

REMIXTAPE: Enriching Narratives about Metrics with Semantic Alignment and Contextual Recommendation

Matthew Brehmer*
University of Waterloo

Margaret Drouhard†
Salesforce

Arjun Srinivasan‡
Databricks

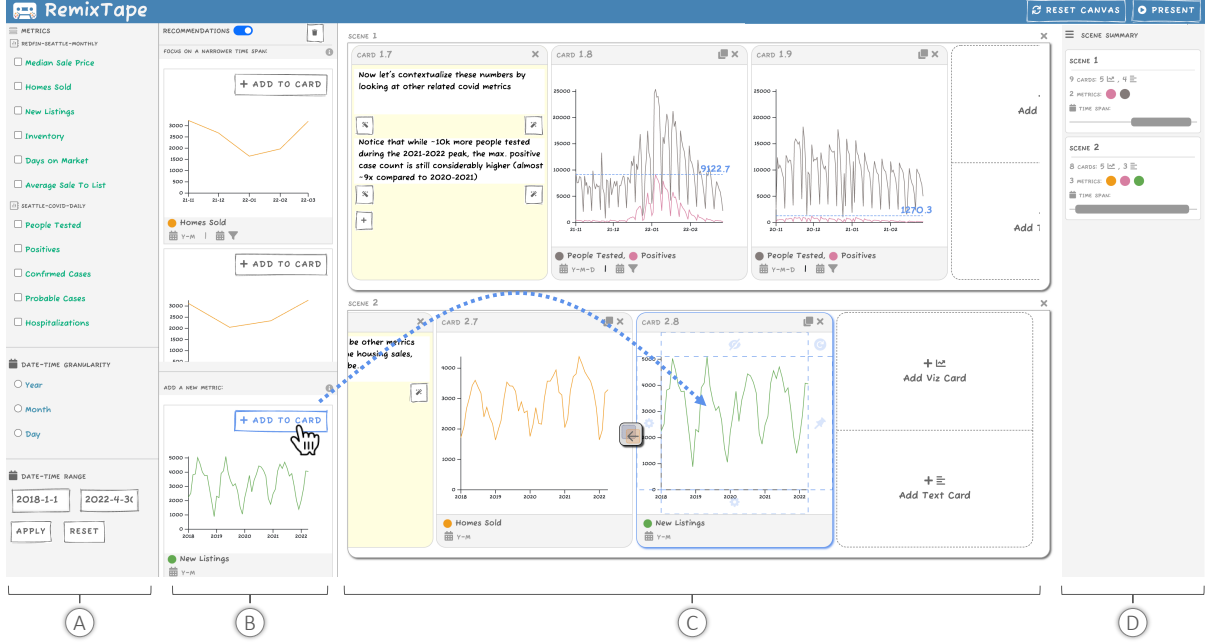


Figure 1: Upon adding a new VizCard (Card 2.8) to a scene in the authoring canvas (C), REMIXTAPE offers contextually-aware recommendations for populating the card (B) using available metrics (A). After selecting one of them (**# New Listings**), REMIXTAPE updates the scene summary (D) and reveals an affordance to semantically align this card with its neighbor (**# Homes Sold**).

ABSTRACT

The temporal dynamics of quantitative metrics or key performance indicators (KPIs) are central to conversations in enterprise organizations. Recently, major business intelligence providers have introduced new infrastructure for defining, sharing, and monitoring metric values. However, these values are often presented in isolation and appropriate context is seldom externalized. In this design study, we present REMIXTAPE, an application for constructing structured narratives around metrics. With design imperatives grounded in prior work and a formative interview study, REMIXTAPE provides a hierarchical canvas for collecting and coordinating sequences of line chart representations of metrics, along with the ability to externalize situational context around them. REMIXTAPE includes affordances to semantically align and annotate juxtaposed charts and text, as well as recommendations of complementary charts based on metrics already present on the canvas. We evaluated REMIXTAPE in a study in which six enterprise data professionals reproduced and extended partial narratives. They appreciated REMIXTAPE as a novel alternative to dashboards, galleries, and slide presentations for supporting conversations about metrics. Finally, we reflect on REMIXTAPE’s

usability and potential utility, and conclude with a call to define a conceptual foundation for remixing in the context of visualization.

Index Terms: Application / design study, business intelligence, metrics, key performance indicators / KPIs, narrative visualization.

1 INTRODUCTION

Conversations about data are ubiquitous in enterprise organizations [74]. Prior research [11] highlights the prevalence of time-oriented data in these conversations with recurring references to quantitative *metrics* (also commonly referred to as *key performance indicators* or KPIs). The temporal dynamics of metrics as depicted in line charts are central to the discussion of trends, the forming of predictions, and the setting of goals. The emergence of unified metrics platforms within enterprise business intelligence (BI) platforms [6, 22, 34, 75, 73] has made tracking and sharing metrics accessible to more people across an organization; metrics no longer need to be bound to dedicated analytical applications or individual dashboards. Despite these new abilities, enterprise communication processes still require ways to emphasize the relationships between metrics and the context around them.

We present REMIXTAPE, an application for enriching conversations around metrics. REMIXTAPE affords coordinating and annotating line chart representations of metrics and arranging them into compelling narrative presentations (Fig. 1). REMIXTAPE reconciles disparate temporal and quantitative value domains with affordances to semantically align juxtaposed metrics, thereby allowing metrics to contextualize one another. It also provides context-driven recommen-

*e-mail: mbrehmer@uwaterloo.ca

†e-mail: mar.drouhard@salesforce.com

‡e-mail: arjun.srinivasan@databricks.com

dations of metrics that complement partially-constructed narratives. Finally, it provides the means to externalize the situational context of metrics via annotation and text-chart synchronization.

The **contributions** of this paper are threefold. First, we draw attention to BI practices involving arranging, coordinating, annotating, and re-contextualizing metrics, reflected in findings from a formative interview study and a synthesis of prior research. Second, we describe REMIXTAPE as a novel combination of semantic alignment and context externalization with text-chart synchronization and chart recommendation. Finally, we interpret findings from a study with six enterprise data professionals, leading to a critical reflection on REMIXTAPE's interface design and plausibility.

2 BACKGROUND AND RELATED WORK

We situate our research within the context of enterprise communication around metrics. We then review prior work relevant to REMIXTAPE's key elements: view coordination, context externalization, chart recommendation, and visualizing time-oriented data.

2.1 Communicating Around Data in the Enterprise

Artifacts originating from a data analysis process may be combined and repurposed across an organization, either by the data analyst or data steward who originally created them, or by those in other roles [74]. Many of these assets are BI dashboards [63], and Tory et al. [74] characterized several related activities that take place to support the circulation of dashboard content, the most prevalent being '*slicing and dicing*' dashboard content, imposing a narrative, changing the representation of the content, and ultimately sharing it. Walchshofer et al. [77] further observed that different specialists undertake data-related activities within an organization, this despite the reality that most BI dashboards are often designed with a single persona in mind. Lee et al. [48] suggest that these unique roles include *analysts*, *editors*, and *presenters*, meaning that those who prepare presentations about data may be distinct from those who constructed the charts and dashboards used in the presentation [11]. The aim of REMIXTAPE is to support the latter role: those who re-purpose and recombine data assets for the purpose of communication.

Many presentations about data taking place within enterprise settings [11] manifest as what Segel and Heer [66] would characterize as *slide shows* of annotated charts. These sequences [37, 38] of charts and text are often arranged in such a way so as to produce a particular rhetorical effect [36], to form an argument [44], or to persuade the viewer to act [59]. There is a considerable breadth of narrative design patterns for presenting these sequences [4], including those specific to presenting time-oriented data [46], and we strived to ensure that this breadth of narrative structure is accessible with REMIXTAPE. Lastly, it is common practice [11] to use slide presentation software such as Microsoft PowerPoint to present sequences of static chart images. In contrast, Ellipsis [64], Timeline Storyteller [12], and Tableau's Story Points [30] have each demonstrated the capability to present sequences of interactive charts that can be dynamically augmented with annotations, a capability that we perpetuate with REMIXTAPE.

REMIXTAPE is also a response to the evolution of enterprise BI deployments, from dedicated applications and dashboards to headless BI and embedded experiences. Part of this shift is the adoption of unified metrics layers (or similarly-described offerings) from BI vendors, including Microsoft's Metrics for Power BI [6] and Tableau's Pulse [73]. While metrics can be monitored or shared individually, they can also be organized into folders, galleries, or collections [35]. Multiple metrics can also be packaged as dashboards, such as in Epperson et al.'s recent Quick Dashboard [28]. The latter is a departure from the convention of bottom-up BI dashboard construction, in that metric definition and individual chart specification are automated, allowing for a top-down dashboard specification. Like Quick Dashboard, REMIXTAPE also obviates the need to define

metrics or specify chart designs; it differs, however, in its reconciliation of metrics with different units and time periods, the arrangement of views into narrative sequences rather than dashboards, and its affordances for externalizing situational context.

2.2 Arranging and Coordinating Multiple Views

Central to REMIXTAPE is the arrangement of line chart representations of metrics through juxtaposition and superposition. While a slideshow is one way to present a sequence of views [12, 25, 33, 64], there can also be value in considering approaches that distribute views over space, whether in a grid of orthogonally aligned views [32], in a loosely structured canvas of data comics [5], or in a more structured dashboard [3, 25, 23, 52, 63]. In each case, the relative size, placement, and salience of views direct a viewer's attention. Another approach [19, 54], one exemplified in VizFlow [71], distributes views over both space and time in scrollable articles. We distribute views in REMIXTAPE across a structured canvas of scenes, so a series of views can be visible concurrently.

