

Transitioning to a Commercial Dashboarding System: Socio-technical Observations and Opportunities

Conny Walchshofer , Vaishal Dhanoa , Marc Streit , and Miriah Meyer 

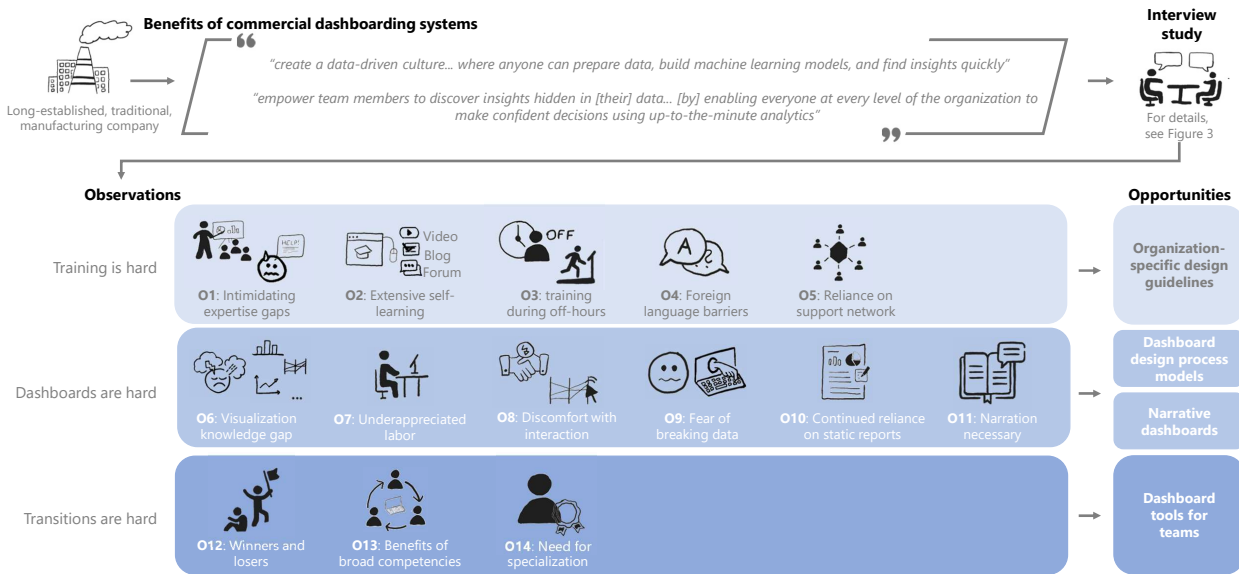


Fig. 1: A summary of the observations and opportunities from our study into the socio-technical challenges of employees at a large, conventional company as they transitioned to using Power BI.

Abstract—Many long-established, traditional manufacturing businesses are becoming more digital and data-driven to improve their production. These companies are embracing visual analytics in these transitions through their adoption of commercial dashboarding systems. Although a number of studies have looked at the technical challenges of adopting these systems, very few have focused on the socio-technical issues that arise. In this paper, we report on the results of an interview study with 17 participants working in a range of roles at a long-established, traditional manufacturing company as they adopted Microsoft Power BI. The results highlight a number of socio-technical challenges the employees faced, including difficulties in training, using and creating dashboards, and transitioning to a modern digital company. Based on these results, we propose a number of opportunities for both companies and visualization researchers to improve these difficult transitions, as well as opportunities for rethinking how we design dashboarding systems for real-world use.

Index Terms—Interview study, socio-technical challenges, visual analytics

1 INTRODUCTION

The success of commercial dashboarding systems like Tableau and Microsoft Power BI has made visual analytics (VA) an increasingly widespread practice. Using these tools, companies and organizations are able to standardize their data analysis processes, integrate their data sources, and use data throughout their decision-making. Visualization and HCI studies have documented the ways that these – and other – data analysis tools and processes have shaped the emergence of data science practices within organizations [3, 10, 18, 31, 32], as well as data-driven decision-making [22, 36]. These studies provide a picture of how

VA processes work in-the-wild and at scale, and elucidate technical challenges that remain [11, 25, 58].

Transitioning to a commercial dashboarding system across an organization is a highly disruptive process, replacing siloed data practices and legacy applications with a unified, self-service model. Research into the design of VA tools, however, has proceeded without explicit consideration of the challenges involved with transitioning to such a system, focusing largely on the technical issues while ignoring the organizational and cultural impacts [19]. What are the challenges workers face when transitioning to a commercial dashboarding system when they are not trained data scientists or visualization experts? What happens to their job responsibilities when long-established work practices drastically change? What support do workers need during the transition? And what opportunities are there for the visualization community to help more people embrace visual analytics?

In this work, we look specifically into the socio-technical effects of transitioning to Power BI in a large, conventional manufacturing company. Members of our team were initially invited to collaborate with an R&D department of the company to develop new, bespoke VA tools to optimize their production. During the collaboration, however, the company decided to adopt Power BI as the standard dashboarding solution across large parts of the organization. This decision provided

- Conny Walchshofer, Johannes Kepler University Linz, conny.walchshofer@jku.at.
- Vaishal Dhanoa, Pro²Future GmbH, vaishali.dhanoa@pro2future.at.
- Marc Streit, Johannes Kepler University Linz, marc.streit@jku.at.
- Miriah Meyer, Linköping University, miriah.meyer@liu.se.

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

us an opportunity to observe the disruptive effects of a commercial dashboarding system first-hand as employees encountered both organizational and cultural challenges. Some of these challenges stem from known issues of data and visualization literacy, but others point to opportunities to rethink the practices we design VA tools to support.

More specifically, we report on the results of an interview study with 17 employees at the company as they transitioned their work practices to use Power BI, summarized in Figure 1. These participants span a range of roles, backgrounds, and skills, bringing forward the experiences of both managers and workers. Our observations complement existing interview studies to bring a new socio-technical perspective to the challenges people face when adopting VA processes. These challenges highlight organizational and cultural issues related to training, effectively using dashboarding systems, and transitioning to new work practices. Based on these observations, we present a number of opportunities for visualization researchers to help data workers embrace VA practices through the development of guidelines and processes for creating dashboards. The observations also point to opportunities for rethinking the types of work that dashboarding systems are meant to support, specifically pointing to the need for more focus on supporting collaborative practices.

2 RELATED WORK

In this section, we briefly summarize the current state of commercial dashboarding systems, as well as the body of work that has studied how data scientists and data workers use those systems.

2.1 Commercial Dashboarding Systems

Until recently, data analysis was almost exclusively carried out by IT professionals and data analysts with specialized skills and knowledge [31]. The increasing need for making informed decisions based on large and complex data, however, has brought interactive VA solutions to the mainstream [61]. The visualization research community has invested significant efforts in understanding how to design VA tools for a wide variety of fields, often through the use of design study [48]. The result is a large number of papers reporting on sophisticated VA tools typically designed for a small group of expert users addressing specific domain problems.

Complementing these specific, sophisticated tools, the research community has also produced knowledge in the engineering of more general systems. Several notable commercial VA systems build off this early academic work: Tableau, stemming from the Polaris system [54]; Spotfire [1]; and Advizor [23]. The success of these systems led other technology companies to invest in building commercial dashboarding systems, including Microsoft, QlikView, Jaspersoft, and SAP.

The use of commercial dashboarding systems is now widespread. The ability of these systems to seamlessly work with other business-relevant systems enables businesses to cover the whole data analysis pipeline, ranging from data preparation to communication [18]. They are marketed as easy-to-use, integrated systems that reduce siloed analysis processes and standardize data practices, targeting a range of workers from analysts to upper management [13]. According to the annual study conducted by Gartner, Inc. (Gartner Magic Quadrant), three analytics and business intelligence platforms currently lead the market: Power BI by Microsoft, Tableau by Salesforce, and Qlik.

2.2 Empirical Studies of Data Analysis in the Wild

Over the last decade, both the visualization and HCI research communities have studied the emerging and evolving work of data scientists, analysts, and workers in a range of organizations and settings. These studies look at the data-driven practices and challenges of people using a multitude of analytic software tools and also characterize the technical work that people do. Across this work, a picture of how data science and VA are practiced in the world emerges, offering new opportunities for VA research and tool design.

Some of the earliest interview studies looking explicitly at data science practices provide initial characterizations of the data analysis process and work responsibilities of analysts. Kandel et al. [31] interviewed 35 analysts from 25 different organizations, documenting the

high-level analysis tasks they performed, the pain points they encountered, and the roles they took on. In a complementary study at a large software company, Fisher et al. [25] interviewed newly designated data scientists working within software development teams. Their findings focus on the challenges the data scientists faced with existing tools and infrastructure, and they call for more attention to these emerging practices by the HCI and UX communities. Kim et al. [32] similarly interviewed data scientists working at a large software company, describing how data-driven insights and decisions were being embraced by software development teams.

