

Beyond Trust Building — Calibrating Trust in Visual Analytics

Wenkai Han*

Hans-Jörg Schulz†

Aarhus University, Denmark

ABSTRACT

Trust is a fundamental factor in how users engage in interactions with Visual Analytics (VA) systems. While the importance of building trust to this end has been pointed out in research, the aspect that trust can also be misplaced is largely ignored in VA so far. This position paper addresses this aspect by putting trust calibration in focus – i.e., the process of aligning the user’s trust with the actual trustworthiness of the VA system. To this end, we present the trust continuum in the context of VA, dissect important trust issues in both VA systems and users, as well as discuss possible approaches that can build and calibrate trust.

Index Terms: Human-centered computing—Visualization—Visualization theory, concepts and paradigms; Human-centered computing—HCI theory, concepts and models—Visualization design and evaluation methods;

1 INTRODUCTION

In one of the most cited paper on Visual Analytics (VA) [32], Keim et al. proposed that VA should *integrate scientific disciplines to improve the division of labor between human and machine*. By integrating human expertise through the human-computer interaction, VA systems aim to enable data experts to explore data graphically and generate insights more easily. However, as users grow dependent on the VA systems, new uncertainties and errors that the VA systems bring in might expose users to the risk of generating ill-informed insights. This would be detrimental for VA system – if users become aware of such uncertainties and errors, they might lose their trust in the VA system and stop using it; if users stay blind to the uncertainties and errors, the ill-informed insights they produced might cause them to make problematic decisions. Such issues coincide with previous trust research – trust is increasingly relevant under the conditions of uncertainty presence in the trustee (VA system), vulnerability to risk for the truster (user) and dependence relationship between the truster and the trustee [33].

Previous research on trust in visualization has mostly focused on the idea of trust building – essentially to improve users’ trust in VA systems [42, 57]. However, VA systems are designed always by humans and subject to potential human errors and subjectivity. Furthermore, one of the fundamental ideas in VA – human-in-the-loop – emphasizes that humans should supervise and steer the analytical process to generate trustworthy insights. Therefore, it is necessary and positive for users to maintain a healthy skepticism towards the VA system. In this position paper, we consequently propose that calibration of the appropriate trust level is equally important as, if not more than, trust building. With these concepts, we mean concretely:

Trust building increases the trust a user puts in a VA system through various means, such as making computations transparent through visualization (showing *what* the system is doing), providing explanations for results (showing *why* the system is doing it), and allowing the user to interject and reparameterize at any point.

*e-mail: wenkaihan@cs.au.dk

†e-mail: hjschulz@cs.au.dk

Trust calibration aligns the trust put into a VA system by the user with the system’s actual trustworthiness through various means, such as communicating uncertainties, providing visual cues and previews of the end result the user can expect from the system, and indicating analysis paths that have shown to work for similar data in the past.

In Sec. 2, we first lay out the trust continuum as a basis for the discourse of trust building and trust calibration. Then, Sec. 3 dissects potential trust issues in both VA systems and users and outlines possible approaches to build and calibrate trust. Sec. 4 subsequently connects some emerging VA approaches with the previous discussions of trust to inspect how they might bring new perspectives for the trust dynamics in VA. At last, Sec. 5 concludes this paper with some overarching insights and recommendations for future research regarding trust calibration for VA.

2 CONTINUUM OF TRUST

Trust building and calibration deal with trust issues from different but complementary angles. Trust building emphasizes increasing users’ trust level in VA systems, while trust calibration focuses on avoiding and mitigating misplaced levels of trust. This difference is illustrated by the trust continuum shown in Figure 1: where trust building aims to increase the trust level from left to right, trust calibration aims to align the trust from bottom to top. The elements of this continuum are introduced below.




Figure 1: The trust continuum extended on the model of Cho et al. [8]

2.1 The Foundation: Trust and Trustworthiness

The definition of *trust* varies in different contexts, but the general concept of trust is defined as “*the belief that the trustee will act in the best interests of the truster in a given situation.*” [41]. This captures the dynamics of the trustee and truster in a social relationship. However, this changes in the context of VA, as one of the two trust parties, the VA system, is a largely non-social actor.