Regardless of whether a sequence of views are arranged across space, over time, or both, it is critical to coordinate these views [60] with consistent visual encoding choices, to clearly harmonize axes and units across adjacent views where appropriate, and to clearly indicate cases where view coordination is impossible or undesirable. It may also be possible to consolidate partially redundant juxtaposed views by superimposing them or by explicitly encoding the differences in their values [53]. While we exclude the derivation of new metrics in REMIXTAPE, as this is the responsibility of the data analyst or data steward, we incorporate functionality for harmonizing juxtaposed line charts as well as superimposing them when they are semantically related, functionality that can be described as semantic alignment. Kristiansen et al.'s characterization of *semantic snapping* [45] is a related concept, one that informs strategies to reduce redundancy and inconsistency across multi-view arrangements of charts, particularly in compact dashboards. However, in narrative presentations of data, compactness is less critical, and some degree of redundancy can be useful for emphasis and message retention.

2.3 Externalizing Situational Context

Many narrative visualization tools (e.g., [12, 26, 64, 30]) allow for the annotation of charts, however it is worth unpacking what annotation entails. Ren et al. [61] describe a design space of annotation as having two dimensions: annotation *forms* (including text, shape, highlight, and image) and annotation *targets* (including data marks, coordinate spaces, prior annotations, and other elements such as axes or legends). While tools like SmartCues [70] and ChartAccent [61] demonstrate coverage of this design space across different types of charts, we consider a subset of ephemeral and persistent annotations for line charts in REMIXTAPE, which include reference line shapes for indicating goals or thresholds and text value annotations for drawing attention to specific peaks, troughs, or inflection points.

The ability to annotate charts with open-ended text commentary allows people to externalize situational context and prior knowledge. Prior work demonstrates several points in the design space for this coordination, one that presents choices in such as the relative size and placement of text commentary and visualization elements as well as the level of interaction between these elements. In a data analysis context, it can be useful to collect externalizations in a prominent dedicated view while retaining an interactive link to their associated chart referents [43]. In a narrative visualization context, one approach involves interleaving text and visualization within an article [19], where visualization elements could appear, update, or disappear in response to the scroll position [54, 71]. Alternatively, Kori [47] and DataTales [72] augment text with interactive widgets that trigger corresponding highlights or updates in adjacent charts. Dashboards can blur the line between analytical and narrative experiences [56], particularly when they are populated with verbose

and responsive captions containing inline filter and parameter controls, in what could be referred to as a *narrative dashboard* [77]. REMIXTAPE similarly places an equal level of emphasis on chart and text elements, and we coordinate them such that a text element can interactively highlight a spatially adjacent line chart.

2.4 Visualization Recommendation

The need to re-purpose and *slice and dice* BI assets [74] implies a need to filter and align views from heterogeneous sources, as well as locate, browse, and identify relevant assets within shared repositories of dashboards, charts, datasets, and metrics. Altogether, these activities may be supported or accelerated by system guidance in the form of contextual recommendations. Prior research offers guidance strategies in visual analysis workflows [17, 14, 13], as well as approaches to visualization [49, 82, 86] and dashboard design [52, 58] recommendation. Line chart recommendations in REMIXTAPE are intended to surface complementary metrics and of logical sequences of exposition for ensuring a coherent narrative. Given this focus, we draw upon design ideas and techniques from visual data exploration recommendation systems that offer guidance by suggesting ‘*next steps*’ during data analysis (e.g., [31, 80, 41, 81, 9, 67]). However, recommendations in REMIXTAPE fundamentally differ from prior systems in that they focus on contributing to a narrative sequence, as opposed to promoting data coverage [80] or supporting an open-ended data analysis process [49].

2.5 Visualizing Metrics over Time

Although there are many ways to represent data over time [1], including those purpose-built for exploratory data analysis (e.g., [84, 83]) and monitoring (e.g., [55, 51]), the line chart is an established and familiar visual idiom, particularly for communicating metrics in a BI context [6, 73]. Apart from line charts, sparklines [76] and BANs [79] are also commonly used in BI contexts to show metric values compactly. However, given their prevalence in practice, REMIXTAPE exclusively represents metrics as line charts.

We observe two special considerations for line charts in REMIXTAPE. First, truncating the vertical axis of a line chart can be useful to draw a viewer’s attention to value differences over time, though without sufficient context or annotation, this choice could be deceptive [20]; REMIXTAPE leaves the choice to truncate the vertical axis up to the discretion of the user. A second consideration is whether to enable the creation of a dual-axis line chart, a convention used when one or more series of values exhibit a different value domain or unit of measure relative to what is already displayed in the chart. Following the empirical findings of Isenberg et al. [39] indicating that dual axis charts can elicit erroneous value judgments, REMIXTAPE restricts authors from creating them, opting instead for the creation of *indexed* line charts (e.g., [8]), relativizing values to a common percentage scale indexed to an initial value.

3 FORMATIVE INTERVIEW STUDY

To complement prior work in the BI domain (Section 2.1), we conducted a series of semi-structured 60-minute interviews to better understand why and how people adapt and recombine BI assets in enterprise contexts. We recruited six professionals who reported customizing, adapting, and repurposing BI assets to support sharing and storytelling across their organizations. The participants represented various sectors of the economy, including finance, higher education, technology, the automotive sector, and the nonprofit sector.

We began by asking participants about their role(s) and work related to BI assets: how they currently adapt, reuse, or share these assets, along with the challenges that arise when doing these activities. In the latter half of each interview, we used a dashboard containing sales metrics as a design probe and asked participants to explain how they might customize and annotate the content to share insights with others. Throughout the scenario, we asked participants

about opportunities for future tool support. The second author led the interviews and initial thematic analysis of the data collected findings, while the other authors assisted with a secondary focused analysis dedicated to the implications for metrics.

Collecting and repurposing charts. Interviewees described combining charts associated with different sources, such as disparate BI dashboards. However, this often entailed downloading raw data from a dashboard and painstakingly reproducing charts in spreadsheet and slide presentation tools in order to make necessary adaptations for a particular context; fittingly, a financial product sales analyst candidly remarked that this time-consuming duplication of work is ‘*stupid*’. A digital marketing producer in the higher education sector described the appeal of an “*additional workspace where I can rearrange and move stuff, [...] pulling stuff from different dashboards and bring it all into one place.*” However, the suggestion of such a workspace prompted a consultant from the automotive sector to insist that the original sources of charts should be clearly indicated. Moreover, participants indicated that such a workspace should allow people to interact with and make changes to the collected charts, suggesting that static screenshots of charts and dashboards are insufficient. According to the consultant, existing processes for embedding live charts from different sources are “*way too difficult.*”

Arranging and coordinating charts. The interviewer’s suggestion of a whiteboard workspace for collecting charts from disparate sources resonated with all of the interviewees. However, the free arrangement of charts across a digital whiteboard seemed ill-suited for some downstream activities, namely giving presentations and sharing reports, as these formats impose some degree of linear structure on the content they contain. The automotive industry consultant suggested that “*a storyboard kind of visual would help,*” as an alternative to a free-form whiteboard organizing structure. A linear structure reaffirmed the need for annotation, as captured by a project manager at a non-profit agency: “*I think just some commentary would be useful if you’re presenting this [content] to your team.*”

Adding context and directing attention. Assuming the possibility of collecting and consolidating charts within a common workspace, another aspect of customization described by interviewees involved adding annotations and enriching the charts with situational context. For instance, the sales analyst described the appeal of free-form annotation tools available when using screen-sharing features in telecommunication tools (e.g., Cisco WebEx Meetings [16]), and that such annotation tools could also be useful in a workspace of collected charts: “*you can use a little flashlight that highlights, and you can draw a box around something, you can draw lines along something, you can write your own text and basically annotate.*” Similarly, the digital marketing producer described a desire to add sticky notes to charts and dashboards, and she also described the ability to ‘*mark up*’ content and to toggle the visibility of this graphical annotation on demand. A financial analyst at a technology company also reminded us of how individual annotations may have varying levels of scope: “*annotations could mean data labels, data points, or it could mean simple one-line / two-line sentences.*”

Focusing on metrics. Participants described a recurring motivation to collect, arrange, and annotate charts depicting metrics, providing audiences with a unified context where they might compare and reason about values over time. For instance, the automotive industry consultant described that “*a lot of clients want to look at the change and the difference, [...] KPI [Key Performance Indicator] type of stuff.*” Similarly, the digital marketing producer described contextualizing a metric by providing year-over-year value comparisons.

Beyond providing the means to draw comparisons between known metrics of interest, the project manager described the potential value of ad-hoc trend comparison once a set of metrics are gathered. In her words, there would be value in seeing an additional “*trend or something that [she] wouldn’t even necessarily know to look for,*” In

other words, there is potential utility in surfacing recommendations that indicate relationships between metrics, recommendations that could complement a partially-constructed narrative.

Design imperatives. We distill our findings into design imperatives:

D1: Support arranging, slicing and dicing [74], and semantically aligning sequences of charts depicting metrics.

D2: Support the enrichment and annotation of metric chart sequences with situational context.

D3: Provide context-aware recommendations of complementary metric charts that enrich a partially-constructed narrative sequence.

4 REMIXTAPE

REMIXTAPE is an application motivated in part by our formative study and design imperatives **D1—D3**, in part by prior research involving communicating around data in enterprise BI settings [11, 40, 74], and in part by the BI industry’s recent investments in applications and services focusing on metrics (see Section 2.1).

Central to REMIXTAPE’s information architecture are collections of metrics as well as a hierarchical canvas of scenes and cards.