More recent studies build off these initial findings, deepening characterizations of how data scientists work and the specific roles they fill. Batch & Elmquist report on a contextual inquiry with analysts working at a large, US, federal agency [11]. Their findings both extend the characterization of analysts' roles and highlight the limited use of visualization throughout the data analysis pipeline. An interview study with professional analysts by Alspaugh et al. [3] similarly reports on the limited use of visualization, while also extending the data analysis process model to include exploration. Through a review of VA and data science literature, Crisan et al. further add pedagogy and collaboration to the data analysis process model and offer a rich and nuanced characterization of data worker roles [18].

Other recent interview studies look more specifically at how data scientists make use of visualization within human-in-the-automation-loop processes [17], the challenges that data workers and dashboard users encounter [58], and how analysts handle alternatives in their analyses [35]. Studies by Liu et al. and Dimara et al. focus on decision-makers within organizations, looking at replication issues within data-driven decision-making practices [36], as well as opportunities for visualization outside of data analysis [22]. Bringing a constructivist lens to their analysis of interviews with data scientists at a large technology company, Muller et al. [38] describe the ways that human knowledge and interpretation within data analysis practices challenge normative views of data-driven decision-making. Similarly, Bartram et al. challenge the research community's focus on sophisticated visualization tools through their study of data workers that highlights the ubiquity of spreadsheets in data analysis practices [10]. A recent study by Zhang et al. [62] highlights socio-technical challenges faced by dashboard creators as they grappled with communicating data about COVID-19 during the pandemic.

Looking more specifically at dashboards themselves, a pair of recent studies analyze dashboards collected from-the-wild to better understand emergent design patterns. The survey by Sarikaya et al. [46] defines a design space for dashboards that goes beyond visual encodings to include other dimensions such as purpose, audience, and data semantics. Building on this work, the survey by Bach et al. [6] resulted in 42 design patterns and 6 unique genres for dashboards. They validate the generative potential of the patterns and genres in a two-week workshop.

These previous studies all focus on *how* data analysis works in-the-wild, who is doing that work, the tools that they use, and the challenges that remain. A few existing, empirical studies have systematically looked instead at factors for the adoption of visual data analysis tools and software systems. Sedlmair et al. [47] reflect on their experiences developing visualization tools for a large automotive company, and discuss several insights for supporting the adoption of new tools. And a pair of surveys look at factors that lead to decisions to adopt commercial dashboarding systems, one based on reviews of leading tools [19], and another based on surveys sent to retailers in different countries [60].

In contrast to this body of literature, the empirical study we present in this paper looks at the disruptive effects of *transitioning* to a commercial dashboarding system. We take a socio-technical lens within our study and focus on how decisions to adopt a new data analysis pipeline across a diverse set of workers disrupt existing work practices and job responsibilities, to the benefit of some and not others. While most previous studies look at similar roles across many organizations, our study presents the effects of transitioning to Power BI across a range of jobs, skills, expertise, and roles within a single, large, conventional manufacturing company to present a picture of how the transition can disrupt and challenge an established organization. We present our findings in Section 4, and a set of opportunities in Section 5 visualization

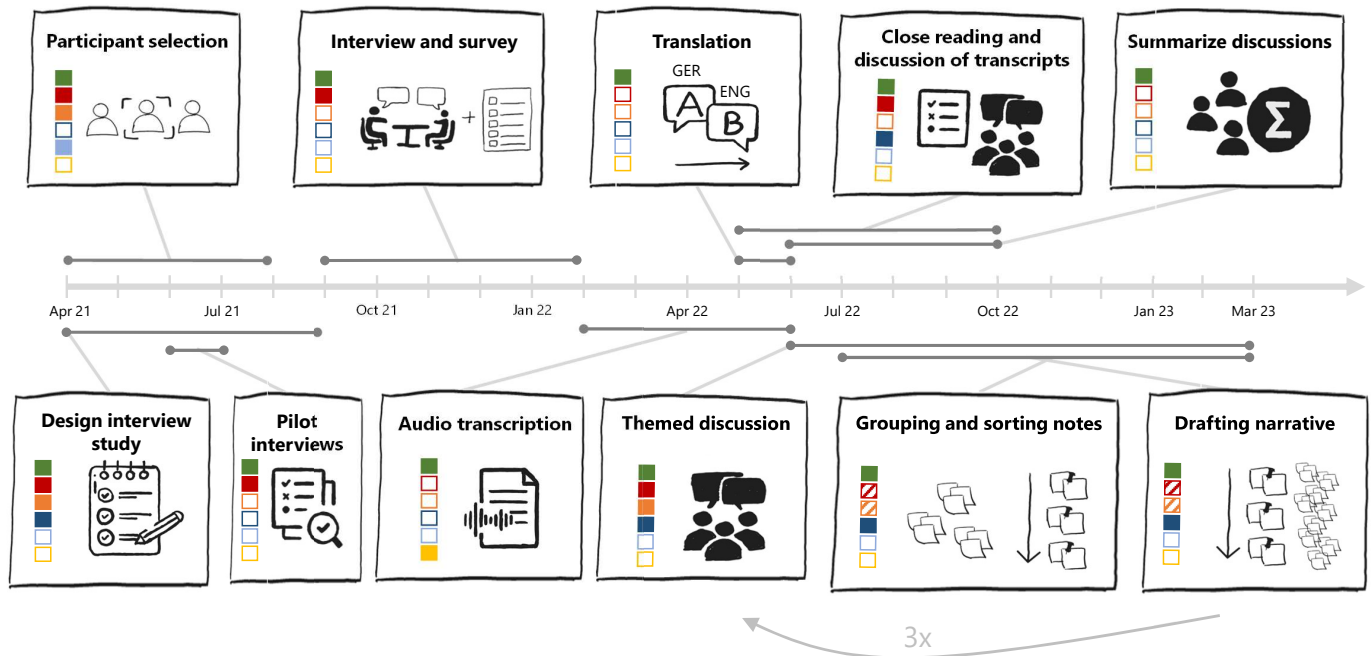


Fig. 2: The interview study described in this paper was conducted over the span of approximately two years. The colored boxes represent the involvement of the respective individuals: ■ Author 1 ■ Author 2 ■ Author 3 ■ Author 4 ■ Collaborator ■ Student. Non-filled boxes represent no participation, and shaded boxes represent a partial involvement.

research to support data workers as they adopt VA tools.

3 METHODS

In 2017, members of our team were invited to a multi-year collaboration with the R&D department of a large, traditional manufacturing company, which we will refer to as TheCompany throughout this paper. The initial purpose of this collaboration was to develop novel VA solutions for optimizing their production processes. We began this work by first exploring how our collaborators might use several tools we developed in previous projects [27, 56]. We also compared our own solutions to off-the-shelf alternatives such as Excel, Power BI, and Tableau for their particular needs. This initial work, however, elucidated the need for better guidance, methods, and processes for onboarding employees to new VA tools. We pivoted to explore onboarding within TheCompany in several follow-up studies [21, 52]. We worked closely with a diverse group of employees throughout the collaboration, getting to know them and their work practices, as well as developing a better understanding of the data analysis pipeline(s) within the organization.

During this collaboration, TheCompany decided to deploy Power BI across large parts of the organization. This deployment was a response to an explosion of siloed data analysis processes, pipelines, and tools that had emerged over the years as the company integrated more and more data into its workflows and decision-making. The transition to Power BI was rooted in an explicit decision by the company management to move towards becoming a more data-driven company, as well as to address the increasing difficulty of maintaining legacy tools as institutional knowledge about these outdated systems dwindled. TheCompany’s decision to deploy Power BI was based in part on an assessment of industry-standard BI tools against their specific needs for scalability, licensing flexibility, increased performance, and improved data quality.

We were given the opportunity to study the company during this transition to better understand both the advantages and challenges of commercial dashboarding systems in a large, conventional company. The interviews we conducted took place 6-12 months into the transition period. We provide an overview of the study in Figure 2.

3.1 Setting

TheCompany is a large, industrial manufacturing company established in the late-1800s, with over 50,000 employees and subsidiaries around

the world. TheCompany is headquartered in Austria and is active in the production of steel and steel-based technologies. The study took place with employees in the steel division, one of four divisions within the company. The study participants work at the company’s headquarters.

Although TheCompany is not an inherent tech company, they produce and use data extensively throughout their processes. TheCompany collects both structured data such as machine and sensor data, but also unstructured data such as text documenting the production process, and imaging data for defect identification within products. They also use big data throughout their business practices, including data about sales, costs, and internal expenses. A centralized IT department serves as a *data steward* [18], controlling access to TheCompany’s data; however, each individual business unit within TheCompany is responsible for its own data management, analysis, and reporting.

