a reference tool. The atlas is supported by a dedicated team of designers and developers based at the Growth Lab which is authoring additional visualizations and storytelling pieces, such as another visualization atlas for economic data on a city level (*Metroverse* [54]).

The Atlas of Sustainable Development Goals 2023 [39] (SDG2023) shows the latest insights into global progress and challenges towards achieving the 17 SDGs [33] defined by the United Nations in 2015. The project was commissioned by the World Bank as a sequel to their previous atlases in 2017 [61], 2018 [62], and 2020 [40]. Data are mainly from the World Development Indicator compiled by the World Bank about global development and living quality. SDG2023 consists of 17 pages, each dedicated to an individual goal and starting with a brief description of the respective goal and its sub-targets. Visualizations include a huge variety of well-designed linecharts, treemaps, small multiples, and others. Many of the visualizations are explained through scrollytelling, using piece-meal explanations and animations.

Our World in Data (OURWORLDINDATA) [42] aims at addressing global challenges using data-driven methods. While not explicitly calling itself an atlas, the project shares significant similarity with the other two examples: it publishes diverse data about poverty, hunger, diseases, under an extensive topic hierarchy (10 topics, 115 subtopics); it makes consistent use of data visualizations; it provides in-depth analysis written by scientists; and it aims to make data and knowledge public. As of March 2024, the atlas features over 4543 charts. Rather than relying on a limited set of data sources, part of the mission of OURWORLDINDATA is to find and aggregate diverse data sources and republish them for dissemination.

3 VISUALIZATION ATLAS DESIGN PATTERNS

To inform our investigation of atlas design and characteristics, we systematically collected atlas examples and identified respective design patterns. Here, we report on the collection and consequent coding.

3.1 Structured Atlas Collection and Design Classification

Phase 1: Seed collection—We started from a seed collection of 8 renowned or award-winning atlases we came across on the internet that focus on global issues, three of which are described in the above section. We found they are usually named as “atlas”, “observatory”, “index”, or “portal” featuring complicated datasets.

Phase 2: Keyword search—Based on the seed collection, we considered two branches of collection strategies: by project names and by creators. We began the keyword search by using common names found in the seed collection like “visualization atlases”, “data observatories”, and kept those focusing on specific domain topics. Meanwhile, we searched related projects created by global data publishers like universities, IGOs and their collaborating design agencies. The collection was further expanded when we interviewed creators who contributed cases from their related fields. This process led to a pool of 41 cases.

Phase 3: Refine—We further examined all the cases and iteratively refined our understanding of visualization atlases. We delineated its scope by considering three exclusion criteria. First, we prioritized web-based atlases in this paper and excluded the printed forms, due to our deliberate focus on open data practices and accessibility to datasets. These printed atlases (e.g., [3, 11, 32]) were limited to affording interactions and exploration. Second, we discarded some cases that were purely online repositories for datasets (e.g., [21]), as we considered visualizations a pivotal component in atlases for exploration and explanation. Third, we also carefully discarded cases lacking an observable structural curation or those ambiguous in selecting a definitive topic (e.g., [52]). As a result, we reached a collection of 33 atlases.

Phase 4: Pattern classification—We went through several iterations for the dimensions (pattern groups) and patterns. Firstly, we explored the contextual background of the collection by coding each example’s domain, organizational creator(s), type of organization and data sources (see supplementary materials). Then, we analyzed the design composition of atlases to explore their design space. We looked at *a*) concepts we knew from other visualization systems such as visualization types, means for storytelling, and onboarding mechanisms, as well as *b*) concepts we found were very specific to atlases such as overview pages, page structure, and navigation. To ensure these dimensions were distinct and complete, all authors independently coded five atlases along all those dimensions. We then resolved ambiguities and, refined, and added dimensions and patterns. For example we discarded the dimension ‘exploration’ and merged ‘explanation features’ with ‘content page styles’. For each dimension, we listed possible patterns with name and detailed description, e.g., the full screen choropleth map with filters are exploratory vis page in OPPORTUNITY [12]. Then, all authors coded the same five atlases again with the refined schemes and pattern descriptions. After several iterations, two authors coded the entire collection and eventually reached agreement (Fig. 2), including 8 dimensions and 45 patterns explained in the following sections.

From our classification, we describe a visualization atlas as made out of *pages*, a notion that roughly corresponds to a webpage. Links on one page can lead to another. Across atlases, we identified two types of

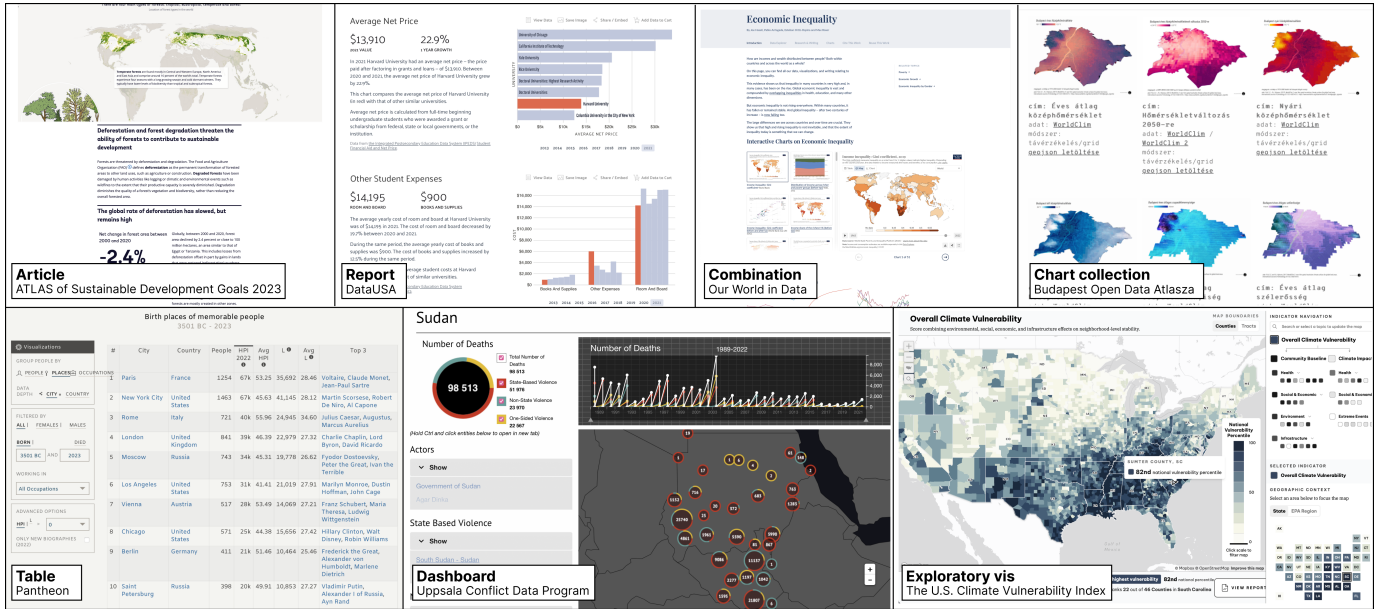


Fig. 3: Examples of content pages styles.

demonstrate the amount and diversity of data available in the atlas, such as the number of pages in a subtopic (DATAUSA). For those that use visualizations for data preview, we identify **vis screenshots** (10; 31%) that cropped directly from the content pages (ECONCOMPLEXITY, Fig. 1); and **cover vis** (6; 19%) that is a teaser specifically designed for the entry page, e.g., the rotating globe in SDG2023 (Fig. 1).