Collections of metrics. A BI metric is a univariate quantitative measure conventionally defined in reference to a temporal dimension and reported at a specified granularity (e.g., day, week, month, quarter). Metric values extend only as far back as to when the metric was defined by an analyst or data steward, and some metrics may no longer be accumulating values; consequently, different metrics will exhibit heterogeneous temporal domains (i.e., the time span between a metric’s starting value and its’ most recent value), as well as heterogeneous granularities. Anticipating this heterogeneity, REMIXTAPE is capable of ingesting multiple metric collections, with each metric containing a temporal attribute (📅), at least one metric (#), and (optionally) one or more categorical dimensions (A).

Canvas, scenes, and cards. Fig. 1 presents an overview of REMIXTAPE’s interface. The input panel (Fig. 1A) lists the available collections of metrics along with options to specify temporal granularities and filters. Within the main canvas (Fig. 1C), content is organized vertically in *scenes*, which are in turn composed of one or more text or visualization *cards* that are sequenced horizontally.


REMIXTAPE allows adding any number of scenes to the canvas, and these scenes can be reordered by dragging them within the canvas. It is similarly possible to append any number of visualization or text cards (VizCards or TextCards) to a scene, and cards can be reordered within a scene via dragging (Section 4.1). REMIXTAPE provides a graphical summary of the canvas (Fig. 1D), indicating each scene’s metrics and its chronological coverage.

VizCards contain a line chart displaying one or more metrics. While initially empty, VizCards can be populated either by manually selecting metrics from the input panel or through system-generated recommendations (Section 4.3). With either approach, it is possible to specify a desired temporal granularity (day, month, year) and constrain the selected metrics’ values to a specific time range (Fig. 1A, bottom). When combining metrics across datasets, REMIXTAPE implicitly finds common temporal granularities and collates time spans to render the appropriate line charts (e.g., showing a gap in the line chart if there are non-overlapping time periods).

TextCards provide text input affordances for adding paragraph-length commentary within a scene (**D2**). Each TextCard also provides affordances to add additional paragraphs as needed.

4.1 Arranging and Semantically Aligning Views (D1)

We offer several ways to arrange and coordinate VizCards.



Slice and dice. Users may want to emphasize a specific time period by only showing the specific time span in a VizCard. To this end, REMIXTAPE allows specifying






Figure 2: Examples of VizCard interaction affordances. (L) Clicking in the top margin provides options to filter, split, or obfuscate selected time spans. (R) Clicking in the right margin annotates the chart with a horizontal reference line to emphasize specific values.

temporal filters in the input panel before adding a VizCard to the canvas, as well as through interactions with a populated VizCard. Specifically, clicking in a VizCard’s top margin reveals controls for splitting the temporal domain of a chart (Fig. 2L), including retaining data points either before or after the split point (👉👈), or moving data points after the split point to a new VizCard (👉👈). Similarly, dragging parallel to the x-axis reveals controls to retain only the selected time span, or to exclude the selected span and retain the spans preceding and following it in two VizCards. Whenever a filter is applied to a VizCard, we annotate it with icons (📅👉👈) to indicate that the card only shows a subset of the data.



Truncate and relativize. Depending on the data context, one may want to show values along the Y-axis starting from the minimum value instead of zero (e.g., in a line chart depicting temperature). Alternatively, in cases such as stock market metrics, it may be useful for a line chart to depict changes relative to a specific point in time. REMIXTAPE allows clicking on either axis within a VizCard to truncate and relativize the value domains of the chart (Fig. 3C-right). Specifically, the domain of the vertical axis can be toggled between the minimum and maximum values within the card or between zero and the maximum value (📉📈). The vertical axis can also be indexed to a percentage scale relative to the first value appearing in the card (📊%). A relative transformation can also be applied to the horizontal axis (📅👉👈), from absolute chronological dates to relative dates beginning from the first date appearing in the card.



Juxtapose and coordinate. Using consistent value domains across charts can enhance sequential narration and aid comprehension [60]. When two VizCards are juxtaposed, REMIXTAPE can coordinate the axis domains of the right card with that of the preceding card by clicking on either axis in the right card and selecting the coordinate option (👉👈, Fig. 3C).



Merge and superimpose. While juxtaposing VizCards is one approach to comparing metrics, another approach is to combine them as a multi-series line chart in a single VizCard. Drawing upon the idea of semantic snapping [45], when two VizCards are juxtaposed, REMIXTAPE presents an option to merge them (👉👈) if they share a comparable value domain and time range (Fig. 3D). When the merge option is selected, REMIXTAPE replaces them with a new VizCard containing metrics from both cards, displaying the metrics’ values for the overlapping time span between the two cards.

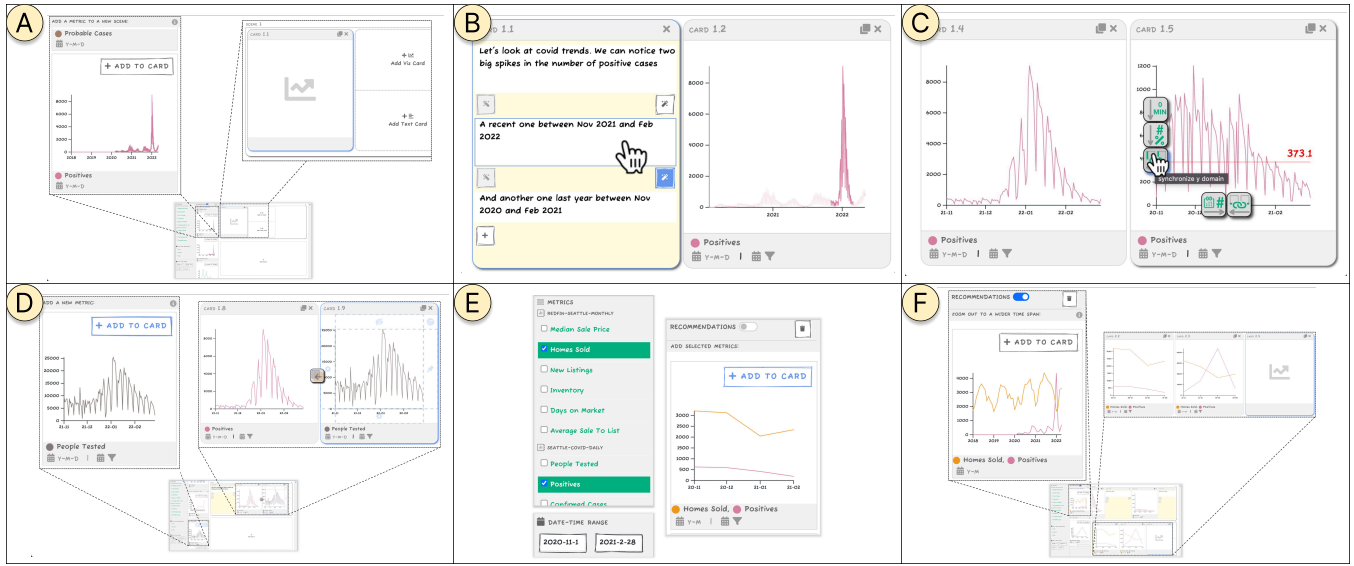



Figure 3: Using REMIXTAPE to construct a narrative about COVID-19 and housing sales. (A) Initial VizCards recommendations for a blank canvas. (B) Hovering over a TextCard element referencing Nov 2021—Feb 2022 ephemeral highlights the corresponding time span in a juxtaposed VizCard (D2). (C) Options for Arranging and Semantically Aligning Views (D1). (D) A recommended VizCard of the **# People Tested** metric exhibits a trend similar to one seen in the preceding VizCard (**# Positives**). An option to merge (↔) the two VizCards appears after adding the **# People Tested** VizCard to the scene. (E) A VizCard depicting two metrics (**# Homes Sold** and **# Positives**) appears when manually selecting metrics from the metric list. (F) A recommendation of an overview VizCard appears when a scene contains VizCards with filtered time spans (D3).

4.2 Adding Situational Context (D2)



Annotate and obfuscate. Annotations can draw the audience’s attention to specific time spans or values. REMIXTAPE supports adding both ephemeral or persistent annotations to VizCards. Ephemeral annotations are bound to the cursor position when mousing over a VizCard, revealing horizontal and vertical reference lines and corresponding date and metric value text annotations. A horizontal reference line and value annotation can also be persisted by clicking in a VizCard’s right margin (Fig. 2R); clicking again clears persistent annotations. Although additive annotations can reveal values, REMIXTAPE also provides support for dramatic subtractive reveals by adding a semi-opaque mask to obfuscate part of a chart, analogous to a fogged glass filter applied to obscure part of an image or a video (🔒 in Fig. 2L). As with persistent annotations, clicking in the card’s margin removes this obfuscatory mask, revealing the values beneath.



Coordinating VizCards & TextCards. When arranged in a sequence, we act on the opportunity to coordinate a TextCard with VizCards that precede or follow it. After commentary is added to a TextCard, each paragraph can be linked (🔗) to an adjacent VizCard (Fig. 3B), invoking the SUTime library [15] for recognizing and normalizing time expressions in the paragraph. Should a time expression intersect with the temporal domain of the linked VizCard, subsequent mouse hover events on that paragraph will ephemeral emphasize the time span mentioned in the paragraph. For instance, in Fig. 3B, the TextCard states “A recent one between Nov 2021 and Feb 2022.” Correspondingly, in the linked VizCard to the right, time periods before and after this span are de-emphasized until the mouse cursor departs from the paragraph.

After commentary is added to a TextCard, each paragraph can be linked (🔗) to an adjacent VizCard (Fig. 3B), invoking the SUTime library [15] for recognizing and normalizing time expressions in the paragraph. Should a time expression intersect with the temporal domain of the linked VizCard, subsequent mouse hover events on that paragraph will ephemeral emphasize the time span mentioned in the paragraph. For instance, in Fig. 3B, the TextCard states “A recent one between Nov 2021 and Feb 2022.” Correspondingly, in the linked VizCard to the right, time periods before and after this span are de-emphasized until the mouse cursor departs from the paragraph.