As a long-established company with data practices that have evolved over more than a century, across TheCompany exists a widerange of data analysis pipelines, tooling, and expertise. Some of the technology stacks in use are built from state-of-the-art commercial tools, some from older commercial tools with limited support, and others from in-house tools developed by individual employees. The decision to adopt Power BI stemmed from the workers’ need to replace siloed practices and legacy tools with a standardized set of processes, databases, and tools, and the top-level management’s desire to unify the reporting system across the company. To ensure proper coordination between Power BI users, one employee was appointed as the Power BI Coordinator.

3.2 Participants

The goal of our study was to understand the effects of transitioning to Power BI within TheCompany across a diverse set of employee roles, backgrounds, skills, and experiences. To find diverse participants, we first asked the Power BI Coordinator for a recommended list of names of employees who had recently started using the tool. This list consisted of workers in various IT and data analysis roles. We augmented this list with a convenience sampling of employees we had worked with over the years of the collaboration who worked in a broader range of groups and roles, and with whom we hoped to elicit open and detailed personal stories due the trust that had been established during the course of the collaborative project. This resulted in a list of 30 possible participants from which we selected 18 who spanned the broadest set of demographics – we previously worked with 8 of the 18. We contacted

Table 1: Demographics of our 17 interview participants. The relevant demographics are: distinction between management or worker role; gender (man/ woman); age; highest degree (H – high school, B – bachelors, M – masters, P – PhD); field of study (B – business/economics, CS – computer science, E – engineering, F – finance/accounting/logistics); total working experience and within the company (in years); level of expertise (proficient, knowledgeable, working, and little to none) in fields relevant to data analysis; data science roles specified by Crisan et al. [18] (*DSt* – data steward, *DSh* – data shaper, *DE* – data engineer, *ML/AI* – ML/AI engineer, *G* – generalist, *RS* – research scientist, *TA* – technical analyst, *M* – moonlighter, *E* – evangelist). The background color of the roles signifies higher order processes (preparation, deployment & engineering, analysis, communication).

| Role | Gender | Age | Degree | Field of study | Working experience (at the company/ total in years) | Level of expertise | | | | Data science roles | | | | | | | | | |
|------|--------|-----|--------|----------------|---|--------------------|------------------|------------------|-----------------------|--------------------|-----|----|-------|---|----|----|---|---|--|
| | | | | | | Statistics | Computer science | Domain knowledge | Human-centered design | DSt | DSh | DE | ML/AI | G | RS | TA | M | E | |
| | | 26 | M | F | 2/2 | | | | | | | | | | | | | | |
| | | 29 | M | F | 3.5/3.5 | | | | | | | | | | | | | | |
| | | 30 | M | B | 3/6 | | | | | | | | | | | | | | |
| | | 32 | M | F | 4/6 | | | | | | | | | | | | | | |
| | | 31 | B | CS | 5/10 | | | | | | | | | | | | | | |
| | | 30 | M | E | 4/11 | | | | | | | | | | | | | | |
| | | 36 | H | F | 14/14 | | | | | | | | | | | | | | |
| | | 34 | B | F | 14/15 | | | | | | | | | | | | | | |
| | | 36 | M | CS | 16/16 | | | | | | | | | | | | | | |
| | | 45 | M | F | 19/19 | | | | | | | | | | | | | | |
| | | 47 | P | E | 10/20 | | | | | | | | | | | | | | |
| | | 35 | M | CS | 3/15 | | | | | | | | | | | | | | |
| | | 44 | M | E | 17/20 | | | | | | | | | | | | | | |
| | | 43 | M | B | 16/21 | | | | | | | | | | | | | | |
| | | 47 | M | E | 20/25 | | | | | | | | | | | | | | |
| | | 52 | P | B | 20/27 | | | | | | | | | | | | | | |
| | | 57 | H | B | 18/30 | | | | | | | | | | | | | | |

these 18 employees via email and all agreed to participate. One of the participants brought a colleague with them to their interview.

We conducted semi-structured interviews with 19 participants. We discarded the two participants that interviewed as a pair as they shared significantly different, less-personal answers compared with the individual interviews. This resulted in 17 participant interviews that we analyzed. We received permission from managers at TheCompany to conduct the interviews. The participants were not compensated financially but were allowed to book as working time.

In Table 1, we detail relevant demographics of our participants including their age, gender, role (manager or worker), as well as their background declared as a field of study. The participants’ backgrounds span a range of fields from accounting to logistics, economics, business informatics, computer science, mathematics, and engineering. Additionally, we detail the participants’ working experience in years-in-total and working experience with TheCompany. It is worth noting that most of our participants have spent the majority of their careers at TheCompany, a common difference among employees at large, conventional companies versus employees in modern technology companies.

To capture the participants’ level of expertise with relevant fields as well as their roles in the data analysis pipeline, we had each participant complete a demographic survey one week prior to their interviews using LimeSurvey¹. Authors 1 and 2 mapped the participants’ responses from the survey to the categorizations of roles detailed by Crisan et al. [18]. The results are shown in Table 1. The survey questions can be found in the supplemental material.

3.3 Interviews

We divided the interview questions into three parts. In the first part, the interviewers introduced themselves, the research project, and the objective of the interview, clarified the compliance with data protection regulations, and asked for permission to audio record the conversation. The second stage gave the floor to the participant by first addressing general information such as their role at the company and their experiences using Power BI. More in-depth questions focused on the workers’ intrinsic motivations and problems faced when making the transition to Power BI. In the final stage of the interview, the interviewers asked the participant for any other things they wished to share, and they wrapped up the interview. The interview guide can be found in the

¹<https://www.limesurvey.org/>

supplementary material. Authors 1 and 2 piloted the interviews with three members of our research lab who work daily with visual data analyses. The pilot interviews allowed us to verify the appropriateness of our questions and the length of the interviews, as well as to practice the pragmatics of the interview process.

We conducted the interviews between September 2021 and January 2022. Authors 1 and 2 carried out the interviews using a pair-interview approach [2, 37]. One interviewer moderated the interview by asking questions from the interview guide and maintaining engagement with the participant. The second interviewer observed the interview, took notes, and asked in-depth follow-up questions. Each interview lasted for about one hour and was conducted either in person or online via Zoom². All interviews were audio recorded, and participants signed a consent form at the start of the interview. The interviews were conducted in German to ensure the participants felt comfortable expressing themselves candidly.

After each interview, the two interviewers reflectively discussed the interviews, sharing thoughts about the answers given by the participant and relating the interview to previous ones. Author 1 summarized these discussions in a reflective memo which was combined with the notes taken during the interview. These documents were included in the analysis of the interviews and discussed between all authors.

We initially tried to use automated services to transcribe the audio recordings, but the strong dialect of our participants was not supported. Thus, we hired an undergraduate student at Johannes Kepler University, Austria to transcribe the interviews; this student was paid an agreed-upon amount of 700 € for the work. Next, we used Deepl.com³ to translate the transcribed interviews into English as not all members of our research team are fluent in German. Author 1, who is a native German-speaker proofread the translated transcripts to fix mistranslations and grammar mistakes.

3.4 Analysis

To analyze the interviews, we followed the process of diffractive reading described by Akbaba et al. [2]. Authors 1, 2, and 4 conducted close readings of the transcripts, annotating them with reflective comments about what each of us found interesting. Taking the transcripts in groups of 2-4, we would individually read and annotate, then come together to

²<https://zoom.us/>

³<https://deepl.com/>

discuss what we found interesting in the transcripts and why. Talking through the interviews provided an opportunity to clarify the many idioms used by our participants, but more importantly, gave us a chance to reflect on the interviews through our different positionalities. This resulted in rich and generative discussions, which were documented and summarized on a digital whiteboard by Author 1.

While conducting the close readings of the transcripts, all authors periodically met to discuss interesting themes that were emerging. These group discussions continued after the close readings and became opportunities for feedback and refinement as Authors 1 and 4 pulled together quotes from the transcripts and drafted several different narrative structures for the results. In all, we produced three narrative drafts before settling on the results we present in Section 4.

When reporting quotes from the interviews in Section 4, we *tidied* them up for readability in English. This included modifying translations to account for local idioms and phrases, and removing filler words. We include the *un-tidied* English translations as well as the original German transcriptions in supplemental materials.

Author 1 additionally reviewed the Power BI website⁴⁵⁶ and categorized the marketing claims present on the website in August 2022. We discussed these claims, which informed our conversations about the challenges that our participants reported and use them to motivate our findings discussed in Section 4. We have included screenshots of the pages we reviewed in the supplemental materials.

3.5 Anonymization

In the interviews, we sought to elicit stories of the organizational and cultural challenges our participants faced. Many of the stories we heard were personal, and many of the participants interacted with each other at work. As a team, we discussed how best to protect the privacy of our participants, not only to the outside world but also within TheCompany. Thus, we have removed participant IDs from the demographic information in Table 1 to not link specific quotes with possibly internally identifiable demographic information. We sorted the table by the employees' role and then by the total working experience. Additionally, we have anonymized information about TheCompany as we report in this paper on internal details of their processes.