A third and last group of components on entry pages aims to foster engagement. **Exploratory vis** (8; 25%) can serve directly as an entry page, e.g., an interactive choropleth map page that links to individual pages (OPPORTUNITY). Some entry pages allow users to **personalize** (11; 34%) their experience and entry point by, enabling the selection of a country of interest (SDG2023) or specifying personal priorities on a set of sliders (BETTERLIFE).

3.4 Atlas Structure Design Patterns

The atlas structure describes how content pages are organized and linked. The page organization can combine different structural options; 12 atlases (37%) in our collection combine two structures.

- In a **parallel structure** (26; 79%), content pages have the same page styles (Sec. 3.2) and roughly follow the same structure and elements inside each content page. In SDG2023, all pages are articles, use scrolly-telling, show similar information about each SDG, and follow the same layout, colors, and other design elements. USCLIMATE is an example that creates reports for each location, based on exactly the same report template. Parallel structures can make it easy for visitors to create a mental model of the atlas (Sec. 3.9, Sec. 3.5) because of repeated and familiar designs. Parallel structures can also help by adding more data and pages in the future. Parallel structures also suggest that the atlas topic is presented from the perspectives of several ‘main characters’, i.e., individual concepts such as indicators, countries, persons, etc.
- In a **hierarchical structure** (7; 22%), content pages are grouped into *subtopics*. Usually, each step of the hierarchy / subtopic is introduced by a dedicated entry page. Content pages in hierarchical structures can follow the same page style or be entirely different. In OURWORLDINDATA, all content pages are organized around subtopics such as life expectancy, gender ratios, and migration, and each topic has a group of pages like articles and exploratory visualization. Another example is CALVINO [45], where the hierarchy is illustrated on the entry page as a map with three itineraries and three stages.
- A **multiple facets** (10; 31%) structure is similar to parallel hierarchical structures but instead of each content page focusing on a *different* part of the data, content pages represent facets “of the same coin”. An example for multiple facets is BETTERLIFE, which offers content

pages about overall indicators, topic breakdowns, country comparisons and user stories. Multiple facets provide complementary views. They can be good solutions for large and multidimensional datasets where simple parallel or hierarchical structures are too constraining. On the flip side, multiple facets might require more advanced navigation features in the atlas (see Sec. 3.5).

If an atlas consists of pages that are only loosely connected, as in ACLED [2] that positions different conflict data resources like country dashboards and case studies without a clear categorization, then we categorize it as **no structure** (2; 7%).

3.5 Atlas Navigation Design Patterns

Navigation across an atlas guides the user to individual pages, helps orient which page they are currently seeing, and may suggest which parts of the atlas they have not yet seen. Navigation within an atlas and across content pages is a prime challenge in atlas design that naturally becomes more tricky the more extensive the atlas is: the more content pages, topics, deeper hierarchy and/or multiple facets. **Search** (18; 55%) is a simple navigation means for large atlases with complex data (METROVERSE). Creating different **visual identities** (8; 25%) per page or subtopic can tell a visitor where they are, e.g., by using colors or icons (SDG2023). **Structural overviews** (24; 73%) include tables-of-content, menus, breadcrumbs, progress bars, and navigation-tabs that show the entire or partial page structures. Eventually, **page links** (14; 43%) directly link to other pages within the atlas (OURWORLDINDATA).

3.6 Visualization Types Design Patterns

Many visualization atlases include multiple visualization types with different levels of complexity and familiarity. The majority of atlases include **basic charts** (25; 76%) (linecharts, barcharts, and variations thereof) and **maps** (29; 88%). However, maps are used to different degrees: mainly maps; alongside other charts; and few maps (shown as differently sized rectangles in Fig. 2). Only 5 (16%) atlases rely entirely on maps as the only visualization type. Six (19%) atlases use some sort of **network** visualization (e.g., node-link diagram, adjacency matrix, sankey diagram). Other atlases include **treemaps** (5, 16%), **projections** (3, 10%) and/or **glyphs** (4, 13%) for high-dimensional data. Visualization types that occur only once or twice in the collection, like radar charts and violin charts, were coded as **other**.

3.7 Visualization Interaction Design Patterns

This dimension describes interaction within a content page to explore data and visualizations. **Creating focus** (29; 88%) draws attention to

or highlights selected data points through, e.g., asking which country is of users' interests (SDG2023). **Switching views** (11; 34%) changes the visualization techniques, e.g., from a treemap to geographic map (ECONCOMPLEXITY), offering multiple complementary perspectives of the data. **Coordinated views** (6; 19%) help compare two or more data sets; e.g., reading university reports side by side (DATAUSA). Some atlases offer to **customize data** (25; 76%) through filtering, brushing, selecting indicators, and visualizing different datasets such as in the example of OPPORTUNITY's advanced tool. Customized views are often intrinsic in exploratory visualization page styles (Sec. 3.2).

3.8 Data Transparency Design Patterns

Patterns in this dimension describe how an atlas' data, analysis, and visualizations are made transparent. It can be done by disclosing **data sources** (31; 94%) through inline or endnote citations or directing to databases and respective websites; by **exports** (29; 88%) to download data or charts as figures; through **open source** code (5; 16%) for data analysis (IPCCATLAS) or visualization implementations (SDG2023). **Quality evaluations** (5; 16%) are explicit explanations of missing data, their reliability, and possible data loss. This can happen through dedicated scores (e.g., low/medium/high data quality per city METROVERSE) or through bespoke visualizations showing data availability (GENDERDATA). Lastly, **methodology** (20; 61%) explanations describe how data is collected and calculated, the criteria for inclusions and exclusion, and if there exist domain research publications.

3.9 Onboarding Design Patterns

Our last dimension describes help resources for visitors to understand how to use an atlas in terms of contextual background, general features, and reading visualizations. Drawing from Stoiber's taxonomy of onboarding types for visual analytics systems [49], we identified three similar types in atlas onboarding: step-by-step **walkthrough** (8; 25%) tutorials that guide users through an atlas' interactive features; **introductory videos** (9; 28%) demonstrating an atlas' usage (IPCCATLAS [26]), structure (CALVINO), or key findings (SDG2023); as well as **textual user guides** (11; 34%). In addition to this taxonomy, we found a new onboarding category used in five atlas cases: **glossaries** (5; 16%) with explanations about an atlas' background, key metrics, and calculations.

4 INTERVIEWS WITH VISUALIZATION ATLAS CREATORS

Complementing our analysis of design patterns, we interviewed atlas creators to learn about the intended purposes of the atlases, overall motivation, design practices and challenges encountered.

4.1 Interview Procedure

Participants: We interviewed eight visualization experts (atlas creators) who were directly involved in the design of seven visualization atlases. Creators' professional backgrounds included 5 visualization designers, 1 full-stack vis engineer (P4), and 1 data journalist (P6). Although some participants have worked on several visualization atlas projects, we intentionally interviewed them based on a specific atlas. Atlas projects focused on during interviews included a range of topics (e.g., literature, sustainable development, and technology), different domains and author organizations (e.g., universities, IGOs, design agencies, and research institutes). 3 participants were internal to the atlas' author organization; 5 were hired as external designers (Table 1).

Interview Structure: One-hour interviews were conducted individually with each participant, except for P2 & P3 who were interviewed together because they worked on the same visualization atlas. The semi-structured interview questions focused on seven aspects: the visualization atlas background, goals & motivations, collaborative process, target audience, design approaches, and a reflection on the nature of visualization atlases, compared to other visualization forms.