4.3 Recommending Charts from Context (D3)

Recommendations in REMIXTAPE are based on both the underlying data (e.g., trends, correlations) and user actions (e.g., metrics added to the canvas). Specifically, REMIXTAPE presents two types of recommendations whenever a VizCard is added to the canvas: **SEQUENTIAL** and **NEW METRIC** recommendations.

For either category of recommendation, REMIXTAPE presents a maximum of five recommendations by default, giving users the option to explicitly request additional recommendations through a “See More...” button. While recommendations are enabled by default, users can override them by manually specifying metrics, granularity, or temporal filters from the input panel (Fig. 1A). Users can also choose to disable recommendations altogether (Fig. 1B).

SEQUENTIAL recommendations. These recommendations retain the narrative’s focus on a single metric while introducing new perspectives on that metric. Specifically, when a VizCard is added to a partially-constructed scene, the system inspects the scene’s content to generate two types of recommendations: *drill-down* and *overview*.

Drill-down recommendations. REMIXTAPE recommends variations of charts already in a scene, adding temporal filters to drill down into prominent peaks and valleys. The “Focus on a narrower time span” recommendation in Fig. 1B is an example of the former: upon parsing the trend in **# Homes Sold** values in the preceding card, REMIXTAPE recommends a card focusing on the drop in values in late 2021 and another focusing on the sudden peak in mid-2019, among others. We employ a robust peak detection algorithm [10] that uses a moving average to detect peaks and valleys. The algorithm uses three input parameters: 1) *lag* controls the size of the moving window; 2) *threshold* (λ) determines the number of standard deviations (i.e., z-score); and 3) *influence* controls the degree to which new data points will affect the moving average and standard deviation (a value in the range of 0 – 1). Based on empirical testing across datasets and temporal granularities, we set these parameters to 5 (*lag*), 3.5 (*threshold*), and 0.5 (*influence*), respectively. For peak detection, at each time step t , a moving average μ_t and standard deviation σ_t are calculated using data within the moving window. A

data point is then considered to be a peak if its value is greater than $\mu_t + \lambda * \sigma_t$ or a valley if its value is less than $\mu_t - \lambda * \sigma_t$.

Overview recommendations. While *drill-down* recommendations follow the overview \rightarrow detail narration pattern, it is possible to arrange a scene with VizCards that show only a small span of the time period for which metric data is available. In these cases, there is an opportunity to provide missing context (i.e., a detail \rightarrow overview transition). Accordingly, REMIXTAPE recommends an overview chart spanning a wider time period (e.g., Fig. 3F).

NEW METRIC recommendations. Metrics can contextualize one another, and while some of these relationships are foreseeable when constructing a narrative, we anticipate that there may be value in suggesting relationships serendipitously. To this end, REMIXTAPE recommends charts featuring metrics that are not yet on the canvas.

Recommendations for a partially constructed scene. In cases where a new VizCard is added to a scene containing other populated VizCards, REMIXTAPE recommends charts of other metrics that complement those already in the scene. Specifically, REMIXTAPE iterates through available metrics and identifies those exhibiting trends that are similar to those in VizCards already in the scene. For instance, in Fig. 1, a chart of **# New Listings** is recommended because its trend is most similar to that of **# Homes Sold**, shown in the scene's penultimate card. To detect similarity, the system computes the Pearson's correlation coefficient (r) between a metric's values and those displayed in the preceding VizCard, recommending metrics with the highest $|r|$.

Recommendations for a new scene. If other scenes containing populated VizCards are present in the canvas, REMIXTAPE suggests broadening the narrative by recommending charts of unused metrics that exhibit a high variance over time.

Recommendations for a blank canvas. While we deem it unlikely that a user will not have a communication goal before using REMIXTAPE, it nevertheless provides recommendations in a cold start scenario. In these instances, REMIXTAPE iterates through available metrics and recommends charts that exhibit a potentially interesting trend (e.g., Fig. 3A) by computing the coefficient of variance (cv) for metrics' values and recommending those with high variability.

4.4 Usage Scenario

The following scenario reflects design imperatives **D1—D3** and is a composite of scenarios described by participants of the formative interview study. It centers on a fictional real estate professional named Sarah who will use REMIXTAPE to prepare a presentation for her colleagues about the dynamics of COVID-19 and the housing market in Seattle, Washington. Her organization's data stewards have defined a collection of metrics tracking Seattle's daily COVID-19 statistics since 2020, as well as another collection tracking Seattle's monthly real estate listings since 2018. Her narrative will reference a few metrics spanning both collections.

Sarah begins by adding a scene and a VizCard to the canvas. She decides that spikes in the number of COVID-19 **# Positives** (Fig. 3A) will establish a familiar context with her audience, so she populates her first VizCard with this metric.

Since **# Positives** does not have values prior to February 2020, Sarah filters the VizCard to focus on what has transpired since then (Fig. 2L). The cards show two prominent spikes in the data: a massive spike at the end of 2021 and a smaller one at the end of 2020. She fills a TextCard to draw attention to these spikes in and links these comments to the VizCard (Fig. 3B); this will highlight these spans when interacting with the TextCard (**D2**).

To draw attention to the two spikes, Sarah creates two duplicates of the VizCard showing **# Positives** and filters them to time periods of the two spikes. Noticing that their vertical axes are misaligned, she coordinates them (Fig. 3C, **D1**) to draw attention to the fact that the recent spike coinciding with New Year's 2022 is almost eight

times greater than that of the earlier spike, which coincided with the 2020 Thanksgiving holiday.

To add context to the two spikes, Sarah duplicates the VizCard showing the New Year's **# Positives** spike and considers REMIXTAPE's *new metric recommendations* (**D3**). She adds the **# People Tested** recommendation to her scene (Fig. 3D). She then merges the two cards and similarly creates another card depicting the dynamics of **# Positives** and **# People Tested** during the Thanksgiving 2020 spike. Next, she adds a TextCard to indicate that the rate of positives was much higher during the recent spike relative to during the previous spike. To emphasize this drastic difference in the number of **# Positives**, she adds reference line annotations orthogonal to the vertical axes (Fig. 2-right, **D2**).

Next, Sarah expands her narrative to include housing metrics. She adds a new scene and selects the **# Homes Sold** and **# Positives** metrics from the metric list in the input panel. To mirror the recent COVID spike featured in the first scene, she explicitly specifies November 2021 to February 2022 as her time period of interest. REMIXTAPE then presents a multi-series line chart that includes the two metrics, automatically aggregating the **# Positives** to match the monthly granularity of the **# Homes Sold** metric (Fig. 3E, **D1**). She adds this chart to the scene and similarly creates a second VizCard spanning the period corresponding to the earlier spike.

Sarah's narrative sequence now includes two VizCards in the second scene that depict a dip in **# Homes Sold** during the last two months of both 2020 and 2021, periods coinciding with rises in **# Positives**. As she adds a new card to contextualize these periods, REMIXTAPE observes that the cards in this scene show a narrow time range; accordingly, it provides an *overview recommendation* spanning the entire available time range for the housing metrics, from January 2018 until April 2022 (Fig. 3F, **D3**). Now her narrative communicates that seasonal drops in **# Homes Sold** in the fall were common even before the pandemic. She then appends a VizCard to the scene showing the values for the months before February 2020 and adds an explanation in a new TextCard.

To expand the narrative to other real estate metrics beyond sales, Sarah adds a new VizCard showing only the **# Home Sales** metric and peruses the *new metric recommendations* that she could add. The **# New Listings** exhibits a similar pattern to sales (Fig. 1, **D3**), so she creates a merged VizCard composed of **# Home Sales** and **# New Listings** (similar to Fig. 3D).

Satisfied with her two scenes that combine metrics from different sources that span partially overlapping time periods, Sarah feels confident in communicating her narrative. Later, in a meeting with her colleagues, she toggles REMIXTAPE's presentation mode as she presents her interpretation of the dynamics of Seattle's housing market before and during the COVID-19 pandemic (the toggle is located in the header at the top right of Fig. 1, and the presentation mode is demonstrated in the supplemental video).

5 USER STUDY

We observed enterprise professional participants perform two activities with REMIXTAPE reflecting **D1—D3**. We adopted a methodology of reproduction and open-ended construction activities established in prior evaluations of interactive construction applications in the visualization literature [2, 62]. Respecting our participants' valuable time, this methodology also allowed to elicit candid feedback pertaining to REMIXTAPE's potential utility, as well as ideas for extending and integrating it into existing workflows. Note that the goal of this study was not to generate objective evidence of utility or usability, but rather to present a qualitative reflection on REMIXTAPE's plausibility [57] informed by multiple subjective expert perspectives.

5.1 Participants

We recruited six professionals (P1—P6) distributed across North America and Europe who routinely communicate narratives concern-

ing data with their colleagues, stakeholders, and customers. They responded to our call for participation that we disseminated through mailing lists and Slack channels at a multinational software corporation and its subsidiary companies. None of these professionals had participated in our initial study described in Section 3. We screened respondents based on their form responses to questions regarding their prior experience with respect to communicating narratives grounded in data, with a focus on time-series metrics. All participants reported using business intelligence (BI) assets created by others, such as dashboards, metrics, or charts. Half of the participants regularly use or refer to content from five or fewer dashboards in the course of their work, while the other half regularly use content from more than ten. While they shared interests and experiences with respect to communicating narratives about metrics, they varied in terms of gender, years of professional experience, and job role (three held management roles, three held customer-facing roles).