4 RESULTS

The marketing promise of Power BI – similar to other commercial dashboarding systems – is to “*create a data-driven culture... where anyone can prepare data, build machine learning models, and find insights quickly*”⁶. Power BI aims to “*empower team members to discover insights hidden in [their] data... [by] enabling everyone at every level of the organization to make confident decisions using up-to-the-minute analytics*”⁵.

At the time of our interviews, many employees that we spoke with were facing challenges in moving their jobs into the new Power BI workflow. Many were not making confident decisions or finding insights quickly. They expressed challenges in learning to use Power BI, in making sophisticated dashboards, and in finding their place in a changing socio-technical landscape. In this section, we report on these challenges and highlight important considerations for other companies that are transitioning to a commercial dashboarding system. We scaffold the section using the 14 observations we summarize in Figure 1 through boldfaced call-outs (**O#**: *observation*). In Section 5, we discuss specific opportunities for the visualization community to address these socio-technical challenges through new visualization research efforts.

4.1 Training is Hard

The transition to Power BI was largely pushed for, and planned by the IT department at TheCompany. This department anticipated a host of technical challenges the transition would bring: establishing licenses across thousands of computers; installing Power BI on employees' various computers; setting up the servers to host and manage data; and

connecting instances of Power BI to the databases. What was less anticipated, however, were the socio-technical challenges employees would face when transitioning to this new workflow. Some of these challenges emerged around training.

When asked about the experience of learning to use Power BI, V4 commented on the mismatch between expectations and reality:

V4: *[Power BI] was sold to us in such a way that it was very simple and very easy... the message was a little bit wrong. At least that's how it came across to us, 'anyone can work with it', and I'm still of the opinion that not everyone can work with it... Some of my colleagues were really overwhelmed.*

Initially, employee training was done by external consultants in a half-day workshop that came with the purchase of Power BI. This training was broad and generic, covering just the basics of the system. After this training, many employees struggled to use Power BI for their specific work tasks. In response to this, management brought in additional external consultants to train employees in using more specific features of the system and also asked members of our team to provide feedback on dashboards.

The management also established internal, peer-to-peer training through a monthly meet-up. The monthly meet-up aimed to be a space for informal exchanges by participants to share their issues and experiences with Power BI. It started with about 10 key employees and had grown to approximately 40 at the time of our interviews. While some of our participants found the monthly meet-ups productive, others reported that over time these meetings became less accessible to beginners. The Power BI Coordinator, V6, who organized the monthly meetings explained the challenges of maintaining effective peer training in light of the increasing gap between power users and beginners (**O1**: *intimidating expertise gaps*):

V6: *There are quite rudimentary and simple questions that come from someone who has not been working with it for so long, but there are of course others, the key users, who are already deeper in, who of course discuss more on an expert level.*

V8, who is not a power user, attends the meetings and described the intimidating power differentials produced by this knowledge gap:

V8: *There are very many who are already very deep inside, and some who have probably only just been introduced to the topic... In front of 30 people a beginner asks a question that probably for 10 people is an eye-roll... It's very difficult, maybe it will smooth out over time.*

The hesitation to publically ask for help in group training sessions extended across employee roles and was also evident at the management level. V7, who has taken on informal roles of supporting and training colleagues through the transition, noted:

V7: *Our management mostly doesn't dare to come to me and say, hey I don't understand that. There are a few that do, and a few I know for a fact that they should, but they don't. I deliberately did a training session once where the whole management sat together without the [worker] employees, so that maybe the fear of them asking questions was taken away a little bit. That actually didn't work too badly. But it's always the same issue: a lot of people don't really dare to say that they don't understand.*

Many of the people we spoke with relied on self-learning – in the form of online video tutorials, blogs, and help forums – to get answers to their specific questions (**O2**: *extensive self-learning*), described by V2 as *learning by doing*:

V2: *Learning by doing means I have the requirement, then I tried to implement it as best as possible. I downloaded*

⁴<https://powerbi.microsoft.com/en-us/>

⁵<https://powerbi.microsoft.com/en-us/why-power-bi/>

⁶<https://powerbi.microsoft.com/en-us/downloads/>

various books from the library. Also, YouTube videos are extremely good. And in my experience, whether it's with Excel or another tool, the problem that I'm facing now I'm sure somebody else has already faced, so Google is your best friend. So if you're good at finding information, you come relatively quickly to an answer.

While some participants thrived with self-learning – for example V5 told us about her enjoyment of spending hours googling for answers – other participants expressed self-learning as a perceived mandate:

V8: I was told to teach myself Power BI. I taught myself with YouTube videos in our fitness center down in the basement. I just watched it there on the side... I had to teach myself all this... That's how I slowly maneuvered myself into the world of Power BI.

A number of other participants similarly recounted their experiences of self-learning as taking place during personal hours (**O3: training during off-hours**). For example, V14 described coming into the office during off-hours to fit in learning:

V14: You find the time to come in, you take time out somewhere to come in... It wasn't all part of the working hours, and it was also a lot of private work, yes, just trying things out and watching videos.

Proficiency and comfort with English as a second language also influenced some participants' learning processes (**O4: foreign language barriers**). While V1 mentioned that googling for help required reading forums in English, V6 went further, detailing changes to his work practices to account for language differences:

V6: I watched a relatively large number of videos and then tried to do something along those lines. Although of course, I found that Power BI had to be switched to English relatively early on. Because if you have an error message in English, then you simply find much, much more on the internet than if you search for it in German.

We also heard multiple reports of participants relying on colleagues to help them learn and use Power BI (**O5: reliance on support networks**). V8 described her reliance on several colleagues – one of whom is a family member – for helping her to learn to create dashboards:

V8: I have to be honest and say that I would have had a hard time doing it on my own. I simply don't have the background or the know-how. But I had two supporters with whom I managed to do it.

These supportive roles were often informal and established through individual employees' personal networks of colleagues. Despite the ad-hoc nature of this support, this personalized and tailored help was deemed efficient:

V17: You are much faster if you can ask someone. With self-study, it takes longer.

Our findings highlight the importance of individualized approaches for helping employees transition to a disruptive, commercial dashboard system, which is in contrast to previous studies that report on the use of MOOCs for data science training [18] or that do not consider the breadth of training options [19]. While some of our participants benefited from group learning, such as in the peer-to-peer meet-up, many others relied on self-learning and the help of colleagues. This individualized learning was often not accounted for in employees' job requirements and required unseen and unrecognized labor. Additionally, employees with a strong social network within the company were able to benefit from the support of people around them.

These findings suggest that companies transitioning to a commercial dashboarding system should consider the time and effort required for

training, and scope this training within the job requirements for employees. This accounting should consider time for self-learning, as well as making supportive roles explicit, accessible, recognized, and valued. For the visualization community, these findings point to opportunities for developing organization-specific design guidelines that are tailored to the unique and specific needs of workers within an organization, which we discuss in Section 5.1.

4.2 Dashboards are Hard

It is common in large organizations to define standard operating procedures with step-by-step instructions to streamline the completion of routine tasks. Although TheCompany had documented procedures for many aspects of the organization, they did not develop one for creating dashboards. The result was that as departments within the company began to transition to Power BI they developed dashboards in ad-hoc ways, resulting in an uncontrolled growth of dashboards with little standardization between them.

As part of our collaboration, TheCompany asked members of our team to act as consultants and provide feedback to employees about their dashboards. We worked closely with about a dozen employees in several small groups to help them transition their existing static reports into Power BI dashboards, providing feedback about how to improve their visualizations, and suggesting different types of charts to use. After this consulting period, however, the employees struggled to create effective dashboards on their own (**O6: visualization knowledge gap**):

V7: What is also missing is the know-how behind [building dashboards]. How do you really visualize the data?... We have already received [feedback] from you and so on, but between reading [the feedback] and really being able to implement them is a gap. Maybe it is a mistake by us that we want to include far too much information in the report. Or we should only use one color but it becomes a rainbow at some point.

We heard similar reactions from both workers and managers who felt that “designing the visualization... that's not so easy” (V3), and expressed that creating effective visualizations in Power BI was like “starting from a blank sheet of paper” (V1). Additionally, participants reported that the work of data modeling – of getting the data “into a representation that you can present well” (V2) – was also a challenging process, and one that left little time for actually crafting visualizations:

V7: At the moment it's the case that we actually spend 80-90% of the time building data models. And maybe 10% on the visualization. I mean, you know our reports, they are not necessarily the finest reports in terms of visuals... We invest a lot of time in getting the numbers together in the first place. And then you have a little bit of time left over to make two or three graphs, and then you're already looking around for the next topic. It is unfortunate.