In the situation of VA, the primary goal, i.e. the “best interests” of the truster (user), is to “*identify and visually distill the most valuable and relevant information content.*” [32] Therefore, we can adapt the definition of trust in the context of VA as *the truster (user)’s belief that the trustee (VA system) will help them correctly identify and visually distill the most valuable and relevant information content.*

Note that trust is slightly different from *trustworthiness*. While trust is a *belief* that is not necessarily based on observed evidence, *trustworthiness* is the verified and objective trust based on observations [62]. In the context of VA, we can think of trust as such belief that users might have about the VA system and that is possibly even preconceived and formed without actually ever having used the system. Whereas *trustworthiness* is based on the observation that the VA system helped users to achieve their goals and the expectation that the VA system will behave consistently in that regard.

2.2 Levels of Trust: Distrust, Untrust, and Undistrust

In addition to the state of full trust, there are three more levels of trust: distrust, untrust, and undistrust.

Distrust measures an active form of negative trust where the truster believes that the trustee will actively work against their interests [41], which can lead to disuse, i.e. abandonment, of digital systems [36]. When users distrust a VA system, they may have found that the VA system repeatedly produces inaccurate visualizations. While this may not be “malintent” by the system or its authors – e.g., in cases where complex data standards are not fully supported [59] – it can still hinder carrying out an analytic task consistently and free of errors. In the worst case, this causes users to no longer deem the system trustworthy and thus abandon it. On the brink of distrust, trust building is vital, as users are likely to disuse the VA system.

Untrust, on the other hand, indicates a state where the truster is *not fully confident* in the trustee, while being at the same time still inclined to trust it for the most part. For VA systems, it is natural for users, especially experts, to be alert and consider if there are any errors in the data, implicit assumptions in the computational process, or overplotted information in the visual representation. In particular, when users are not yet fully acquainted with a VA system, such considerations might help them to be aware of the implications of their analytic choices they may not yet be aware of.

Undistrust means the lack of trust [7], where the truster becomes *suspicious* of the VA system, but has not fully distrusted it. Compared to untrust, undistrust leans more towards the negative side, where the truster contemplates more to distrust the trustee. In the state of undistrust, users have serious doubts about the VA system and its trustworthiness, but they can still perform most of their intended tasks and generate some insights when using it with caution.

2.3 Misplaced Trust: Mistrust and Misdistrust

Mistrust and mistrust denote situations where trust or mistrust is *misplaced* compared to the trustee’s actual trustworthiness. In other words, the level of trust brought forth by the truster and the trustworthiness of the trustee are miscalibrated and the user’s expectations of the system do not align with what the system can actually provide.

Deriving from the notion of misinformation, *mistrust* is often called *misplaced trust* [40]. It arises when the truster gives a positive estimation of the trustee that later proves to be misplaced. This is particularly problematic, as mistrust can lead to misuse, i.e., users generating inaccurate results and gaining false insights, which works against their interests of using VA systems. Furthermore, later when users find out such mistakes, it is more possible for them to feel “betrayed” or “cheated” by the VA system and start distrusting it.

Defined as *misplaced distrust*, *misdistrust* is the counterpart of mistrust, where a truster distrusts a trustworthy trustee [43]. Misdistrust originates from miscommunication or misunderstanding between the user and VA system. Misdistrust is detrimental to the interaction, as users might disuse the VA system, when in fact the VA system can be trusted. Once mistrust has formed, it can eliminate the possibility of the VA system to later “redeem” itself.

Along the same lines, de Visser et al. proposed a trust calibration model between the level of trust and the actual trustworthiness [12]. When the trust level is higher than the actual trustworthiness, they speak of *over-trust*, whereas a lower trust level than the trustworthiness is termed *under-trust*. Note that there is a key difference between over-trust/under-trust and mistrust/misdistrust. Over-trust and under-trust can refer to any situation where the user’s trust is higher/lower than the actual trustworthiness of the VA system, even if not by much. This would be the case, when users generally untrust a system that may in fact not be fully trustworthy, but that could still be used with caution – and whose actual trustworthiness is thus on the undistrust level. However, mistrust and mistrust pinpoint the specific problematic scenarios where users trust or distrust a system that should not be trusted or distrusted, respectively.

2.4 The Bounds of Trust: Cooperation Threshold and Limit of Forgivability

The concepts of “*cooperation threshold*” and “*limit of forgivability*” were introduced by Marsh and Briggs [40]. They delineate trust and untrust, as well as distrust and undistrust, respectively.

Cooperation threshold refers to the point beyond which trust is established and the two parties will jointly proceed towards the same goal [40]. In a social context, *cooperation* means the action of different people working together, whereas in the context of VA system and the user, we define cooperation as fluent, reliable, and convergent interaction between system and user that work towards jointly identifying and distilling valuable and relevant information. Note that human usage of a VA system alone does not constitute as full cooperation, but that it requires the mutually dependent nature of the iterative human-in-the-loop interaction with each other.