5.2 Procedure

We conducted study sessions remotely via video conferencing. One author moderated the sessions while the other authors took notes. We recorded and transcribed each session, which included screen capture recordings of REMIXTAPE usage. Each session lasted between 75 and 90 minutes and included four phases, with the majority of this time split between the two hands-on activities.

Introduction. We began each session with a brief interview to understand current workflows with respect to metrics in communication contexts and any associated challenges. Next, we played a three minute-long video tutorial of REMIXTAPE to familiarize participants with its major features. We then gave participants an opportunity to ask clarifying questions and provided them with a user interface reference document that they could refer to during the activities.

Activity 1: Canvas reproduction. The first activity served as hands-on training as well as a preliminary assessment of REMIXTAPE's usability and learnability. We asked participants to screen-share their web browser and proceed to a URL where we had deployed an instance of REMIXTAPE populated with two collections of metrics: the monthly Seattle housing metrics introduced in Section 4.4 and a set of daily Seattle weather metrics, with both collections of metrics spanning January 2018 to June 2022. Next, we provided participants with a screenshot of a REMIXTAPE canvas populated with two scenes, each containing two VizCards and two TextCards: a sequence representing a partial narrative about the dynamics of Seattle's housing and weather. The activity required participants to reproduce the canvas as illustrated in the screenshot beginning from an empty canvas, which involved selecting metrics and specifying their temporal granularities, arranging and coordinating VizCards via filtering, relativizing, and merging them (D1), as well as annotating VizCards and adding TextCards (D2). This activity allowed flexibility in terms of the sequences of interactions leading to a correctly reproduced canvas. The only restriction that we imposed was the avoidance of the recommendation feature, which we set aside for the second activity. Finally, we encouraged participants to think aloud and indicate when they encountered difficulties with the interface so that we could resolve any confusion.

Activity 2: Open-ended construction. The second activity was a more free-form construction scenario [62]. While continuing to screen-share their web browser, we directed participants to another URL where we had deployed another instance of REMIXTAPE featuring the COVID-19 and housing metrics introduced in Section 4.4. Unlike in the first activity, we pre-populated the canvas with a scene containing two TextCards and three VizCards that served to introduce a narrative about two major spikes in COVID-19 cases. This activity required that participants either extend this initial narrative or establish their own narrative about the potential dynamics between these two sets of metrics. Also unlike the first activity, we

encouraged participants to make use of the recommendations feature as they completed the activity (D3). Otherwise, we similarly encouraged participants to think aloud as they populated the canvas with VizCards and TextCards.

Debrief. Finally, we asked participants to comment on how REMIXTAPE could be extended or integrated into their workflows.

5.3 Findings

The lead author performed an initial thematic analysis of our observations in reference to design imperatives D1—D3, later engaging the other authors in discussion and focused analysis.

Arranging sequences of cards (D1). Every participant remarked upon the unfamiliar scene- and card-based organization of REMIXTAPE's canvas, acknowledging that REMIXTAPE is neither a conventional BI dashboard interface nor a slide presentation tool. Anticipating the former, P2 described how the introduction of scenes and cards had them *"feeling a little lost"*, while by anticipating the latter, P5 felt similarly lost without the interface conventions of slide presentation tools. Looking past these differences, P6 appreciated our decision to place equal weight on VizCards and TextCards, allowing people to place them in any order in the canvas: *"The text just makes it more approachable and gives people the confidence to be really sure what they're supposed to be taking away."*


While the juxtaposition of TextCards and VizCards could help to keep teams *'on the same page'* [P6], P1 suggested adding a way to toggle between the display of a single VizCard and a juxtaposed display of multiple cards when using REMIXTAPE as a synchronous presentation tool. Similarly, P3 suggested collapsing scenes so that they could be revealed on demand, such as scenes prepared in anticipation of audience questions during a live presentation. Going a step further, P4 saw the selective reveal of charts as being critical for persuading the audience and guiding attention: *"if the idea is to persuade, then that would be the best way to do it because I can't control what [the audience is] viewing before I get there."* P4 also appreciated the affordances for arranging and coordinating VizCards in REMIXTAPE and envisioned selectively revealing content at finer levels of specificity, such as revealing individual metrics appearing in a multi-series VizCard.

P4 also anticipated situations where a presenter might limit the scope of where and how their audience could interact with charts appearing on a shared canvas, describing how an audience could *"click on a particular TextCard that exposes a VizCard once I've unlocked access to that, because I'm at that part of the scene,"* likening the process to the selective muting and unmuting of instrument channels in live musical performance software.

Scene and card hierarchies elicited several alternative suggestions with respect to arranging cards within a REMIXTAPE canvas. First, P2 questioned why we limited the interface to a two-level hierarchy and suggested the possibility of nesting cards or creating those with a focal chart with multiple peripheral charts, which could be useful for displaying a single aggregate metric alongside a categorical breakdown of values contributing to that aggregate. Second, P3 appreciated the logical separation of cards into scenes but found the vertical ordering of scenes to be limiting, describing use cases where two scenes each containing two cards are placed to be horizontally adjacent. P5 provided a similar suggestion of *"a vector canvas where you can just layer the cards [and] freely place them [...] or toggle between floating and tiled, hierarchical and linear."* A free-form canvas environment could reinforce the informality of some conversations around data, with P5 predicting that it would *"be really good for a rehearsal or dry run"* of a presentation.

To P2, the seemingly informal process of generating VizCards and arranging them on the canvas is one that could accelerate individual storyboarding workflows: *"VizCards are very cheap and I can spend the whole day creating and leaving them until I find something interesting, but I cannot do that in [existing*

tools].” However, when considering different audiences, contexts of use, and the lifespan of communicative artifacts, expectations regarding information architecture and required functionality may shift. In particular, to support ephemeral and synchronous conversations around metrics with colleagues, P2 remarked that the canvas / scene structure along with functionality for annotating and semantically aligning cards may not be warranted. In contrast, such functionality would be important to retain for longer-living communication artifacts consumed by audiences asynchronously.

Slicing and dicing cards (D1) and annotating them (D2). Participants’ reactions to the affordances for coordinating and arranging cards were mixed. While P5 felt that the “*The design mode works very fluidly*”, P2, P3, and P6 needed reminders of interactions invoked from a VizCard’s margins (Fig. 3C, Fig. 2), often right-clicking on the chart area to invoke a context menu, seeking options for filtering and adding annotations. Furthermore, when performing filter operations like card splitting, P6 failed to realize that a filter was applied to the VizCard (i.e., they did not notice the  status icons on the card). However, these concerns subsided as participants became more accustomed to the interactions during the activities. Citing tutorial design in video games, P5 suggested including in-situ onboarding features to aid discoverability, as these could reduce the initial learning curve; these features could manifest either as a stepped walkthrough of portions of the interface or via animated tooltips revealed when placing the cursor over specific elements.

All of the participants appreciated the temporal comparison REMIXTAPE supports by allowing them to spatially juxtapose, semantically align, and especially merge metrics into a single VizCard. However, it was not always clear to participants whether or how they could semantically align or merge two VizCards. For instance, four participants tried to merge VizCards that exhibited different time granularities or time spans; they wanted to align temporal patterns at different scales or across non-overlapping time periods, such as to highlight seasonal patterns. Upon realizing that merging VizCards was not possible in these cases, participants asked “*How would I allocate for this lag [between two metrics]? Can I offset [one metric]?*” [P4] or if the system could “*union them in a way*” [P5]. Such points of confusion with respect to semantic alignment interactions suggest that the system needs to provide more clarity, not only with respect to when VizCards can be merged, but also to explain cases where a merge is not feasible.

Irrespective of their experience with the card coordination and annotation interactions in REMIXTAPE’s design mode, all participants remarked positively on the ability to directly interact with content in the presentation mode. For example, P6 appreciated the “*dynamic nature*” of the presentation mode, providing affordances to tell “*more of a story than a static screenshot*.” P6 also commented that linking text with charts dynamically would support synchronous communication with audiences. However, P4 expressed a desire for additional control over text positioning and appearance (e.g., through tooltips or the progressive reveal of TextCards in presentation mode). Finally, P4 and P6 alluded to both presenters and audiences having the ability to interact with cards in REMIXTAPE’s presentation mode, which suggests that REMIXTAPE could be extended with real-time multi-user collaboration functionality.

Recommending contextually relevant VizCards (D3). The value proposition of recommendations varied based on prior knowledge of a dataset and whether one already has a comprehensive narrative in mind prior to using the tool, with P6 stating: “*I like the recommendation engine for people like me [who] are more business users. I don’t really know the story I’m trying to tell, and I don’t know tools as intimately as analysts do.*” P2 referred to recommendations as their “*favorite feature*” after seeing REMIXTAPE suggest charts across a collection of metrics spanning various temporal granularities. Participants also saw the potential value in recommendations as a form of guidance when initially forming a narrative, or as a source

of contextually-relevant ideas during narrative construction. For instance, P5 stated that “*the recommendations feature is interesting in that it’s really looking at what else is currently in the story.*”

In contrast, P3 questioned the role of recommendations when working with data that they were already familiar with. Given that REMIXTAPE “*feels more like a presentation tool*” and does not offer rich data exploration affordances (noting that “*this experience doesn’t scream data exploration to me*”), its recommendations may have limited value for those who already have a narrative in mind after conducting analysis with other tools.

Participants also commented on the types of recommendations (Section 4.3). P2 explained that in some instances, they care more about a particular metric and they will “*dig deep into it and explore it*”, suggesting the value of SEQUENTIAL recommendations and particularly those of the *drill-down* variety. In other instances, they might care more about the relationship between metrics, suggesting the value of NEW METRIC recommendations. Currently, the organization of scenes on the canvas affect the type of recommendations shown, but P2’s comments suggest an ability to take more explicit control over the type of recommendations generated by REMIXTAPE. This sentiment was echoed in another session, in which P3 sought a way to “*pre-filter*” the type of recommendations that REMIXTAPE would generate.