With little time to craft visualizations, and a lack of guidance as to how to do it well, some participants reported that they take their “standard visuals” (V5) from Excel – or what V14 called “the classics” – and remake them in Power BI:

V16: The biggest problem is, we all only know Excel and we're now doing the same thing in Power BI that we did in Excel. We still take bar charts, we still take pie charts, we still take line charts... The main problem is, we don't know what visualizations there are... What there actually is, we can't imagine.

And for participants that *did* engage with the myriad of visualization options in Power BI, the extent of available visualizations sometimes left them overwhelmed:

V14: I find the number of visuals that are now available a little bit, yes, that the selection is just so immense. I mean it's great on the one hand, but on the other hand, you can't decide how best to represent your data so that it is usable, because there are just so many possibilities.

We heard from multiple workers that the time-intensive labor of creating effective dashboards was not always visible or appreciated (**O7: underappreciated labor**), including V6 who hears about many struggles from employees in the monthly training meet-ups he organizes:

V6: *[Managers] always just see the visualizations and say 'the beautiful visualizations'. But they don't see the workload behind it. . . And I think that's still a sticking point for the management, to [allocate] resources to create dashboards. That it's not one hour a week, but I need, I don't know, three person-days a week where I can deal with it exclusively, otherwise it simply won't work.*

Furthermore, using interactive dashboards introduced yet other challenges. In particular, the interactivity of the dashboards was unnerving for some due to the updating and changing nature of the views, leading to a lack of trust and adoption (**O8: discomfort with interaction**). V4 lamented that colleagues were still relying heavily on Excel, in part due to their discomfort with interactive views:

V4: *We still have the issue that [dashboards] are not used enough at the management level. . . Somehow there is always this blockade to getting used because they are too interactive. I often get this feedback that, 'wow, I click around there, and around there, and then everything changes'.*

One particularly tricky interaction that participants commented on was that of filtering – a finding also reported by Sarikaya et al. [46] – the effects of which were not always well understood:

V8: *I think the hardest thing for me to understand, as a non-IT-savvy person, was the filters. So if I filter something, does that only apply to that one visualization? Or if I select a menu setting for another visualization, will that apply to all? . . . That for me, at the beginning, was not so easily comprehensible.*

Our participants also explained that interactions in a dashboard were sometimes confounded with data modeling (**O9: fear of breaking data**). V14 told us that she observed colleagues being fearful of their interactions “breaking” the data, which was also reported by V9:

V9: *They first have to get into the habit and realize that it's not bad if they press three times, because [they are] not changing anything. What I often get questions about is whether they really believe they are changing something in the data just because they are now filtering.*

The challenges of using interactive dashboards meant that at the time of our interviews, most teams and employees were still relying on static reports – either in a PDF or a PowerPoint slide deck – in meetings and decision-making contexts, despite having access to Power BI for 6-12 months (**O10: continued reliance on static reports**). These static reports are often a series of complex plots generated from Excel sheets that are cut-and-paste into a PDF for sending around, or into a PowerPoint slide deck for presenting in a meeting. We heard from participants that sometimes these reports become unwieldy and hundreds of pages long, and also require significant overhead to keep updated with the most recent data:

V7: *[It] was really the case that every day in the morning, for the first hour and a half to two hours, we did nothing else but update reports.*

Despite the inefficiency of static reports, V12 speculated that only 30% of the reports were using Power BI dashboards:

V12: *We still have a lot of static reports, PDF reports that are generated daily. . . It's really still the way we did it 20 years ago. Of course, perhaps made more elegant in terms of look and feel. But I do believe that the most dominant part of all reports is certainly still a static PDF.*

Our participants told us that their static reports are often annotated with important context, opinions, and comments that support understanding why the charts (and the data) are meaningful and what those insights mean for decision-making. Often they would print these static charts out on paper to add their own notes and thoughts before and during meetings. V17, who is a manager, speculated that because the new dashboards did not include similar support for annotation and guided narratives many of his colleagues were continuing to use the company's older, static reporting formats:

V17: *So who structures [the information in a dashboard] for us? In which dashboard do I find what? Where is the clever description? What is the dashboard supposed to do? What is it supposed to be good for? What does it all represent? . . . It becomes so confusing that [managers] say, okay I don't know now what [this dashboard is meant to say], so I'll take my old familiar report.*

On the other hand, V4 told us about a dashboard that was being used productively in meetings, providing an information-rich view of the status of production. This success, though, was dependent on the dashboard creator guiding colleagues through the visualizations (**O11: narration necessary**):

V4: *One [successful dashboard] is from my colleague, where you really look at deviations in production. And that is something that is actually [designed for interaction], that you click through a bit and look at it a bit in detail. That you can look at the plants, where do any deviations come from now. And this is really used, but only in meetings where my colleague shows it. It is not the case that anyone who can or wants to work with it independently is likely to do so. It is shown and he must then navigate through for us.*

The findings we bring forward around the creation and use of dashboards indicate that dashboards – and visualizations – remain difficult to produce and consume [24, 28, 46]. While initial consultations with visualization experts helped employees to turn existing Excel charts into Power BI dashboards, creating new dashboards from scratch was difficult. Using the dashboards was also difficult, requiring assumed knowledge about how interactive mechanisms worked and a guide to narrate through the views, similar to findings by others [22, 43].

Despite promises to empower users, companies should not underestimate the complexity of commercial dashboarding systems. Creating and using dashboards requires significant training, time, and attention. For visualization researchers, the difficulty of dashboards is an opportunity to develop specific guidance for designing dashboards within an organization, and to bring storytelling capabilities into dashboarding systems. We further discuss these opportunities in Section 5.2 and 5.3.

4.3 Transitions are Hard

Transitioning a conventional company that has siloed, legacy data practices to one that instead is integrated and standardized is a disruptive process [55]. For some of our participants, the transition signaled that TheCompany was “becoming a bit more modern” and “arriving in the 21st century” (V14). Other participants took this sentiment further, seeing the transition as a generational shift in work practices:

V12: *There will always be those people who simply rely on the old traditions: we've always done that, that's the way it's going to be, and I print out my report and that's the way I want it, and in the worst case I take a screenshot and I print it out in Word. So that will certainly exist, but these dinosaurs will probably also die out at some point. But now, I would say, from the newer generation, where I would count myself among them, I rather believe that this will be received positively. Because it simply has more advantages. . . I also believe that this is a revolution, and we won't be able to stop it, even if, as individuals, one or the other might fight it tooth and nail.*

And with this transition, some employees thrived: in particular, those with an affinity for, confidence with, and interest in technology (**O12: winners and losers**). V9, for example, who has an educational background in business analytics and finance, described how her technical skills and interests allowed her to transition to Power BI more effectively than her peers:

V9: *Power BI, it's a very powerful tool. I always find something like that very exciting when you can look at what could work and how... If I now compare [my technical skills] with a classic accountant in our company, I would definitely say yes, I learn quickly... Power BI interests me and that's why I believe that my approach is different from the standard finance colleague, I would say.*

Similarly, V2 – who has a background in computer science – fondly described his work with Power BI as feeling “a bit like my home”.

But given the large number of employees in the company, and its core business in long-established manufacturing processes, the diversity of employees presented the full range of affinity for new technology. V7 noted the challenges he saw many of his colleagues had with the transition to Power BI:

V7: *It's so difficult because we have so many different people: people who are technically more affine like me, but then you have people who are already overwhelmed when they have to open a PowerPoint file.*

It was also noted by multiple participants that with the transition came a competitive pressure among workers. For example, V4 reported pressure to achieve quick results:

V4: *[The pressure came] more from the management. They said, 'okay, now we have to see that we somehow get something together very quickly.' ... This a bit of a competition between the departments.*

Across many of our interviews, we heard about the different ways in which the transition to Power BI was requiring a new level of analytical and technical competency across the company (**O13: benefits of broad competencies**). This competency was deemed by some managers as critical for maintaining an engaged and empowered workforce as the company moved to rely more on data:

V9: *I see that a lot comes from acceptance when they understand what's going on behind the scenes and when it's no longer this black hole where everything falls in at the top and then my beautiful dashboard comes out at the bottom. The more they understand, the more they question. That's quite observable on the management floor... So I think that on the technical side, not everyone has to be able to build dashboards themselves, but at least understand the basics.*

Many workers, however, described the challenges of being expected to understand the full visual analysis pipeline and the continued need for specialized skills (**O14: need for specialization**). For example V7, who is a *data shaper* [18], regularly helped colleagues above and beyond his explicit job responsibilities to get data models into formats for their dashboards. He reflected on this work:

V7: *Somehow we are currently in the mindset that everyone must be able to work with Power BI and should be able to do everything themselves... If I'm talking about the most complex data models, where I have to, I don't know, connect five different systems behind it and, I don't know, understand the Snowflake Schema and blah blah blah. I mean, I can't expect everybody to do that.*

We also heard about the ways that domain knowledge continued to play an important role in making sense of data from V5:

V5: *It is the case that especially with these very large reports, the salespeople have a better sense of what could be true and what not. They look at their customers' data and can immediately tell if the number is realistic. I can't do that at all... I got a call once, for example, with a salesperson saying I have the same numbers in there today as I had yesterday in relation to this one graphic. And I thought to myself, I don't know what numbers were in there yesterday... I wouldn't notice something like that.*

Our findings confirm that transitioning a conventional company like TheCompany to a VA system like Power BI is indeed disruptive, not only to the technical infrastructure but also to the social, cultural, and organizational fabric of a company. Despite promises of self-service and empowerment of employees, we instead found that data modeling, dashboard creation, and decision-making all required specialized knowledge and training, which reflects findings by others that data work in large companies continues to rely on cooperation between specialized employees [47, 58]. We see the need for companies to consider building teams of specialists [40], rather than expect that employees can become VA generalists. For visualization researchers, we see an opportunity to explicitly design visual analytics and dashboarding systems for collaborative teams, which we discuss further in Section 5.4.