Data Collection & Analysis: All interviews were audio-recorded and transcribed to enable an in-depth thematic analysis [7]. The seven questions from the interviews directly informed our coding; for each of the questions we created a coding category using NVivo [28]. Then, we compared the coded transcripts for each category across the interviews

for common sub-themes, e.g., 5 out of the 8 interviewed mentioned *data insights* as a goal, forming the basis for the following findings.

4.2 Motivation & Goals for Creating Visualization Atlases

Promote Selected Datasets for Exploration & for Taking Action. 4/8 creators highlighted the promotion of a selected dataset or range of data that had been collected and curated by an associated organization as one of the core motivations for creating a visualization atlas. To that end, interactive visualizations are created to make this data explorable by people within and outside of the organization. For example, ECONCOMPLEXITY: *"started as a side product of a book and [the Growth Lab] hoped to make the data visualizations more interactive, exploratory for people."* [P8]. Similarly, the SDG SERIES showcased and communicated data owned by the World Bank in the form of atlases since the 70s, first in print, then digitally as PDFs and now in the form of webpages that feature interactive visualizations.

Some atlas projects aim to act as a call for action. For example, BUDAPEST is based on public data about Budapest painstakingly put together by a data journalist (P6). The atlas can be seen as a call for action for making recent national development data open access, to raise awareness on urban development issues, and to hold public administration accountable for data transparency. Similarly, the driving force of USCLIMATE, is to *"inform local communities and give them some advocacy power to apply for funding in more appropriate ways"* [P7].

Publish Insights & Stories. 5/8 creators reported the dissemination of insights in the form of stories and *"demonstrations"* [P8] to be a central goal of atlas projects. For example, CALVINO was driven by the idea of *"producing analytical narrations using visualization in literature criticism"* [P1]. The atlas provides three narrative strands that combine literature analysis with interactive visualizations. The SDG2023 emphasizes latest insights *"on each [sustainability] goal through a lens of current events"* [P2]. Interestingly, in 3/7 projects, the goal of highlighting insights and stories arose during the atlases' design process. For example, while designing ECONCOMPLEXITY the creators' team *"realized that people just simply do not understand [the economic complexity theory]"* [P8], so they included narrations in the country profiles that explain the terminologies used.

Promote Topics in a Systematic, Data-Driven Way. For 4/7 visualization atlases, the data was in place from the beginning, directly informing design processes. In the remaining three atlas projects, the systematic identification and curation of data sources on the topic in focus motivated the atlas project. For example, the goal of AIPOLICY [34] was to *"build insights that are not visible from any single data sources over time"* [P4]. The team therefore not only tried to identify data sources relevant to this specific atlas, but also came up with updatable and reproducible data sourcing methods that can be applied across different topics and atlases: *"[F]or each observatory there are generic data sets that work on any topics like [fetching latest] research papers [on this topic]."* [P4]. Customized data sources are then included *"that are specific to the topic, [...], such as climate data for climate-related atlases, or data from Stack Overflow for tech-related atlases"* [P4].

Keeping data up-to-date is key for some atlases. The team of RARE DISEASE [51] considered the accessibility, update frequency, and data source reliability to ensure data tracking and updates and that the visualization atlas remains relevant in the long run. They even collaborate *"with users [in] that we constantly add new data strings"* [P5].

4.3 Intended Audiences of Visualization Atlases

Closely linked to their goals, many visualization atlases mention a wide range of target audiences, including researchers, policy makers, the media, and the general public. We interviewed creators about the intended audiences and how this influenced their design process.

Internal & External Experts. 5/7 atlases target expert audiences internal to the atlas team or external domain experts. *Internal experts* work at the institution issuing the visualization atlas and are often closely involved in its design process (e.g., by informing data stories, as in the case of SDG2023 and ECONCOMPLEXITY). They have extensive domain knowledge and are familiar with the underlying data. While their visualization expertise may vary, their involvement in the

Participant & Role	Vis Atlas Name	Domain	Organization
P1 - Vis designer	Atlante Calvino (CALVINO) [45]	Literature criticism	University of Geneva, the Politecnico di Milano
P2, P3 - Vis designers	ATLAS of Sustainable Development Goals 2023 (SDG2023) [39]	Sustainable development	World Bank
P4 - Full-stack engineer	OECD AI Policy Observatory (AIPOLICY) [34]	Artificial intelligence	OECD
P5 - Vis designer	Rare Diseases Observatory (RAREDISEASE) [51]	Health	The Jožef Stefan Institute
P6 - Data journalist	Budapest Open Data Atlasza (BUDAPEST) [10]	Urban development	ATLO
P7 - Vis designer	The U.S. Climate Vulnerability Index (USCLIMATE) [19]	Climate	Darkhorse Analytics, EDF, Texas A&M University
P8 - Vis/UX designer	Atlas of Economic Complexity (ECONCOMPLEXITY) [50]	Economics	Harvard Growth Lab

Table 1: Background of eight participants from seven projects, including their roles, associated projects, project domains, and initiating organizations.

creation process grants them the ability to inform visualizations and related atlas features. Internal experts use visualization atlases to inform research on the topics presented (ECONCOMPLEXITY & AIPOLICY), make presentations and/or demonstrations. Visualization atlases may also be used for teaching: “[The PI] would request some of the features [...] in order to teach a certain topic in his class” [P1; CALVINO].

External experts are academics and analysts external to the institution or project team issuing the visualization atlas. They have extensive domain knowledge, but may not be as familiar with the presented data and included visualizations. They typically will use the visualization atlas as a reference tool (e.g., in the case of CALVINO), or to inform their research: e.g., RAREDISEASE also targets doctors and clinicians, who are “working on developing new treatments for rare diseases” [P5].

Policy makers. 3/7 atlas projects specifically target “policymakers” which can include a range of roles and usage scenarios. In the case of USCLIMATE, they are “US-based decision makers at different levels [...] who] aim to make positive systematic changes” [P7]. In ECONCOMPLEXITY, a pilot user study identified policy makers as “mainly people working in the data office, both locally and nationally, [who] care about specific issues” [P8]. In AIPOLICY policy makers were defined as “not technical experts in AI, but [people who] would like to be informed about AI progress, especially in their own countries, and other similar countries” [P4]. Our creators agreed that senior policy makers, depending on their background, expertise, and interest, may not directly engage with the visualizations included in the atlases, but, similar to external experts, they may use a visualization atlas as a reference tool that can provide multiple resources on a topic to comprehensively understand its challenges.

The General Public. 5/7 atlases were designed with a “general public” audience in mind. Again, vaguely defined, this can include people without or with only little knowledge about the topic covered in the atlas, as well as people without expert knowledge in data analysis or visualization [8]. Some creators stated that they address people who have “at least [seen] a chart before” [P2, P3; SDG2023] and that have an invested interest in the respective topic, e.g., because they are directly affected. For example, P7 framed the intended audience of USCLIMATE as “people from the trenches on the ground level, [including] people engaged in local communities who care about climate issues deeply and inform funding proposals.” RAREDISEASE targets the parents and children who are directly affected by rare diseases.

4.4 Design Practices & Challenges

Design Inspirations. Creators mentioned a range of inspirations that influenced their design, sometimes visual approaches, related projects, or broader practices from related communities. P8, for example, mentioned that they checked related atlas OPPORTUNITY [12] or open data project like Trase [53] “to see the new trends in the economic visualization world” [ECONCOMPLEXITY]. Other designers drew from science communication and data journalism in their design approach: “data journalism was the main inspiration, just because they [the visualizations] look really engaging. Whereas science communication is sometimes more dry looking.” [P2; SDG2023].