One potential concern pertains to the specificity of recommendations, as voiced by P4: “*Is it almost unethical to say ‘Hey, recommendation model, I’m trying to prove a correlation between ‘this’ and ‘this’? [...] Will [REMIXTAPE] populate recommendations that support my hypothesis?’*” Setting aside the issue of whether such a degree of specification is possible or ethical, P4 also suggested a potential mitigating strategy that could be applied to VizCards: the annotation of individual VizCards to indicate that they were either manually specified by the user or generated via REMIXTAPE’s recommendation engine. These annotations could serve to communicate the provenance of a narrative and establish a degree of trust in both the user who constructed the canvas as well as in the system.

6 DISCUSSION

We critically reflect on our design imperatives and consider alternatives suggested by our study findings, as well as on aspects of REMIXTAPE that could apply beyond metrics and BI.

6.1 Reflecting on Study Findings & Key Takeaways

Deviating from dashboards, slides, and galleries. In our formative interview study (Section 3), interviewees described a need to present and contrast metrics, echoing a similar need articulated in an earlier interview study by Brehmer and Kosara [11]. In their study, they documented the use case of revealing metric values in such a way that allows a presenter to make a time-over-time contrast, such as juxtaposing or superimposing an organization’s most recent quarterly sales performance against the same quarter of the previous year. While such a contrast could, in some instances, be shown within a single BI dashboard, Brehmer and Kosara’s interviewees remarked that dashboard applications were not well-suited for narrative presentations, and that attempting to use them in such a way gave the impression of a tool demonstration rather than as a medium for storytelling. With the advent of unified metrics layers (e.g., [6, 73]), metrics no longer need to be housed in individual dashboards, but can rather be aggregated into collections displayed as galleries [35]; however, such an interface is still better suited for browsing and monitoring. If dashboards and galleries are not effective presentation tools, people turn to slide presentations. We know from Tory et al. [74] that BI presentations often involve “*slicing and dicing*” charts and other BI assets and combining them into linear slide presentations, often as static screenshots. REMIXTAPE’s scrollable storyboard canvas therefore represents an alternative to

dashboards, galleries, and slide presentations for preparing and delivering narratives about metrics. However, as reported in Section 5.3, our alternative organizing structure was disorienting for several participants, suggesting either a new onboarding experience or an ability to toggle to more familiar or simplified modes of content organization. **Takeaway:** *Exercise caution when proposing alternative narrative BI experiences that deviate from the expected conventions of dashboard or slide presentation interfaces.*

Extending the scene + card canvas for narratives about metrics. We did not reject the conventions of slide presentation software altogether: cards in REMIXTAPE are analogous to individual slides, with each slide containing either an interactive line chart or an interactive block of text. While our study participants generally appreciated the simplicity of cards, the equal treatment of text and visualization, and the chunking of cards into scenes, their comments also suggest that we could relax our imposition of a linear order of cards. To this end, we could assign a level of importance to individual metrics. For example, in Elshehaly et al.'s QualDash application [27], a card-based dashboard specification involved assigning healthcare quality metrics to either a primary or secondary importance; while hidden by default, secondary supporting metrics in QualDash were revealed in an expanded card of a primary metric, an approach that could satisfy P3 and P4's related suggestions (Section 5.3). The flexible structure of comic strips [5, 85] is another (and possibly simpler) organizing scheme for metrics on a canvas. This approach is not incompatible with assigning importance levels to metrics, in that panels associated with secondary-class metrics could be smaller or less visually prominent, using other conventions for showing metrics beyond line charts, such as sparklines [76] or BANs [79]. **Takeaway:** *Consider scene + card-based presentations about metrics with cards of varying visibility and prominence.*

Rethinking content recommendations for BI presentations. Given our study findings, we acknowledge a potential mismatch between a desire for serendipitous content recommendations while formulating a narrative and REMIXTAPE's hierarchical scene + card-based canvas. These recommendations might be better suited when initially placing cards on an unbounded whiteboard canvas, as we suggested to participants of our formative interview study. Moreover, content recommendations might be altogether unnecessary if one already has a clearly-formulated narrative. Alternatively, recommendations could be offered *after* a narrative sequence is drafted: REMIXTAPE could recommend alternative card sequences exhibiting different rhetoric or argument structures [37, 46], or suggest perception-based intermediary VizCards that make for smoother transitions between points in a narrative [42, 50].

Another opportunity is recommending further synchronization between visualization and text. Currently, REMIXTAPE provides the ability for users to compose text and link this text to spatially adjacent VizCards, interactively highlighting portions of the latter when the text is brought into focus. Beyond this synchronization, there is an opportunity to leverage so-called 'auto-insight' features from recent work on natural language generation for describing charts [78, 68, 21]; for example, when adding a TextCard to the canvas, REMIXTAPE could generate an initial description of its adjacent VizCards, which the user could opt either to retain, edit, or extend. The converse is also possible [7]: for instance, when composing a TextCard with nouns (e.g., 'spike'), adjectives (e.g., 'dramatic'), verbs (e.g., 'falling'), or adverbs (e.g., 'gradually') that are typically associated with line charts, the system could recommend corresponding highlights and annotations on adjacent VizCards, or it could recommend the addition of other VizCards that depict similar patterns to those described in natural language. **Takeaway:** *Narrow the scope of content recommendations (e.g., alternative sequential structures, tighter text-chart alignment); otherwise, recommendation features give an impression that the application is intended for exploratory data analysis rather than for communicating narratives.*

6.2 Limitations & Future Work

Beyond BI and metrics. Our focus on metrics reflects the broader BI software community's recent investments in cloud infrastructure for shared metric repositories [6, 18, 22, 34, 75, 73]. Although we situate REMIXTAPE within an enterprise BI context [11, 40, 74], aspects of this project may be applicable in other domains where there is a need to construct and present narratives around data. STEM education or policy communication are two that come to mind, particularly as public and civic data portals provide opportunities to collect metrics originating from different institutions and agencies.

A design space for visualization remixing. The name REMIXTAPE is a portmanteau of *remix* and *mixtape*, with both terms being allusions to music production and dissemination. Additionally, the latter imparts a nuance of continuity: a sequential presentation of re-contextualized content. While the immediate purpose of REMIXTAPE is to support presentations about sequences of metrics, it nevertheless reflects a broader desire to *slice and dice* BI assets and enrich them with situational context, one emanating both from our own formative study (Section 3) and from interviews with enterprise professionals in earlier work [11, 74]. We also acknowledge that a need to remix the artifacts of knowledge work transcends the BI domain [29]. REMIXTAPE therefore represents a point in a design space for *visualization remixing*, and we call upon the research community to define a conceptual foundation for this design space and to differentiate it from remixing in other contexts.

6.3 Additional Evaluation Opportunities

Additional evaluation opportunities. Selecting an appropriate evaluation methodology for open-ended narrative visualization construction applications like REMIXTAPE can be challenging [62], particularly with many constructs to consider (e.g., learnability, usability, workflow integration, expressiveness [2]). Additionally, comparative studies are often logistically infeasible for these applications due to the confounding differences in feature sets and interaction paradigms [65]. Redesigning aspects of REMIXTAPE following the takeaway lessons of Section 6.1 and conducting a longitudinal post-deployment study centered around collections of metrics already familiar to participants will allow us to better understand REMIXTAPE's potential for communicating narratives about metrics and the role of these narratives in decision making [24].

7 CONCLUSION

The relationship between visualization research and business intelligence (BI) software has evolved over time, from the former's contributions to the conventions of interactive self-service visual analysis two decades ago [69] to the research community's renewed interest in BI dashboards since the end of the 2010s [3, 63, 74, 77], which in turn yielded new application-based research on approaches to dashboard construction (e.g., [52, 58]). As many dashboards could be simply described as containers for time-oriented metrics, our formative interview study highlights the frustration that enterprise professionals feel when attempting to collect, recontextualize, and present metrics from disparate dashboards. REMIXTAPE is a response both to these frustrations and to the emergence of shared metrics platforms within BI ecosystems [28, 18, 6, 22, 34, 75, 73], wherein sets of metrics may be associated with heterogeneous data sources and time frames. Lastly, REMIXTAPE represents a point in the nascent design space of visualization remixing, one that transcends time-oriented metrics and BI contexts, a design space that should be systematically explored and defined in future research.

ACKNOWLEDGMENTS

The authors conducted this research while at Tableau; we thank Zach Morrissey, Marlen Promann, Carolyn Tweedy.