5 RESEARCH OPPORTUNITIES

Here we outline several opportunities for the visualization research community, motivated by our findings presented in Section 4.

5.1 Organization-specific Design Guidelines

We heard from many of our participants that generalized training was not sufficient for learning how to use Power BI; most of them reported a significant amount of self-learning. In line with insights reported by Bako et al. [9] to have struggles in designing visualizations, our participants confirmed that creating visualizations remains a challenge despite decades of research, books, articles, and tools meant to bring visualization capabilities to the masses. This made us wonder: what if visualization design guidelines are *too* general?

The visualization research community has developed a plethora of design guidelines for creating domain-agnostic charts. More specific work on dashboards proposes considerations for audience, purpose, and data semantics [46], as well as design patterns and genres [6]. But these guidelines are general and not tailored to the specific needs and constraints of individual domains and organizations, nor do they fully account for the diversity of *novice* users [14].

In contrast, work done by practitioners within different organizations provides examples of organization-specific design guidelines⁷. For example, the Financial Times Visual Vocabulary⁸ suggests suitable visualizations for specific tasks; the IBM design guide includes dedicated sections on data visualization and infographics⁹; and the Novartis graphics principles include a compact visualization cheat sheet¹⁰. These guides are tailored to the needs, domain problems, and use cases of the organizations in which they are deployed.

Unlike studies that reveal the desire of professional designers to have fine-level control of their visualization design processes [35, 41, 42], our participants were neither trained nor experienced in design and struggled to create effective and consistent dashboards. Instead, we point to the success of design studies and custom visual analysis tools as highlighting the opportunity to study and create organization-specific visualization design guidelines. The availability of organization-specific guidelines could support dashboard authors in creating effective visualizations that speak the same design language. This support could also speed up the creation process by freeing authors from the burden of thinking about low-level design decisions and scoping some of the more expansive decisions [46]. Similar to design studies, we could publish organization- or domain-specific design guidelines that have

⁷<https://policyviz.com/2016/11/30/style-guides/>

⁸<https://www.ft.com/vocabulary/>

⁹<https://www.ibm.com/design/language/>

¹⁰<https://graphicsprinciples.github.io/>

been developed together with collaborators from various organizations and domains. Such guidelines could include:

- characterizations of the data and tasks of the core, target users;
- suggested visualization charts and techniques for the most common data and task combinations;
- recommendations for using color and other encoding channels;
- defined behaviors for selections and filters between the components of a dashboard; and
- a design language that follows an organization's corporate design, including color palettes and fonts.

These types of design guidelines are not typically part of the standard documentation that comes with dashboarding tools. The help content mostly includes technical instructions on how to achieve something in the tool – *how to set and change the filtering behavior* – rather than guiding users to make the good choices in the first place – *which filtering behavior is beneficial in certain scenarios and why*. We see an opportunity for the visualization research community to fill this gap.

5.2 Dashboard Design Process Models

Our interviews confirmed that creating effective dashboards is a challenging and time-consuming undertaking and one that is often unstructured and tedious. Creating dashboards includes many challenging steps beyond designing effective visualizations: collecting data analysis needs and requirements; iteratively implementing a dashboard design; and deploying, monitoring, and continuously improving a released dashboard. The tutorials and documentation that come with commercial dashboards largely focus on introducing the features for creating dashboards rather than the complete design process itself.

Recent work by Bach et al. calls out the need for dashboard design processes to help structure the design process and articulate best practices [6]. They propose one such process model that includes steps for determining the context, data, structure, visual representations, layout, and interactivity. But we suggest these models should go further, including both formative and summative design work in assessing needs and testing results. Like our call for organization-specific design guidelines in Section 5.1, these process models could be tailored to organizations to reduce much of the work of characterizing the domain, the workforce, workflows, and communication paths. These process models must strike a balance between being a lightweight process that does not add unnecessary bureaucracy while containing sufficient detail to guide dashboard creators through the process. We are currently exploring an organization-specific dashboard design model as part of our collaboration with TheCompany.

5.3 Narrative Dashboards

Across our interviews, we heard from participants about their challenges using dashboards. They felt intimidated by filtering and using interactive features. They got confused when a click would change multiple views. They did not know where to look at or why to care. They needed narration to make sense of a myriad of charts and complex data. And in the end, many returned to their traditional, static reports that were rich with explanations and structured into a clear, linear story, findings also reported by Tory et al. [58].

The visualization community, however, has developed an extensive body of knowledge of how to design, author, and present data-driven stories [16, 33, 44, 49]. Also referred to as *narrative visualization*, this work investigates how to combine visualizations with annotations, narrations, and interaction techniques to communicate ideas from and with data. Building from the initial characterization of data-driven storytelling by Segel & Heer [49], more recent work in the space has explored combining comics and data [7, 8], accessibility through automatic audio narrations [51], scrollytelling with interactive visualizations [39, 53], and techniques for virtual narration of charts [15, 30]. This line of research has further extended into *data videos* that augment visualizations with animations, audio and text narrations, music, and sound effects to elicit engagements with data [4, 5, 45, 50, 59].

In short, the visualization research community knows a lot about how to communicate with data, and has innovated a broad array of

techniques to do so. We see an opportunity to explicitly bring these ideas into dashboarding systems – and other VA tools – to allow data workers to enrich dashboards with their insights and important context. A step in this direction could be to build off with provenance tracking techniques that capture exploration steps and then provide opportunities for analysts to author stories about their sessions [29].

5.4 Dashboard Tools for Teams

Despite the promises of self-service and empowerment that come with commercial dashboarding systems, we heard from many of our participants that the process of making decisions from complex data within a large organization requires specialized knowledge and skills. This echoes arguments that encourage organizations to hire teams of specialists rather than chase after unicorns [12, 20, 40]. Dashboarding systems, however, are traditionally designed with single users in mind. What might we design dashboards to do and support if we assumed that different specialists work on different parts of the analysis pipeline?

The visualization and HCI communities have a long history of studying collaborative work practices and designing technology to support these practices. For example, early work on the LARK system [57] allowed multiple users to concurrently interact with the visualization pipeline on a tabletop display. More recently, the emerging field of immersive analytics follows similar goals but is supported by AR/VR technology [26]. And ideas such as *data hunches* point to opportunities for externalizing and communicating expert knowledge, facilitated by visualizations [34]. We suggest that it is worth considering how these existing bodies of work might inspire new approaches to making dashboard systems facilitate inherently collaborative working practices.

6 CONCLUSIONS AND FUTURE WORK

In this paper, we present the results of an interview study with employees transitioning their work practices to Power BI at a large, conventional company. These results provide new observations into the socio-technical challenges people face when adopting commercial dashboarding systems and complement existing studies that focus on the technical challenges of data workers. Our findings point to a number of opportunities for the research community to develop new guidelines and rethink how dashboard tools are designed. An interesting subject for future research involves conducting a comparative analysis of our empirical discoveries, alongside similar studies, with best practices recommended by dashboarding software providers and practitioners through their product documentation, blog posts, and video tutorials.

The limitations of this work are similar to those of other interview studies: our detailed accounts of a small group of people within a single organization may not generalize to other data workers and other settings. Further, our interviews represent a snapshot of time during a lengthy transition process. As such, we are interested in conducting follow-up studies with our participants to see how their experiences change (or not) with time. Also, we encourage further research into socio-technical considerations in a broader range of settings and populations.

ACKNOWLEDGMENTS

We would like to thank Jimmy Moore for the input-rich discussions at the beginning of the project, our interviewees for their time and candid participation in the study, and our reviewers for their feedback. This work was supported in part by funding from the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation, the Austrian Science Fund (FWF DFH 23–N), and the Austrian Research Promotion Agency (FFG 881844). Pro²Future is funded within the Austrian COMET Program under the auspices of the Austrian Federal Ministry for Climate Action, Environment, Energy, Mobility, Innovation and Technology, the Austrian Federal Ministry for Digital and Economic Affairs, and of the States of Upper Austria and Styria. COMET is managed by the Austrian Research Promotion Agency FFG.