Audience-Centred Design. Most atlas projects target a range of diverse audiences (see above), and creators acknowledged this as a major design challenge. However, only 3/7 (RAREDISEASE, AIPOLICY, and ECONCOMPLEXITY) of the atlas projects in focus throughout our interviewed applied user-centered methods such as surveys, focus

groups or interviews with target audience representatives as part of their design process. Some atlas creators acknowledged the lack of user-centered approaches as a potential problem (SDG2023), some highlighted intentions in studying the use and impact of their visualization atlases after it was finalized: “over time, we can do user interviews [...], it is usually post projects” [P7, USCLIMATE].

Modes of Collaboration. Interviews show that designing visualization atlases is highly interdisciplinary; 6/7 atlases were created by interdisciplinary teams that included visualization designers, web developers, and domain experts. Only one atlas (BUDAPEST) was created by a sole designer. *Parallel Collaboration* was mentioned in the context of 3/7 atlases (CALVINO, ECONCOMPLEXITY, SDG2023) with distinct roles assigned to domain experts and vis designers, in that domain experts would provide curated data alongside written stories, and the designers focused on the visual design and user experience. *Close Collaboration between Domain experts & Vis Designers* was mentioned in the context of atlases (AIPOLICY, RAREDISEASE) where data was gradually sourced across the projects, and insight generation was still evolving during the design process. Here, domain experts and designers worked closely together to decide (a) what data and data stories should be included in the atlas and (b) how to design visualizations and the atlas’ general interface. Our interviews indicate that this kind of collaboration allows for a gradual and experimental process of atlas creation, where different versions are explored, before considering a final design that may then be expanded to a wider range of data or topics.

4.5 Creators’ Voices: What Makes a Visualization Atlas?

The majority of projects covered in our interviews are named an “atlas” (CALVINO, SDG2023, BUDAPEST, ECONCOMPLEXITY), but other, related terms such as “Observatory” (AIPOLICY, RAREDISEASE), or “index” (USCLIMATE) are also used. We therefore interviewed creators about the key aspects that, to them, make for a visualization atlas in comparison to other closely related approaches.

Defined Topic + Breadth of Perspectives & Data Sources. Our interviews reveal that key to visualization atlases and comparable projects is the focus on a *specific topic*, presented in a data-driven way: “I think you need a topic. [...] So the sustainable Development Goals served as our topic.” [P2; SDG2023]. “The main idea is to kind of collect all the data and show this overview for whoever is new [or] interested in the topic and wants to explore the details.” [P4; AIPOLICY]. Creators also agreed that this topic should be approached from multiple perspectives and data sources: “You try to cover [the topic] throughout, not just like one angle, but different views of it [including coverage from research as well as the news].” [P4; AIPOLICY].

Topic & Visual Curation. Curation was highlighted as another key aspect of visualization atlases, concerning the topic, data and visuals: “You need a topic, you need a method of inclusion and exclusion, and you need a stylistic theme, that is coherent across the atlas.” [P2; SDG2023]. Topic and visual curation are linked: “every page helps you in understanding a topic or an aspect of a topic with the aid of visual materials.” [P1-CALVINO]. 5/8 creators acknowledged **data curation**, i.e., deciding the in- or exclusion of data sources, as a key part of the atlas design process: “We had lots of discussions and trial-and-error getting the data, looking at it and [deciding] if we are missing something here, or if this is too much to present in a good way.” [P5; RAREDISEASE]. The need to balance the granularity of data was also highlighted: “It’s great to look at the broad trends, but can we get to

more granular [perspectives]? So this broad view and then a specific area profile view marry each other very well.” [P7; USCLIMATE].

Topical curatorial decisions can be political, reflecting on the authoring organization’s identity and/or intentions. “We decided against [showing all of the targets in SDGs], because the World Bank is not the organization defining those targets. So they didn’t want to overemphasize the whole structure [which] was more of a UN responsibility.” [P2; SDG2023]. At the same time, creators highlighted the need for transparency regarding curatorial decisions and data limitations because “for a traditional atlas, you give your trust to the cartographer or the statistical bureau that everything is 100%. [...In] the data cut from international databases for Hungary, you can see some imperfections, inconsistencies in the data with our local knowledge” [P6; BUDAPEST].

Our interviews further indicate that atlases demand **visual curation** to ensure visual consistency across the different sub-topics that may be presented within an atlas, which affects the design of corresponding visualizations: “you should design a sort of a graphical language that is reoccurring through the entire Atlas and with like color palettes and fonts and margins, layouts, grids and those kinds of things. Because you should be able to recognize that different visualizations are belonging to the same project.” [P1; CALVINO].

Debates on Comprehensiveness. An aspect that creators debated was the comprehensiveness of topic coverage. Creators stated that “you need to have a wide range of data, [...] it has to cover the whole width of one specific aspect that is being worked on.” [P5; RAREDDISEASE]. However, some creators argued that atlas topics can be so big that comprehensiveness is not achievable. For example, P2 stated that the SDG2023 is “less comprehensive in terms of how each goal is covered. I was surprised by that in the beginning, because I think my notion of an ‘atlas’ would have meant that we are talking about every single target, [...] I understand why this is not feasible, because then you cannot go in depth enough at the same time. So they [the World Bank] wanted some depth in one aspect of each goal, rather than showing very surface-level insights.” [P2; SDG2023].

Global & Geospatial Perspectives. There was quite some debate about the necessity of a global, geospatial perspective as a defining element of visualization atlases: “I take the atlas in a very literal way, defining a geographic coverage that could be any kind of data encoding on top of it, like we map the economic data.” [P8; ECONCOMPLEXITY]. However, other creators questioned the global nature of visualization atlases as a necessity: USCLIMATE “is a national level data set. It doesn’t matter which country that you have, but different levels of geographic granularity” [P7]. P6 questioned the necessity of geospatial aspects: “You want to see and understand visually what’s happening, and it’s not related only to maps or locations.” [BUDAPEST].

Up-to-Date Information. Creators further discussed the necessity for keeping visualization atlases up-to-date, which is linked to the requirement of working with dynamic data sources. For example, a new edition of the Atlas of Sustainable Goals is released every couple of years, often with a new set of targets in focus. Therefore, previous issues are not updated but show a snapshot in time. In contrast, presenting up-to-date information is a focus of RAREDDISEASE and AIPOLICY. We found that naming conventions in the space of atlas-related projects matter. For example, the creators of SDG2023 highlighted the term ‘atlas’ “emphasizes the comprehensiveness [while] I think by calling it ‘observatory’, they want to draw attention to the dynamic updates” [P2; SDG2023]. This was confirmed by the creators of AIPOLICY and RAREDDISEASE: “I would like to think about an ‘observatory’ as [something] having multiple streams of data, but also data gets updated. So it’s about [...] observing through time.” [P4; AIPOLICY].

The Role of Interactivity. The importance of interactivity came across in 5/7 atlas projects mentioned in interviews. Although interactivity was not a key aspect in their design, the creators of the SDG2023, for example, mentioned that they included interactive elements in the visualization atlas to separate it from previous print-based issues. In other cases the support of interactive exploration is a key aspect: “Those visualizations really need to be interactive in order to make sense, [...], when you look at that feasibility chart, there are just so many nodes that it doesn’t make sense without tooltip.” [P8; ECONCOMPLEXITY]. Sim-

ilarly, for the USCLIMATE the “engagement factor is crucial. That ability to click around and drive it myself.” [P7].