REFERENCES

- [1] W. Aigner, S. Miksch, H. Schumann, and C. Tominski. *Visualization of Time-Oriented Data*. Springer, 2011.
- [2] F. Amini, M. Brehmer, G. Bolduan, C. Elmer, and B. Wiederkehr. Evaluating data-driven stories & storytelling tools. In N. Henry Riche, C. Hurter, N. Diakopoulos, and S. Carpendale, eds., *Data-Driven Storytelling*. A K Peters / CRC Press, 2018. <https://tinyurl.com/amini2018>.
- [3] B. Bach, E. Freeman, A. Abdul-Rahman, C. Turkay, S. Khan, Y. Fan, and M. Chen. Dashboard design patterns. *IEEE TVCG (Proc. VIS)*, 29(1), 2023. doi: 10.1109/TVCG.2022.3209448
- [4] B. Bach, M. Stefaner, J. Boy, S. Drucker, L. Bartram, J. Wood, P. Ciuccarelli, Y. Engelhardt, U. Koeppen, and B. Tversky. Narrative design patterns. In N. Henry Riche, C. Hurter, N. Diakopoulos, and S. Carpendale, eds., *Data-Driven Storytelling*. A K Peters / CRC Press, 2018.
- [5] B. Bach, Z. Wang, M. Farinella, D. Murray-Rust, and N. Henry Riche. Design patterns for data comics. In *Proc. ACM CHI*, 2018. doi: 10.1109/MCG.2017.33
- [6] P. Baumgartner. Goals in Power BI is becoming ‘metrics’ and is moving to GA!, 2022. Microsoft Power BI blog. <https://tinyurl.com/powerbimetrics>.
- [7] A. Bendeck, D. Bromley, and V. Setlur. SlopeSeeker: A search tool for exploring a dataset of quantifiable trends. In *Proc. ACM IUI*, 2024. doi: doi.org/10.1145/3640543.3645208
- [8] M. Bostock. Line chart, index chart, 2018. Observable notebook: <https://observablehq.com/@d3/index-chart>.
- [9] F. Bouali, A. Guettala, and G. Venturini. VizAssist: an interactive user assistant for visual data mining. *The Visual Computer*, 32(11), 2016. doi: 10.1007/s00371-015-1132-9
- [10] J. v. Brakel. Robust peak detection algorithm using z-scores. <https://stackoverflow.com/questions/22583391/peak-signal-detection-in-realtime-timeseries-data/22640362>, 2014.
- [11] M. Brehmer and R. Kosara. From jam session to recital: Synchronous communication and collaboration around data in organizations. *IEEE TVCG (Proc. VIS)*, 28(1), 2022. doi: 10.1109/TVCG.2021.3114760
- [12] M. Brehmer, B. Lee, N. Henry Riche, D. Tittsworth, K. Lytvynets, D. Edge, and C. White. Timeline Storyteller: The design & deployment of an interactive authoring tool for expressive timeline narratives. In *Proc. Computation + Journalism Symp. (C+J)*, 2019. <https://tinyurl.com/TSCJ2019>.
- [13] D. Ceneda, T. Gschwandtner, T. May, S. Miksch, H.-J. Schulz, M. Streit, and C. Tominski. Characterizing guidance in visual analytics. *IEEE TVCG (Proc. VAST)*, 23(1), 2017. doi: 10.1109/TVCG.2016.2598468
- [14] D. Ceneda, T. Gschwandtner, and S. Miksch. A review of guidance approaches in visual data analysis: A multifocal perspective. *CGF (Proc. EuroVis)*, 38(3), 2019. doi: 10.1111/cgf.13730
- [15] A. X. Chang and C. D. Manning. SUTime: A library for recognizing and normalizing time expressions. In *Proc. Language Resources and Evaluation (LREC)*, 2012.
- [16] Cisco. WebEx Meetings, 2022. <https://www.webex.com/video-conferencing>.
- [17] C. Collins, N. Andrienko, T. Schreck, J. Yang, J. Choo, U. Engelke, A. Jena, and T. Dwyer. Guidance in the human-machine analytics process. *Visual Informatics*, 2(3), 2018. doi: 10.1016/j.visinf.2018.09.003
- [18] R. Colvin and A. Kritzer. Introducing Metrics, headlines for your data, 2020. Tableau blog. <https://tinyurl.com/tableaumetrics>.
- [19] M. Conlen and J. Heer. Idyll: A markup language for authoring and publishing interactive articles on the web. In *Proc. ACM UIST*, 2018. doi: 10.1145/3242587.3242600
- [20] M. Correll, E. Bertini, and S. Franconeri. Truncating the y-axis: Threat or menace? In *Proc. ACM CHI*, 2020. doi: 10.1145/3313831.3376222
- [21] Z. Cui, S. K. Badam, M. A. Yalçın, and N. Elmqvist. Datasite: Proactive visual data exploration with computation of insight-based recommendations. *Information Visualization*, 18(2), 2019. doi: 10.1177/1473871618806555
- [22] databrain. Databrain: Metrics layer for the modern data stack, 2022. <https://www.usedatabrain.com/>.
- [23] D. Deng, W. Cui, X. Meng, M. Xu, Y. Liao, H. Zhang, and Y. Wu. Revisiting the design patterns of composite visualizations. *IEEE TVCG*, 29, 2023. doi: 10.1109/TVCG.2022.3213565
- [24] E. Dimara, H. Zhang, M. Tory, and S. Franconeri. The unmet data visualization needs of decision makers within organizations. *IEEE TVCG*, 28(12), 2022. doi: 10.1109/TVCG.2021.3074023
- [25] M. Elias, M.-A. Aufaure, and A. Bezerianos. Storytelling in visual analytics tools for business intelligence. In *Proc. IFIP Human-Computer Interaction*, 2013. doi: 10.1007/978-3-642-40477-1_18
- [26] M. Elias and A. Bezerianos. Annotating BI visualization dashboards: Needs & challenges. In *Proc. ACM CHI*, 2012. doi: 10.1145/2207676.2208288
- [27] M. Elshehaly, R. Randell, M. Brehmer, L. McVey, N. Alvarado, C. P. Gale, and R. A. Ruddell. Qualdash: Adaptable generation of visualisation dashboards for healthcare quality improvement. *IEEE TVCG (Proc. InfoVis)*, 27(2), 2021. doi: 10.1109/TVCG.2020.3030424
- [28] W. Epperson, K. Wongsuphasawat, A. Whilden, F. Du, and J. Talbot. A declarative specification for authoring metrics dashboards. In *Proc. IEEE Symp. on Visualization in Data Science (VDS)*, 2023. doi: 10.1109/VDS60365.2023.00006
- [29] J. L. Feuston and S. E. Lindley. How social dynamics and the context of digital content impact workplace remix. In *Proc. ACM CHI*, 2018. doi: 10.1145/3173574.3174171
- [30] E. Fields. Build narratives in your dashboards with Story Points in Tableau 8.2, 2014. Tableau blog. <https://tinyurl.com/story-points-2014>.
- [31] D. Gotz and Z. Wen. Behavior-driven visualization recommendation. In *Proc. ACM IUI*, 2009. doi: 10.1145/1502650.1502695
- [32] S. Gratzl, N. Gehlenborg, A. Lex, H. Pfister, and M. Streit. Domino: Extracting, comparing, and manipulating subsets across multiple tabular datasets. *IEEE TVCG (Proc. InfoVis)*, 20(12), 2014. doi: 10.1109/TVCG.2014.2346260
- [33] S. Gratzl, A. Lex, N. Gehlenborg, N. Cosgrove, and M. Streit. From visual exploration to storytelling and back again. *CGF (Proc. EuroVis)*, 35(3), 2016. doi: 10.1111/cgf.12925
- [34] T. Hardy. The next layer of the modern data stack, 2022. dbt blog. <https://tinyurl.com/dbtmetrics>.
- [35] A. Ho and M. Shulman. Tune in to use collections, playlists for your data, 2021. Tableau blog. <https://www.tableau.com/about/blog/2021/07/tune-use-collections-playlists-your-data>.
- [36] J. Hullman and N. Diakopoulos. Visualization rhetoric: Framing effects in narrative visualization. *IEEE TVCG (Proc. InfoVis)*, 17(12), 2011. doi: 10.1109/TVCG.2011.255
- [37] J. Hullman, S. Drucker, N. Henry Riche, B. Lee, D. Fisher, and E. Adar. A deeper understanding of sequence in narrative visualization. *IEEE TVCG (Proc. InfoVis)*, 19(12), 2013. doi: 10.1109/TVCG.2013.119
- [38] J. Hullman, R. Kosara, and H. Lam. Finding a clear path: Structuring strategies for visualization sequences. *CGF (Proc. EuroVis)*, 36(3), 2017. doi: 10.1111/cgf.13194
- [39] P. Isenberg, A. Bezerianos, P. Dragicevic, and J.-D. Fekete. A study on dual-scale data charts. *IEEE TVCG (Proc. InfoVis)*, 17(12), 2011. doi: 10.1109/TVCG.2011.160
- [40] S. Kandel, A. Paepcke, J. M. Hellerstein, and J. Heer. Enterprise data analysis and visualization: An interview study. *IEEE TVCG (Proc. VAST)*, 18(12), 2012. doi: 10.1109/TVCG.2012.219
- [41] A. Key, B. Howe, D. Perry, and C. Aragon. Vizdeck: self-organizing dashboards for visual analytics. In *Proc. ACM SIGMOD*, 2012. doi: 10.1145/2213836.2213931
- [42] Y. Kim, K. Wongsuphasawat, J. Hullman, and J. Heer. Graphscape: A model for automated reasoning about visualization similarity and sequencing. In *Proc. ACM CHI*, 2017. doi: 10.1145/3025453.3025866
- [43] Y.-S. Kim, N. Henry Riche, B. Lee, M. Brehmer, M. Pahud, K. Hinckley, and J. Hullman. Inking your insights: Investigating digital externalization behaviors during data analysis. In *Proc. ACM ISS*, 2019. doi: 10.1145/3343055.3359714
- [44] R. Kosara. An argument structure for data stories. In *Short Paper Proc. Eurographics / IEEE VGTC Visualization (EuroVis)*, 2017. doi: 10.2312/eurovisshort.20171129
- [45] Y. S. Kristiansen, L. Garrison, and S. Bruckner. Semantic snapping for guided multi-view visualization design. *IEEE TVCG (Proc. VIS)*, 28(1), 2022. doi: 10.1109/TVCG.2021.3114860
- [46] X. Lan, X. Xu, and N. Cao. Understanding narrative linearity for telling