REFERENCES

- [1] C. Ahlberg. Spotfire: An Information Exploration Environment. *ACM SIGMOD Record*, 25(4):25–29, 1996. doi: 10.1145/245882.245893 2

- [2] D. Akbaba, D. Lange, M. Correll, A. Lex, and M. Meyer. Troubling Collaboration: Matters of Care for Visualization Design. In *ACM CHI Conference on Human Factors in Computing Systems*, 2023. doi: 10.1145/3544548.3581168 4
- [3] S. Alspaugh, N. Zokaie, A. Liu, C. Jin, and M. A. Hearst. Futzng and Moseying: Interviews with Professional Data Analysts on Exploration Practices. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):22–31, 2019. doi: 10.1109/TVCG.2018.2865040 1, 2
- [4] F. Amini, N. Henry Riche, B. Lee, C. Hurter, and P. Irani. Understanding Data Videos: Looking at Narrative Visualization through the Cinematography Lens. In *Conference on Human Factors in Computing Systems*, pp. 1459–1468. ACM, 2015. doi: 10.1145/2702123.2702431 9
- [5] F. Amini, N. H. Riche, B. Lee, A. Monroy-Hernandez, and P. Irani. Authoring Data-Driven Videos with DataClips. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):501–510, 2017. doi: 10.1109/TVCG.2016.2598647 9
- [6] 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*, pp. 1–11, 2022. doi: 10.1109/TVCG.2022.3209448 2, 8, 9
- [7] B. Bach, N. Kerracher, K. W. Hall, S. Carpendale, J. Kennedy, and N. Henry Riche. Telling Stories about Dynamic Networks with Graph Comics. In *Conference on Human Factors in Computing Systems*, pp. 3670–3682. ACM, 2016. doi: 10.1145/2858036.2858387 9
- [8] B. Bach, Z. Wang, M. Farinella, D. Murray-Rust, and N. Henry Riche. Design Patterns for Data Comics. In *ACM CHI Conference on Human Factors in Computing Systems*, pp. 38:1–38:12, 2018. doi: 10.1145/3173574.3173612 9
- [9] H. K. Bako, X. Liu, L. Battle, and Z. Liu. Understanding How Designers Find and Use Data Visualization Examples. *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–11, 2022. doi: 10.1109/TVCG.2022.3209490 8
- [10] L. Bartram, M. Correll, and M. Tory. Untidy Data: The Unreasonable Effectiveness of Tables. *IEEE Transactions on Visualization and Computer Graphics*, 28(01):686–696, 2022. doi: 10.1109/TVCG.2021.3114830 1, 2
- [11] A. Batch and N. Elmqvist. The Interactive Visualization Gap in Initial Exploratory Data Analysis. *IEEE Transactions on Visualization and Computer Graphics*, 24:278–287, 2018. doi: 10.1109/TVCG.2017.2743990 1, 2
- [12] S. Baškarada and A. Koronios. Unicorn data scientist: the rarest of breeds. 51(1):65–74, 2017. doi: 10.1108/PROG-07-2016-0053 9
- [13] M. Behrisch, D. Streeb, F. Stoffel, D. Seebacher, B. Matejek, S. H. Weber, S. Mittelstaedt, H. Pfister, and D. Keim. Commercial Visual Analytics Systems-Advances in the Big Data Analytics Field. *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–1, 2018. doi: 10.1109/TVCG.2018.2859973 2
- [14] A. Burns, C. Lee, R. Chawla, E. Peck, and N. Mahyar. Who Do We Mean When We Talk About Visualization Novices? In *ACM CHI Conference on Human Factors in Computing Systems*, pp. 1–16, 2023. doi: 10.1145/3544548.3581524 8
- [15] Y.-R. Cao, J.-Y. Pan, and W.-C. Lin. User-oriented Generation of Contextual Visualization Sequences. In *ACM CHI Conference on Human Factors in Computing Systems*, pp. 1–8, 2020. doi: 10.1145/3334480.3383057 9
- [16] Q. Chen, S. Cao, J. Wang, and N. Cao. How Does Automation Shape the Process of Narrative Visualization: A Survey of Tools. 2022. doi: 10.48550/ARXIV.2206.12118 9
- [17] A. Crisan and B. Fiore-Gartland. Fits and Starts: Enterprise Use of AutoML and the Role of Humans in the Loop. In *ACM CHI Conference on Human Factors in Computing Systems*, pp. 1–15, 2021. doi: 10.1145/3411764.3445775 2
- [18] A. Crisan, B. Fiore-Gartland, and M. Tory. Passing the Data Baton : A Retrospective Analysis on Data Science Work and Workers. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1860–1870, 2021. doi: 10.1109/TVCG.2020.3030340 1, 2, 3, 4, 6, 8
- [19] M. K. Daradkeh. Determinants of visual analytics adoption in organizations: Knowledge discovery through content analysis of online evaluation reviews. *Information Technology & People*, 32(3):668–695, 2019. doi: 10.1108/ITP-10-2017-0359 1, 2, 6
- [20] T. Davenport. Beyond Unicorns: Educating, Classifying, and Certifying Business Data Scientists. *Harvard Data Science Review*, 2020. doi: 10.1162/99608f92.55546b4a 9
- [21] V. Dhanoa, C. Walchshofer, A. Hinterreiter, H. Stitz, E. Groeller, and M. Streit. A Process Model for Dashboard Onboarding. *Computer Graphics Forum*, 41(3):501–513, 2022. doi: 10.1111/cgf.14558 3
- [22] E. Dimara, H. Zhang, M. Tory, and S. Franconeri. The Unmet Data Visualization Needs of Decision Makers Within Organizations. *IEEE Transactions on Visualization and Computer Graphics*, 28(12):4101–4112, 2022. doi: 10.1109/TVCG.2021.3074023 1, 2, 7
- [23] S. Eick. Visual discovery and analysis. *IEEE Transactions on Visualization and Computer Graphics*, 6(1):44–58, 2000. doi: 10.1109/2945.841120 2
- [24] M. Elias and A. Bezerianos. Exploration Views: Understanding Dashboard Creation and Customization for Visualization Novices. In *Human-Computer Interaction*, vol. 6949, pp. 274–291. Springer Berlin Heidelberg, 2011. doi: 10.1007/978-3-642-23768-3_23 7
- [25] D. Fisher, R. DeLine, M. Czerwinski, and S. Drucker. Interactions with big data analytics. *Interactions*, 19(3):50–59, 2012. doi: 10.1145/2168931.2168943 1, 2
- [26] A. Fonet and Y. Prié. Survey of Immersive Analytics. *IEEE Transactions on Visualization and Computer Graphics*, 27(3):2101–2122, 2021. doi: 10.1109/TVCG.2019.2929033 9
- [27] K. Furmanova, S. Gratzl, H. Stitz, T. Zichner, M. Jaresova, A. Lex, and M. Streit. Taggle: Combining overview and details in tabular data visualizations. *Information Visualization*, 19(2):114–136, 2020. doi: 10.1177/1473871619878085 3
- [28] L. Grammel, M. Tory, and M. Storey. How Information Visualization Novices Construct Visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):943–952, 2010. doi: 10.1109/TVCG.2010.164 7
- [29] S. Gratzl, A. Lex, N. Gehlenborg, N. Cosgrove, and M. Streit. From Visual Exploration to Storytelling and Back Again. *Computer Graphics Forum*, 35(3):491–500, 2016. doi: 10.1111/cgf.12925 9
- [30] B. D. Hall, L. Bartram, and M. Brehmer. Augmented Chironomia for Presenting Data to Remote Audiences. In *ACM Symposium on User Interface Software and Technology*, pp. 1–14. ACM, 2022. doi: 10.1145/3526113.3545614 9
- [31] S. Kandel, A. Paepcke, J. M. Hellerstein, and J. Heer. Enterprise Data Analysis and Visualization: An Interview Study. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2917–2926, 2012. doi: 10.1109/TVCG.2012.219 1, 2
- [32] M. Kim, T. Zimmermann, R. DeLine, and A. Begel. The emerging role of data scientists on software development teams. In *International Conference on Software Engineering*, pp. 96–107. ACM, 2016. doi: 10.1145/2884781.2884783 1, 2
- [33] R. Kosara and J. Mackinlay. Storytelling: The next step for visualization. *Computer*, 46(5):44–50, 2013. doi: 10.1109/MC.2013.36 9
- [34] H. Lin, D. Akbaba, M. Meyer, and A. Lex. Data Hunches: Incorporating Personal Knowledge into Visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):504–514, 2023. doi: 10.1109/TVCG.2022.3209451 9
- [35] J. Liu, N. Boukhelifa, and J. R. Eagan. Understanding the Role of Alternatives in Data Analysis Practices. *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–1, 2019. doi: 10.1109/TVCG.2019.2934593 2, 8
- [36] Y. Liu, T. Althoff, and J. Heer. Paths Explored, Paths Omitted, Paths Obscured: Decision Points & Selective Reporting in End-to-End Data Analysis. In *ACM CHI Conference on Human Factors in Computing Systems*, pp. 1–14, 2020. doi: 10.1145/3313831.3376533 1, 2
- [37] J. Monforte and J. Úbeda Colomer. Tinkering with the two-to-one interview: Reflections on the use of two interviewers in qualitative constructionist inquiry. *Methods in Psychology*, 5:100082, 2021. doi: 10.1016/j.metip.2021.100082 4
- [38] M. Muller, I. Lange, D. Wang, D. Piorkowski, J. Tsay, Q. V. Liao, C. Dugan, and T. Erickson. How Data Science Workers Work with Data: Discovery, Capture, Curation, Design, Creation. In *ACM CHI Conference on Human Factors in Computing Systems*, pp. 126:1–126:15, 2019. doi: 10.1145/3290605.3300356 2
- [39] E. Mörth, S. Bruckner, and N. N. Smit. ScrollyVis: Interactive visual authoring of guided dynamic narratives for scientific scrollytelling. 2022. doi: 10.48550/ARXIV.2207.03616 9
- [40] C. O’Neil and R. Schutt. *Doing Data Science: Straight Talk from the Frontline*. O’Reilly Media, Inc., 2013. 8, 9
- [41] P. Parsons. Understanding Data Visualization Design Practice. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):665–675, 2022. doi: 10.1109/TVCG.2021.3114959 8
- [42] P. Parsons, C. M. Gray, A. Baigelenov, and I. Carr. Design Judgment in Data Visualization Practice. In *IEEE Conference on Information Visual-*