Comparison to Other Forms of Visualization. The interviews also brought up debates about how visualization atlases relate to other types or forms of visualizations because of their difference in monitoring or storytelling purposes. The designer for RAREDDISEASE argued against infographics “because the idea is to to have it like a exploratory tool”, and she acknowledged an atlas can be made of “several dashboards put together with some storytelling”. Yet P1 shared the opposite opinion that “the idea of narration is something that is very important to have in a project that is called Atlas. Otherwise, for me it’s just a dashboard to compare, combine and visualize complex data. But it doesn’t convey you the story”. Similar ideas came across from the creators of SDG2023 that “A dashboard is meant to monitor, usually updating data. So we don’t have that [in the SDG2023]” [P2].

5 VISUALIZATION ATLASES GENRES

The design patterns in section 3 provide low-level design solutions and common practices while our interviews provide very detailed information about the rationales behind atlas design. In search of a practical middle ground, we now describe atlas *genres*. Inspired by previous work on genres in narrative visualization [44] and dashboards [4], those genres can serve as templates for specific purposes and scenarios while informing the design decisions such as the choice of design patterns.

Informed by previous methodologies [4, 31, 44], we searched for similarities between atlases based on their design dimensions. In this process and inspired by our interviews, we added one new dimension to the atlas coding—**data**—which was not part of our visual design pattern analysis. The data dimension captures if (a) data in an atlas are updated or static, (b) if data come from internal or external sources, and (c) if an atlas is based on a single or multiple datasets.

We used Bertifier [38] and its table ordering techniques [6] to order rows (atlases) in Fig. 2 by similarity of patterns across all dimensions, resulting in three visual clusters. Then we integrate interview findings in terms of motivations and topics to refine the clusters into 5 genres: Storybanks, Observatories, Series, Monographs, and Exploratoria.

Genre: Storybank aims to promote a defined, yet broad topic in a systematic, data-driven way. They are driven by multiple data sources, often external, that are updated regularly. Updates to a storybank may include the expansion of the topic and subtopics covered. Key atlas examples include OURWORLDINDATA and DATAUSA.

From a design perspective, storybanks are often defined by a large collection of articles or reports followed by an extendable structure of hierarchy/multiple facets plus parallel pages. In OURWORLDINDATA, the topic hierarchy creates placeholders for inserting new articles, insights, or charts. Multiple entry points are provided, including featured content to enable visitors to pick-up recent highlights. However, storybanks have limited onboarding features, which, combined with their breadth and amount of content, can make the navigation more difficult.

Genre: Observatory shares similar design features as storybanks, but the topics are more narrowly defined and systematically tracked based on key performance indicators. The AIPOLICY is a key example, where indicators are sourced externally, processed, and presented together by-country and -topic to track the latest updates in AI policy.

Genre: Series are published in more or less regular intervals. The visualization atlas series focuses on a high-level theme and each issue will not be updated systematically. Instead each issue features a selected (range of) topic(s) under that overarching theme, usually pertinent to its specific time periods. The topics are driven by internal data sources that frame in-depth stories, usually with a clear message. From a design perspective, atlas issues feature a simple, parallel structure, consistent page style and different visualizations. However, the overall style, featured pages can be unique to each atlas issue and evolve over time. The SDG2023 is the only example of a visualization atlas issue we discovered, belonging to SDG SERIES issued by the World Bank. Over time, the series has featured more visualizations with more interactivity.

Genre: Monograph is a visualization atlas focusing on a highly specific topic with the purpose of promoting these for exploration and taking action. The corresponding data sources can be quite unique to

the topic. Monographs are typically finished pieces without regular updates. As such, visualization atlas monographs are typically less expansive with their focus in depth than breath, featuring detailed stories and insights. From a design perspective, atlas monographs can follow tailored overall designs and highly customized visualizations. Key examples include the CALVINO and the BUDAPEST atlas.

Genre: Exploratorium promotes the active and open-ended exploration of particular topics by making extensive use of interactivity. Our exploratoria examples (such as the ECONCOMPLEXITY) are driven by internal datasets that are maintained and updated regularly by the author organization. Exploratoria can have quite complicated designs: they typically feature at least one *exploratory vis* page as the core exploration tool, sometimes, but not always with supporting pages of either *articles* or *reports* to explain different perspectives on the topic in *parallel*. There are also often multiple entry points and navigation strategies which are communicated through a range of onboarding strategies.

6 KEY CHARACTERISTICS OF VISUALIZATION ATLASES

In this section, we review our findings in detail and unpick 9 key characteristics grouped into three high-level aspects: TOPICS, CURATION, and VISUALIZATION that best describe the nature of visualization atlases. The characteristics then inform a condensed and practical definition that a visualization atlas is a *compendium of (web) pages aimed at explaining and supporting exploration of data about a dedicated topic through data, visualizations and narration*. To examine whether a project qualifies as a visualization atlas, the condensed definition captures the purpose and context of visualization atlases while the characteristics we describe in this section are specific criteria for inclusion and exclusion of what constitutes a visualization atlas.

TOPICS—Visualization atlases present complex topics, targeting a wide range of audiences. Topics are complex, data-driven and comprehensive.

Complex—Atlas topics are often of wide interest and target diverse audiences that include domain experts, either internal or external to the project, for research and further analysis, audiences in policy for general information and dissemination (rather than directly for decision making), as well as members of the general public including those in “the trenches”: community workers and those directly concerned by the topic of the atlas. As such, the usage scenarios of atlases are very broad, including promoting important global topics in a data-driven way, promoting these data sets for exploration and action, and publishing insights and stories about both the data and topics. They aim to combine breadth and depth (horizontal and vertical exploration [16]).

Data-driven—Visualization atlases present topics through data. They report and analyze typically multiple data sources around the topic. More specifically, they report on selected findings from these data, in the form of visualization-driven stories, reports or more exploratory interactive visualizations. By drawing from and carefully selecting data, visualization atlases promote the topic and enhance understanding of the available data—or the eventual absence of data.

Comprehensive—Atlases aim to cover the respective topic in a meaningful and well-informed breadth and depth. For example, USCLIMATE covers the entire United States on the level of tracts, counties, states and EPA regions, reporting climate vulnerability from two groups of indicators: climate impact and community baseline. Likewise, SDG2023 covers all 17 sustainable development goals.

CURATION—Visualization atlases are highly curated in terms of content, data and presentation; they are scoped, structured, and visually curated.

The choices about topics and data require deliberate curation throughout the entire process of the atlas creation. Additional choices are required in combination with the targeted range of audiences and intended usage scenarios. To this end, the author organizations—global NGOs, large or small research institutes, or design agencies—often bring together interdisciplinary teams with expertise in the topic, data, visualization and user interface design.

Scoped—Visualization atlases are highly scoped rather than being open repositories for arbitrary kinds of data. Content curation concerns the framing of the topic, the choice of data, and the framing of stories. For example, in some atlases, certain data and related stories are deliberately excluded to maintain scope, to keep the content accessible to large audiences, but also to reflect on the nature, responsibilities and agendas of the organizations involved. The transparency and nature of these decisions can vary. In visualization atlas projects topic and data curation involve a handful of internal domain experts (top-down), whereas others involve the whole project team (including visualization designers) or even inform curatorial decisions through user-centred research (bottom-up).

Structured—Atlases are structured collections of (web) pages with clear purposes and delineation for each page. Atlases usually provide entry pages, cross-links between pages, clear recurring (parallel) page structures, and aids for navigation and search. Structure aims to represent the topic (e.g., subtopics) and help understanding. Onboarding mechanisms help introduce those new to the topic and atlas.