- expressive time-oriented stories. In *Proc. ACM CHI*, 2021. doi: 10.1145/3411764.3445344
- [47] S. Latif, Z. Zhou, Y. Kim, F. Beck, and N. W. Kim. Kori: Interactive synthesis of text and charts in data documents. *IEEE TVCG (Proc. VIS)*, 28(1), 2022. doi: 10.1109/TVCG.2021.3114802
- [48] B. Lee, N. Henry Riche, P. Isenberg, and S. Carpendale. More than telling a story: Transforming data into visually shared stories. *IEEE CG&A*, 35(5), 2015. doi: 10.1109/MCG.2015.99
- [49] D. J.-L. Lee, V. Setlur, M. Tory, K. G. Karahalios, and A. Parameswaran. Deconstructing categorization in visualization recommendation: A taxonomy and comparative study. *IEEE TVCG*, 28(12), 2022. doi: 10.1109/TVCG.2021.3085751
- [50] H. Lin, D. Moritz, and J. Heer. Dziban: Balancing agency & automation in visualization design via anchored recommendations. In *Proc. ACM CHI*, 2020. doi: 10.1145/3313831.3376880
- [51] J. Lin, E. Keogh, S. Lonardi, J. P. Lankford, and D. M. Nystrom. Viztree: a tool for visually mining and monitoring massive time series databases. In *Proc. VLDB*, 2004. doi: 10.5555/1316689.1316811
- [52] Y. Lin, H. Li, A. Wu, Y. Wang, and H. Qu. DMiner: Dashboard design mining and recommendation. *IEEE TVCG*, 30(7), 2023. doi: 10.1109/TVCG.2023.3251344
- [53] S. LYi, J. Jo, and J. Seo. Comparative layouts revisited: Design space, guidelines, and future directions. *IEEE TVCG (Proc. VIS)*, 27(2), 2021. doi: 10.1109/TVCG.2020.3030419
- [54] S. McKenna, N. Henry Riche, B. Lee, J. Boy, and M. Meyer. Visual narrative flow: Exploring factors shaping data visualization story reading experiences. *CGF (Proc. EuroVis)*, 36(3), 2017. doi: 10.1111/cgf.13195
- [55] P. McLachlan, T. Munzner, E. Koutsofios, and S. North. LiveRAC: interactive visual exploration of system management time-series data. In *Proc. ACM CHI*, 2008. doi: 10.1145/1357054.1357286
- [56] E. Meeks. Third Wave Data Visualization, 2018. Tapestry Conf. keynote; <https://youtu.be/iChfcTx7ao>.
- [57] M. Meyer and J. Dykes. Criteria for rigor in visualization design study. *IEEE TVCG (Proc. InfoVis)*, 26(1), 2019. doi: 10.1109/TVCG.2019.2934539
- [58] A. Pandey, A. Srinivasan, and V. Setlur. Medley: Intent-based recommendations to support dashboard composition. *IEEE TVCG (Proc. VIS)*, 29(1), 2023. doi: 10.1109/TVCG.2022.3209421
- [59] A. V. Pandey, A. Manivannan, O. Nov, M. Satterthwaite, and E. Bertini. The persuasive power of data visualization. *IEEE TVCG (Proc. InfoVis)*, 20(12), 2014. doi: 10.1109/TVCG.2014.2346419
- [60] Z. Qu and J. Hullman. Keeping multiple views consistent: Constraints, validations, and exceptions in visualization authoring. *IEEE TVCG (Proc. InfoVis)*, 24(1), 2018. doi: 10.1109/TVCG.2017.2744198
- [61] D. Ren, M. Brehmer, B. Lee, T. Höllerer, and E. K. Choe. ChartAccent: Annotation for data-driven storytelling. In *Proc. IEEE PacificVis*, 2017. doi: 10.1109/PACIFICVIS.2017.8031599
- [62] D. Ren, B. Lee, M. Brehmer, and N. Henry Riche. Reflecting on the evaluation of visualization authoring systems. In *Proc. BELIV Workshop*, 2018. doi: 10.1109/BELIV.2018.8634297
- [63] A. Sarikaya, M. Correll, L. Bartram, M. Tory, and D. Fisher. What do we talk about when we talk about dashboards? *IEEE TVCG (Proc. InfoVis)*, 25(1), 2018. doi: 10.1109/TVCG.2018.2864903
- [64] A. Satyanarayan and J. Heer. Authoring narrative visualizations with Ellipsis. *CGF (Proc. EuroVis)*, 33(3), 2014. doi: 10.1111/cgf.12392
- [65] A. Satyanarayan, B. Lee, D. Ren, J. Heer, J. Stasko, J. Thompson, M. Brehmer, and Z. Liu. Critical reflections on visualization authoring systems. *IEEE TVCG (Proc. InfoVis)*, 26(1), 2020. doi: 10.1109/TVCG.2019.2934281
- [66] E. Segel and J. Heer. Narrative visualization: Telling stories with data. *IEEE TVCG (Proc. InfoVis)*, 16(6), 2010. doi: 10.1109/TVCG.2010.179
- [67] T. Siddiqui, A. Kim, J. Lee, K. Karahalios, and A. Parameswaran. Effortless data exploration with zenvisage: an expressive and interactive visual analytics system. *Proc. VLDB*, 10(4), 2016. doi: 10.14778/3025111.3025126
- [68] A. Srinivasan, S. M. Drucker, A. Endert, and J. Stasko. Augmenting visualizations with interactive data facts to facilitate interpretation and communication. *IEEE TVCG (Proc. InfoVis)*, 25(1), 2019. doi: 10.1109/TVCG.2018.2865145
- [69] C. Stolte, D. Tang, and P. Hanrahan. Polaris: A system for query, analysis, and visualization of multidimensional relational databases. *IEEE TVCG*, 8(1), 2002. doi: 10.1109/2945.981851
- [70] H. Subramonyam and E. Adar. SmartCues: A multitouch query approach for details-on-demand through dynamically computed overlays. *IEEE TVCG (Proc. InfoVis)*, 25(1), 2018. doi: 10.1109/TVCG.2018.2865231
- [71] N. Sultanum, F. Chevalier, Z. Bylinskii, and Z. Liu. Leveraging text-chart links to support authoring of data-driven articles with VizFlow. In *Proc. ACM CHI*, 2021. doi: 10.1145/3411764.3445354
- [72] N. Sultanum and A. Srinivasan. DATATALES: Investigating the use of large language models for authoring data-driven articles. In *Short Paper Proc. IEEE Visualization Conf. (VIS)*, 2023. doi: 10.1109/VIS4172.2023.00055
- [73] Tableau. Pulse, 2024. <https://www.tableau.com/products/tableau-pulse>.
- [74] M. Tory, L. Bartram, B. Fiore-Gartland, and A. Crisan. Finding their data voice: Practices and challenges of dashboard users. *IEEE CG&A*, 43(1), 2023. doi: 10.1109/MCG.2021.3136545
- [75] Transform. Transform data: Official metrics store for the modern data stack, 2022. <https://transform.co/>.
- [76] E. R. Tufte. Sparklines. In *Beautiful Evidence*. Graphics Press, 2006.
- [77] C. Walchshofer, V. Dhanoa, M. Streit, and M. Meyer. Transitioning to a commercial dashboarding system: Socio-technical observations and opportunities. *IEEE TVCG (Proc. VIS)*, 30(1), 2024. doi: 10.1109/TVCG.2023.3326525
- [78] Y. Wang, Z. Sun, H. Zhang, W. Cui, K. Xu, X. Ma, and D. Zhang. Datashot: Automatic generation of fact sheets from tabular data. *IEEE TVCG (Proc. VIS)*, 26(1), 2020. doi: 10.1109/TVCG.2019.2934398
- [79] S. Wexler. In Praise of BANs (Big-Ass Numbers), 2017. <https://datarevelations.com/bans/>.
- [80] K. Wongsuphasawat, D. Moritz, A. Anand, J. Mackinlay, B. Howe, and J. Heer. Voyager: Exploratory analysis via faceted browsing of visualization recommendations. *IEEE TVCG (Proc. InfoVis)*, 22(1), 2016. doi: 10.1109/TVCG.2015.2467191
- [81] K. Wongsuphasawat, Z. Qu, D. Moritz, R. Chang, F. Ouk, A. Anand, J. Mackinlay, B. Howe, and J. Heer. Voyager 2: Augmenting visual analysis with partial view specifications. In *Proc. ACM CHI*, 2017. doi: 10.1145/3025453.3025768
- [82] A. Wu, Y. Wang, X. Shu, D. Moritz, W. Cui, H. Zhang, D. Zhang, and H. Qu. AI4VIS: Survey on artificial intelligence approaches for data visualization. *IEEE TVCG*, 28(12), 2022. doi: 10.1109/TVCG.2021.3099002
- [83] J. Zhao, F. Chevalier, and R. Balakrishnan. Kronominer: using multi-foci navigation for the visual exploration of time-series data. In *Proc. ACM CHI*, 2011. doi: 10.1145/1978942.1979195
- [84] J. Zhao, F. Chevalier, E. Pietriga, and R. Balakrishnan. Exploratory analysis of time-series with ChronoLenses. *IEEE TVCG (Proc. InfoVis)*, 17(12):2422–2431, 2011. doi: 10.1109/TVCG.2011.195
- [85] Z. Zhao, R. Marr, J. Shaffer, and N. Elmqvist. Understanding partitioning and sequence in data-driven storytelling. In *Information (iConference)*, 2019. doi: 10.1007/978-3-030-15742-5_32
- [86] S. Zhu, G. Sun, Q. Jiang, M. Zha, and R. Liang. A survey on automatic infographics and visualization recommendations. *Visual Informatics*, 4(3), 2020. doi: 10.1016/j.visinf.2020.07.002