- ization, pp. 176–180, 2020. doi: [10.1109/VIS47514.2020.00042](https://doi.org/10.1109/VIS47514.2020.00042) 8
- [43] L. M. Perkhofor, P. Hofer, C. Walchshofer, T. Plank, and H.-C. Jetter. Interactive visualization of big data in the field of accounting: A survey of current practice and potential barriers for adoption. *Journal of Applied Accounting Research*, 20(4):497–525, 2019. doi: [10.1108/JAAR-10-2017-0114](https://doi.org/10.1108/JAAR-10-2017-0114) 7
- [44] N. H. Riche, C. Hurter, N. Diakopoulos, and S. Carpendale. *Data-Driven Storytelling*. CRC Press, 2018. 9
- [45] S. Sallam, Y. Sakamoto, J. Leboe-McGowan, C. Latulipe, and P. Irani. Towards Design Guidelines for Effective Health-Related Data Videos: An Empirical Investigation of Affect, Personality, and Video Content. In *ACM CHI Conference on Human Factors in Computing Systems*, pp. 1–22, 2022. doi: [10.1145/3491102.3517727](https://doi.org/10.1145/3491102.3517727) 9
- [46] A. Sarikaya, M. Correll, L. Bartram, M. Tory, and D. Fisher. What Do We Talk About When We Talk About Dashboards? *IEEE Transactions on Visualization and Computer Graphics*, 25(1):682–692, 2019. doi: [10.1109/TVCG.2018.2864903](https://doi.org/10.1109/TVCG.2018.2864903) 2, 7, 8
- [47] M. Sedlmair, P. Isenberg, D. Baur, and A. Butz. Evaluating information visualization in large companies: challenges, experiences and recommendations. In *Workshop on Beyond Time and Errors: Novel Evaluation Methods for Visualization*, pp. 79–86. ACM, 2010. doi: [10.1145/2110192.2110204](https://doi.org/10.1145/2110192.2110204) 2, 8
- [48] M. Sedlmair, M. Meyer, and T. Munzner. Design Study Methodology: Reflections from the Trenches and the Stacks. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2431–2440, 2012. doi: [10.1109/TVCG.2012.213](https://doi.org/10.1109/TVCG.2012.213) 2
- [49] E. Segel and J. Heer. Narrative Visualization: Telling Stories with Data. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1139–1148, 2010. doi: [10.1109/TVCG.2010.179](https://doi.org/10.1109/TVCG.2010.179) 9
- [50] D. Shi, F. Sun, X. Xu, X. Lan, D. Gotz, and N. Cao. AutoClips: An Automatic Approach to Video Generation from Data Facts. *Computer Graphics Forum*, 40(3):495–505, 2021. doi: [10.1111/cgf.14324](https://doi.org/10.1111/cgf.14324) 9
- [51] A. Siu, G. S-H Kim, S. O’Modhrain, and S. Follmer. Supporting Accessible Data Visualization Through Audio Data Narratives. In *ACM CHI Conference on Human Factors in Computing Systems*, pp. 1–19, 2022. doi: [10.1145/3491102.3517678](https://doi.org/10.1145/3491102.3517678) 9
- [52] C. Stoiber, F. Grassinger, M. Pohl, H. Stitz, M. Streit, and W. Aigner. Visualization Onboarding: Learning How to Read and Use Visualizations. In *IEEE Workshop on Visualization for Communication*. IEEE, 2019. 3
- [53] C. Stoiber, C. Walchshofer, F. Grassinger, H. Stitz, M. Streit, and W. Aigner. Design and Comparative Evaluation of Visualization Onboarding Methods. In *Proceedings of the Symposium on Visual Information Communication and Interaction*, pp. 1–5. ACM, 2021. doi: [10.1145/3481549.3481558](https://doi.org/10.1145/3481549.3481558) 9
- [54] C. Stolte, D. Tang, and P. Hanrahan. Polaris: A System for Query, Analysis, and Visualization of Multidimensional Relational Databases. *IEEE Transactions on Visualization and Computer Graphics*, 8(1):52–65, 2002. doi: [10.1109/2945.981851](https://doi.org/10.1109/2945.981851) 2
- [55] M. Storm and H. Borgman. Understanding Challenges and Success Factors in Creating a Data-Driven Culture. In *Hawaii International Conference on System Sciences*, 2020. doi: [10.24251/HICSS.2020.663](https://doi.org/10.24251/HICSS.2020.663) 7
- [56] M. Streit, S. Gratzl, H. Stitz, A. Wernitznig, T. Zichner, and C. Haslinger. Ordino: a visual cancer analysis tool for ranking and exploring genes, cell lines and tissue samples. *Bioinformatics*, 35(17):3140–3142, 2019. doi: [10.1093/bioinformatics/btz009](https://doi.org/10.1093/bioinformatics/btz009) 3
- [57] M. Tobiasz, P. Isenberg, and S. Carpendale. Lark: Coordinating Co-located Collaboration with Information Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1065–1072, 2009. doi: [10.1109/TVCG.2009.162](https://doi.org/10.1109/TVCG.2009.162) 9
- [58] M. Tory, L. Bartram, B. Fiore-Gartland, and A. Crisan. Finding Their Data Voice: Practices and Challenges of Dashboard Users. *IEEE Computer Graphics and Applications*, pp. 1–1, 2022. doi: [10.1109/MCG.2021.3136545](https://doi.org/10.1109/MCG.2021.3136545) 1, 2, 8, 9
- [59] X. Xu, L. Yang, D. Yip, M. Fan, Z. Wei, and H. Qu. From ‘Wow’ to ‘Why’: Guidelines for Creating the Opening of a Data Video with Cinematic Styles. In *ACM CHI Conference on Human Factors in Computing Systems*, pp. 1–20, 2022. doi: [10.1145/3491102.3501896](https://doi.org/10.1145/3491102.3501896) 9
- [60] M. A. E.-A. Youssef, R. Eid, and G. Agag. Cross-national differences in big data analytics adoption in the retail industry. *Journal of Retailing and Consumer Services*, 64:102827, 2022. doi: [10.1016/j.jretconser.2021.102827](https://doi.org/10.1016/j.jretconser.2021.102827) 2
- [61] L. Zhang, A. Stoffel, M. Behrisch, S. Mittelstadt, T. Schreck, R. Pompl, S. Weber, H. Last, and D. Keim. Visual analytics for the big data era - A comparative review of state-of-the-art commercial systems. In *IEEE Conference on Visual Analytics Science and Technology*, pp. 173–182, 2012. doi: [10.1109/VAST.2012.6400554](https://doi.org/10.1109/VAST.2012.6400554) 2
- [62] Y. Zhang, Y. Sun, J. D. Gaggiano, N. Kumar, C. Andris, and A. G. Parker. Visualization Design Practices in a Crisis: Behind the Scenes With COVID-19 Dashboard Creators. *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–11, 2022. doi: [10.1109/TVCG.2022.3209493](https://doi.org/10.1109/TVCG.2022.3209493) 2