Visually curated—Visual curation concerns how the topic, data and stories are visually (re-)presented and made explorable. This includes a sense of visual consistency across the atlas, while also recognizing the individuality of sub-topics and related visualizations and pages. It also includes establishing a coherent visual identity across pages, clearly communicating navigation, and facilitating user onboarding. Visual identity can be created through using consistent and unique (combinations of) colors, font types, page layout, visualization style, etc. A visual identity is important to make the atlas a coherent piece and to support its authority. This also applies to general accessibility to meet the needs of readers with bodily impairments, responsiveness to different device contexts, and audiences with low levels of general literacy, data literacy, numeracy, or visualization literacy.

This high degree of curation ensures a focused engagement with the specific topic of the visualization atlas. Without the criteria of CURATION and TOPIC, atlases would be more like other repositories, or collections of data and visualizations. Naturally, curation processes are heavily influenced by the creator organizations, their profiles, and intentions. We will discuss this further in [Sec. 7](#).

VISUALIZATION—Visualization atlases heavily rely on data visualization; they are visualization-driven, explanatory and exploratory.

Visualization-driven—Data visualizations serve as the main medium to engage with the topic and data presented in a visualization atlas. This characteristic aligns well with open data initiatives that improve public engagement through open data and data visualizations such as discussion platforms [58] and public discourse [30]. Visualization atlases are not bound to geospatial visualization, but can incorporate a variety of charts and related visualization techniques, sometimes even quite non-standard (CALVINO). The use of consistent visualization templates can emphasize the holistic nature of the visualization atlas, while individualized visualization designs can emphasize the nature of specific sub-topics. Importantly, visualizations in atlases can follow both reader-driven and author-driven approaches [44]; they are both exploratory and explanatory:

Explanatory—Visualizations in atlases serve explanation through e.g., visual curation, annotations, textual explanations. Visualizations for storytelling are integrated into articles and reports to communicate and emphasize stories and insights (author-driven). They complement text and annotations through visual access to the data; they can reiterate the information in the text, or serve as an additional information source.

Exploratory—Visualizations also serve exploration through being interactive, providing detailed and complementary views of the data. They may attract curiosity and function as illustrations on entry pages to communicate the overall mission and topic of an atlas. Those pages that allow for more flexible, free-form exploration may take the style of exploratory visualization, dashboards, and chart collections that are rich in interaction (reader-driven).

7 DISCUSSION

Uniqueness of visualization atlases. Our definition and its key characteristics helps us discuss visualization atlases' uniqueness with respect to other forms of visualization and collections. Visualization atlases are higher-level structures that incorporate other genres and forms of visualizations in their content pages: analytic interfaces (interactive visualizations), dashboards [4] or narrative visualizations [44] such as articles [22], comics [5], videos [1], etc. Some atlases use a limited number of content page design patterns, while others include a wider range and mix components from news articles, dashboards, and slide-shows. Atlases are also different from mere collections of articles (e.g., blogs, news websites) as these collections usually lack framed *structure* of those pages, making each page a rather independent contribution without *scoped* content curation. Likewise, collections of visualizations like VISAP or survey repositories are typically not *data-driven* but rather compilations of visualizations that are neither *explanatory* nor *exploratory* in the context of the survey. Dashboards are an interesting case since they share many of an atlas' characteristics; *data-driven*, *scoped*, *structured*, *visually-curated*, and *visualization-driven*. However, dashboards focus on monitoring and at-a-glance-overviews rather than supporting explanation or being comprehensive, especially since they usually lack mechanisms for onboarding and narration. Instead, atlases can be made of pages of dashboards, one of the content page types. Finally, atlases are different from encyclopedias (e.g., Wikipedia) or online galleries (Google Arts & Culture, Pinterest), as these are not *data-driven* nor *visualization-driven* within a topic.

Limitations of methods and definitions. Visualization atlases are of considerable complexity and richness. Our definition for visualization atlases is derived from shared qualities in the survey collection. This paper aims to describe these qualities and understand visualization atlases in the context of emphasizing notions of complex topics, accessibility and structure. The boundaries between the qualities and notions can be fuzzy as each atlas case encapsulates human ideas and inventions. Our identified genres (Sec. 5) capture groups of common qualities and usage contexts. We took an explorable approach in collecting atlas cases, meaning that because of the lack of a common terminology or concept of visualization atlases, we might have missed similar projects in nature which use different names. However, we believe our definitions and characteristics describe this emerging class of projects sufficiently well to justify the existence of this new format.

Extending atlases. We did not encounter visualization atlases incorporating visualizations such as data comics [5, 59], data-gifs [47], or data videos [1, 46]. These and other forms/genres of data representation (e.g., sonification or data physicalization [29]), as well as related approaches (e.g., interactive notebooks, explainable explorables [57], immersive technology [27], or streaming [63]) offer a vast range of design opportunities for visualization atlases. All these building blocks alongside our 33 examples can inform the design of future atlas projects, especially those involving teams. They can inform design card decks [5, 43], design activities [25]. Future research could look into specific toolkits and web-design platforms for making atlas design, development, and maintenance easier. We hope this paper inspires future research to imagine new genres of visualization atlases that may also expand the design dimensions and patterns presented here.

Designing visualization atlases in a user-centred way. Our interviews have highlighted that visualization atlases potentially speak to a wide range of audiences, including researchers, domain experts, educators, policy makers, and the 'general public'. This imposes several design challenges. Some of these audience groups are notoriously difficult to define and characterize, which, in turn, makes it difficult to define usage scenarios and do justice to the diverse goals and needs of different audiences. This echoes recent research [24] that highlights the existence of varying levels of audiences' openness to public data and visualizations. Given these challenges and the significant real-world impact of visualization atlases as tools of authority, the integration of user-centered methods in the design process can more effectively address specific usage scenarios and audience perceptions. While our interviews indicate that creators are aware of this issue, only 3/7 of the atlas projects in our interview actually applied user research due

to limited resources, and more bottom-up approaches such as co- or participatory design were not mentioned at all.

We therefore argue for future design practices in the context of visualization atlases to put a stronger emphasis on allowing *a*) target audiences and *b*) people who are at the center of or potentially affected by the presented topics and data to actively participate in atlas design processes (as early as the stage of data creation/selection and curation). How this can be achieved in practice, given the complexity of visualization atlases that often already involve large interdisciplinary teams, and how to best document and share best practices is an interesting challenge but also an opportunity for future research.

Real-world impact of visualization atlases. Visualization atlases potentially have a huge real-world impact, due to the fact that they often focus on topics of global relevance and because atlases are deliberately designed to advance access to these topics. At the same time, a lot of visualization atlas projects are initiated by impactful and large organizations which give these atlases reach and credibility. Besides educating us about certain topics, atlases may promote a culture of data-driven evidence and decision-making and increase data and visualization literacy. However, the use and impact of visualization atlases in different contexts and on different audiences is largely unexplored. The creators we interviewed acknowledged that this is usually not explored as part of the atlas project. Future research in this area is crucial to *a*) inform visualization atlas design decisions, and *b*) explore their potential impact on research, education, policymaking, etc. Respective tools to investigate the real-world impact of visualization atlases could include a range of qualitative research methods such as interviews and observations. However, we may also draw from theory such as feminism (e.g., [18, 41]), or critical perspectives on visualization [13, 17] to interrogate the context and design of existing visualization atlases and the power structures into which they potentially feed.

Transparency of curation processes. As we have highlighted in this paper, curation processes are a key characteristic of visualization atlases. Content and visual curation ensures consistency and coherence in ways a visualization atlas presents and discusses the topic in focus. However, given the fact that visualization atlases potentially have a big real-world impact, making these curation processes transparent is of utmost importance and a huge challenge at the same time. Transparency not only concerns data provenance [9], i.e., the origin of data sources and the data transformation processes that preceded the visualization and narratives included in the atlas. It also concerns transparency about the goals and intentions behind the atlas, including the organizations involved and their (political) profile, the atlas team, also acknowledging the labor of external people, and the decision-making processes that influenced the selection of topic(s), inclusion/exclusion of data (sources), design processes, etc. Interviews with creators show an awareness of the importance of transparency, but also highlight that it remains a challenge (SDG2023).

Highlighting the limitations of an atlas can be quite important in this regard. Recent work on provenance in visualization [9, 56], uncertainty visualization [36, 37], and visualizing missingness in data [20, 48] may be a starting point here. Correll [13] has reminded us of the ethical dimensions of data visualization and of our responsibilities as visualization researchers and designers. In this context of visualization atlases which we can consider tools of power, just like their traditional eponyms, establishing rigorous, reliable and ethical methods is an important and interesting challenge.

8 CONCLUSION

This is the first paper to describe visualization atlases and to propose a definition for this emerging genre. We investigated 33 atlases from different angles, yielding design patterns, genres, and first-hand reflection from atlas designers. We believe visualization atlases are "here to stay" as they present solutions for how we can deal with large, complex, and urgent global topics understood through data—providing access to data, promoting knowledge, educating people, supporting further analyses, make data open. We believe visualization atlases are crucial for a data-literate society and a culture of data-driven (inclusive) decision making, balancing access to data and critical engagement.

ACKNOWLEDGMENTS

We thank our reviewers for their invaluable insights and critiques that greatly enhanced our work. Special appreciation goes to the following atlas creators for their participation and unique perspectives in interviews: Dominikus Baur, Attila Bátorfy, Tommaso Elli, Alenka Guček, Rob Korzan, M. Beshar Massri, Alice Thudt, Nil Tuzcu (listed in alphabetical order by surname).

REFERENCES

- [1] F. Amini, N. Henry Riche, B. Lee, C. Hurter, and P. Irani. Understanding data videos: Looking at narrative visualization through the cinematography lens. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, 10 pages, p. 1459–1468. ACM, New York, 2015. doi: 10.1145/2702123.2702431
- [2] Armed Conflict Location & Event Data Project. Acled (armed conflict location & event data project). <https://acleddata.com/>. Last accessed: 2024-03-31.
- [3] F. Atamny, L. Eusebio, and W. Atamny. *Atlas of Palestinian Communities In Israel*. Fachri Atamny, 2017. With 202 maps and 10 information graphics. Linen binding.
- [4] B. Bach, E. Freeman, A. Abdul-Rahman, C. Turkay, S. Khan, Y. Fan, and M. Chen. Dashboard design patterns. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):342–352, 2023. doi: 10.1109/TVCG.2022.3209448
- [5] B. Bach, Z. Wang, M. Farinella, D. Murray-Rust, and N. Henry Riche. Design patterns for data comics. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 12 pages, p. 1–12. ACM, New York, 2018. doi: 10.1145/3173574.3173612
- [6] M. Behrisch, B. Bach, N. Henry Riche, T. Schreck, and J.-D. Fekete. Matrix reordering methods for table and network visualization. In *Computer Graphics Forum*, vol. 35, pp. 693–716, 2016. doi: 10.1111/cgf.12935
- [7] V. Braun and V. Clarke. *Thematic Analysis: A Practical Guide*. American Psychological Association, 2012.
- [8] A. Burns, C. Lee, R. Chawla, E. Peck, and N. Mahyar. Who do we mean when we talk about visualization novices? In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, article no. 819, 16 pages. ACM, New York, 2023. doi: 10.1145/3544548.3581524
- [9] A. Burns, C. Lee, T. On, C. Xiong, E. Peck, and N. Mahyar. From invisible to visible: Impacts of metadata in communicative data visualization. *IEEE Transactions on Visualization and Computer Graphics*, 30:3427–3443, 2024. doi: 10.1109/TVCG.2022.3231716
- [10] A. Bátorfy. Attila batorfy: Budapest open data atlasza, 2023. <https://atlo.team/boda/>, Last accessed: 2024-03-31.
- [11] J. Cheshire and O. Uberti. *Atlas of the Invisible: Maps & Graphics that Will Change how You See the World*. Particular Books, 2021.
- [12] R. Chetty, N. Hendren, J. Friedman, M. R. Jones, and S. R. Porter. The opportunity atlas: Mapping the childhood roots of social mobility. <https://www.opportunityatlas.org>. Last accessed: 2024-03-31.
- [13] M. Correll. Ethical dimensions of visualization research. 13 pages, p. 1–13. ACM, New York, 2019. doi: 10.1145/3290605.3300418
- [14] Datawheel. Pantheon. <https://pantheon.world/>, Last accessed: 2024-03-31.
- [15] Datawheel, Deloitte Global. Data usa. <https://datausa.io/>. Last accessed: 2024-03-31.
- [16] M. Dörk, S. Carpendale, and C. Williamson. The information flaneur: A fresh look at information seeking. In *Proceedings of the SIGCHI conference on human factors in computing systems*, 10 pages, pp. 1215–1224. ACM, New York, 2011. doi: 10.1145/1978942.1979124
- [17] M. Dörk, P. Feng, C. Collins, and S. Carpendale. Critical infovis: exploring the politics of visualization. In *CHI '13 Extended Abstracts on Human Factors in Computing Systems*, 10 pages, p. 2189–2198. ACM, New York, 2013. doi: 10.1145/2468356.2468739
- [18] C. d’Ignazio and L. F. Klein. Feminist data visualization. Workshop on Visualization for the Digital Humanities (VIS4DH), IEEE, Baltimore, 2016.
- [19] Environmental Defense Fund, Texas A&M University, Darkhorse Analytics. The u.s. climate vulnerability index. <https://climatevulnerabilityindex.org>. Last accessed: 2024-03-31.
- [20] S. J. Fernstad. To identify what is not there: A definition of missingness patterns and evaluation of missing value visualization. *Information Visualization*, 18(2):230–250, 2019. doi: 10.1177/1473871618785387
- [21] Google. Explore - google arts & culture, 2024. <https://artsandculture.google.com/explore>. Last accessed: 2024-03-31.
- [22] S. Hao, Z. Wang, B. Bach, and L. Pschetz. Design patterns for data-driven news articles. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, article no. 231. ACM, New York, 2024. doi: 10.1145/3613904.3641916
- [23] R. Hausmann, C. A. Hidalgo, S. Bustos, M. Coscia, S. Chung, J. Jimenez, A. Simoes, and M. A. Yildirim. *The Atlas of Economic Complexity*. MIT Press, Cambridge, MA, 2013.
- [24] H. A. He, J. Walny, S. Thoma, S. Carpendale, and W. Willett. Enthusiastic and grounded, avoidant and cautious: Understanding public receptivity to data and visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 30(1):1435–1445, 2024. doi: 10.1109/TVCG.2023.3326917
- [25] S. Huron, B. Bach, U. Hinrichs, M. Keck, and J. C. Roberts. Ieee vis workshop on data vis activities to facilitate learning, reflecting, discussing, and designing. In *IEEE VIS 2020*, 2020. doi: hal-02887742
- [26] IPCC. Ippc wgi interactive atlas. <https://interactive-atlas.ipcc.ch/>. Last accessed: 2024-03-31.
- [27] P. Isenberg, B. Lee, H. Qu, and M. Cordeil. *Immersive visual data stories*, pp. 165–184, 2018. doi: 10.1007/978-3-030-01388-2_6
- [28] K. Jackson and P. Bazeley. Nvivo 14, 2023. Version 14.
- [29] Y. Jansen, P. Dragicevic, P. Isenberg, J. Alexander, A. Karnik, J. Kildal, S. Subramanian, and K. Hornbæk. Opportunities and challenges for data physicalization. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, p. 3227–3236. ACM, New York, 2015. doi: 10.1145/2702123.2702180
- [30] S. Knudsen, J. Vermeulen, D. Kosminsky, J. Walny, M. West, C. Frisson, B. Adriel Aseniero, L. MacDonald Vermeulen, C. Perin, L. Quach, P. Buk, K. Tabuli, S. Chopra, W. Willett, and S. Carpendale. Democratizing open energy data for public discourse using visualization. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*, 4 pages, p. 1–4. ACM, New York, 2018. doi: 10.1145/3170427.3186539
- [31] X. Lan, Y. Wu, and N. Cao. Affective visualization design: Leveraging the emotional impact of data. *IEEE Transactions on Visualization and Computer Graphics*, 30(01):1–11, 2024. doi: 10.1109/TVCG.2023.3327385
- [32] L. North. *Scotland: An Atlas of Opportunity*. Common Weal, 2019. <https://lateralnorth.com/portfolio/an-atlas-of-opportunity/>. Last accessed: 2024-03-31.
- [33] U. N. D. of Economic and S. Affairs. *The Sustainable Development Goals Report 2023: Special Edition*. United Nations, New York, 2023. <https://unstats.un.org/sdgs/report/2023/>. Last accessed: 2024-03-13.
- [34] Organisation for Economic Co-operation and Development. Oced.ai policy observatory. <https://oecd.ai/en/>. Last accessed: 2024-03-31.
- [35] Organisation for Economic Co-operation and Development. Oecd better life index. <https://www.oecdbetterlifeindex.org/>. Last accessed: 2024-03-31.
- [36] L. Padilla, M. Kay, and J. Hullman. *Computational Statistics in Data Science*, chap. Uncertainty visualization. Wiley, 2020.
- [37] G. Panagiotidou, H. Lamqaddam, J. Poblome, K. Brosens, K. Verbert, and A. Vande Moere. Communicating uncertainty in digital humanities visualization research. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):635–645, 2023. doi: 10.1109/TVCG.2022.3209436
- [38] C. Perin, P. Dragicevic, and J.-D. Fekete. Revisiting bertin matrices: New interactions for crafting tabular visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):2082–2091, 2014. doi: 10.1109/TVCG.2014.2346279
- [39] A. F. Pirlea, U. Serajuddin, D. Wadhwa, and M. Welch. Atlas of sustainable development goals 2023. World Bank, 2023. <https://datatopics.worldbank.org/sdgateatlas/>. Last accessed: 2024-03-13.
- [40] A. F. Pirlea, U. Serajuddin, D. Wadhwa, M. Welch, and A. Whitby. *Atlas of the Sustainable Development Goals 2020: From World Development Indicators*. World Bank, Washington, DC, 2020. <https://datatopics.worldbank.org/sdgateatlas/>. Last accessed: 2024-03-13.
- [41] R. Risam. 35 intersectional feminist data visualization precepts and practices. *The Routledge Companion to Intersectionalities*, p. 28, 2023.
- [42] M. Roser. Our world in data. <https://ourworldindata.org/>, n.d. Accessed: 2024-02-20.
- [43] R. Roy and J. P. Warren. Card-based design tools: A review and analysis of 155 card decks for designers and designing. *Design Studies*, 63:125–154, 2019. doi: 10.1016/j.destud.2019.04.002
- [44] E. Segel and J. Heer. Narrative visualization: Telling stories with data. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1139–

1148. doi: 10.1109/TVCG.2010.179

- [45] F. Serra, V. Cavalloro, V. Giustetto, M. Parigini, M. Mauri, T. Elli, Briones, B. Gobbo, M. Pappalepore, and P. Ciuccarelli. Atlante calvino: Literature and visualization. <https://atlantecalvino.unige.ch>. Last accessed: 2024-03-31.
- [46] L. Shen, Y. Zhang, H. Zhang, and Y. Wang. Data player: Automatic generation of data videos with narration-animation interplay. *IEEE Transactions on Visualization and Computer Graphics*, 30(1):109–119, 2024. doi: 10.1109/TVCG.2023.3327197
- [47] X. Shu, A. Wu, J. Tang, B. Bach, Y. Wu, and H. Qu. What makes a data-gif understandable? *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1492–1502, 2021. doi: 10.1109/TVCG.2020.3030396
- [48] H. Song and D. A. Szafrir. Where’s my data? evaluating visualizations with missing data. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):914–924, 2019. doi: 10.1109/TVCG.2018.2864914
- [49] C. Stoiber, F. Grassinger, M. Pohl, H. Stitz, M. Streit, and W. Aigner. Visualization onboarding: Learning how to read and use visualizations. 2019. doi: 10.31219/osf.io/c38ab
- [50] The Growth Lab at Harvard University. The atlas of economic complexity. <http://www.atlas.cid.harvard.edu>. Last accessed: 2024-03-13.
- [51] The Jožef Stefan Institute. Rare diseases observatory, 2024. <https://rarediseases.ijs.si/index>. Last accessed: 2024-03-31.
- [52] Thomson Reuters. Reuters graphics, 2024. <https://www.reuters.com/graphics/>. Last accessed: 2024-03-31.
- [53] Trase. Trase: Mapping the childhood roots of social mobility, 2023. <https://trase.earth>. Last accessed: 2024-03-31.
- [54] N. Tuzcu, A. White, B. Leonard, and S. Geoffrey. Unraveling the complexity: A user-centered design process for narrative visualization. In *Extended Abstracts of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1–7. ACM, New York, 2023. doi: 10.1145/3544549.3573866
- [55] Uppsala Conflict Data Program. Ucdp conflict encyclopedia. <https://ucdp.uu.se>. Date of retrieval: 24/03/31.
- [56] T. Vancisin, L. Clarke, M. Orr, and U. Hinrichs. Provenance visualization: Tracing people, processes, and practices through a data-driven approach to provenance. *Digital Scholarship in the Humanities*, 38(3):1322–1339, 2023. doi: 10.1093/lc/fqad020
- [57] B. Victor. Explorable explanations, 2011. <http://worrydream.com/ExplorableExplanations>. Last accessed: 2024-03-13.
- [58] F. B. Viegas, M. Wattenberg, F. van Ham, J. Kriss, and M. McKeon. Manyeyes: a site for visualization at internet scale. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1121–1128, 2007. doi: 10.1109/TVCG.2007.70577
- [59] Z. Wang, H. Romat, F. Chevalier, N. H. Riche, D. Murray-Rust, and B. Bach. Interactive data comics. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):944–954, 2022. doi: 10.1109/TVCG.2021.3114849
- [60] World Bank. Gender data portal. <https://genderdata.worldbank.org/>, Last accessed: 2024-03-31.
- [61] World Bank. *Atlas of Sustainable Development Goals 2017: World Development Indicators*. World Bank, Washington, DC, 2017. <https://datatopics.worldbank.org/sdgatlas/archive/2017/>. Last accessed: 2024-03-13. doi: 10.1596/978-1-4648-1080-0
- [62] World Bank. *Atlas of Sustainable Development Goals 2018: From World Development Indicators*. World Bank, Washington, DC, 2018. <https://datatopics.worldbank.org/sdgatlas/archive/2018/>. Last accessed: 2024-03-13.
- [63] Z. Zhao and N. Elmqvist. Datatv: Streaming data videos for storytelling. *arXiv preprint*, 2022. doi: arXiv:2210.08175