Limit of forgivability refers to the limit beyond which the trustee is truly distrusted and can be considered only as acting against the truster’s best interests. According to Marsh and Briggs [40], this limit determines the worth of the trustee entering into redemption strategies to seek forgiveness from the truster. In the context of VA, we can see this as the limit beyond which a deeply disappointed user would abandon and disuse a VA system.

3 SHOULD I TRUST, AND WHAT TO TRUST?

VA provides users with powerful tools for understanding and reasoning. However, VA systems also confront users with computed results and mined patterns that stand in conflict with the user’s previous knowledge, experiences, and beliefs. This leads to a series of questions: “Should I trust myself or the VA system?”, “How much should I trust the VA system?”, and “Which part of the VA system should I trust more?” To answer these questions, analysts must know about the strengths and the weaknesses of both sides – the VA system and themselves – to know whom to trust in which situation. Thus, the following dissects potential trust issues on both sides, provides pointers to existing research for each, and details what can be done to build and calibrate trust in each case.

3.1 Should I Trust the VA System?

VA systems are designed by humans and therefore subject to potential human errors and subjectivity. Moreover, VA systems rarely have access to the “big picture” of the context behind a given analytic task. For example, a VA system does not know that reporting a computed result to the 10th digit after the comma miscommunicates a level of certainty and detail that is not warranted when averaging 5 roughly estimated numbers, leading to mistrust/over-trust in that result. In this section, we dissect how trust issues emerge in different parts of VA systems – data, computational process, visualization, and interaction – as well as what can be done to address these issues.

3.1.1 Should I trust the data?

“Garbage in, garbage out.” This principle captures the observation that the quality of the input to a digital system is directly reflected in the quality of the produced output. It also holds true for trust in VA: if the input data to a VA system are not trustworthy, then this lack of trustworthiness will propagate all the way to the derived insights.

As much as people label datasets as “raw”, such data are still collected through certain technical and social lenses. The “raw data” we obtained “are always already cooked and never entirely raw” [18], and thus raise questions of trust. National population census data in some countries are collected through investigators going into every household and might be subject to various human errors. Natural sciences researchers place sensors with varying accuracy in locations that they deem as reasonable to gather data for their research. Tech companies collect user data through their own algorithms, selecting data that are relevant to their field of business, easy to access, and legal to be collected. As such, even the data in their most original

forms are conceived before the collection process and limited by various technical and social constraints. When such conceptions and constraints are not communicated to the users of the data, inconsistencies in the “raw data” can be easily overlooked and lead to mistrust, or even be misconstrued as intentional manipulations and lead to misdistrust.

In a review paper on trust in digital information, Kelton et al. concluded that people tend to put more trust in accurate, up-to-date, complete information without deception and distortion, and is persistently obtainable with responsible methodology [33]. Therefore, to calibrate the trust to be placed on the input data, the inclusion of related information about the data source and communication of uncertainties in the collected data are essential [59]. Such metadata can inform users about where data discrepancies stem from and make users aware of the impact these discrepancies have on their analysis. Metadata make the processes transparent by which the data were gathered and further processed. However, such metadata are not proof of this process being the most suitable and they rarely explain why a particular process was chosen. Adding this reasoning behind them would further help to judge the data’s trustworthiness.

Yet we also need to communicate the metadata to the user to make a judgment of trust. Uncertainty visualization is a frequently mentioned approach to communicate quantitative uncertainties [4]. In theory, communicating such metadata should allow for better judgment of the data and thus of any processing result based on that data. In practice, though, it turns out that most users have a hard time to reason with uncertainties, let alone to parse the provided visualizations [29]. As for qualitative uncertainties originating from the process of data gathering and preprocessing, communicating the data provenance is an established approach [23]. Given the data provenance reflects a systematic and responsible methodology behind it, it has the potential to instill trust in users. In addition, such openness about the process behind the data can give an impression of “we have nothing to hide” and increase the trust level in general.

3.1.2 Should I trust the computational process?

The computational process in VA systems is like a black box ingesting data and producing results to be subsequently visualized. As such, it provides little to no internal status to understand its inner workings. Having little insight in and understanding of the computational process, it is almost inevitable for the users to start assuming “intentions” of a VA system – likely negative ones. Harboring such assumptions, users will actively look for instances where the system appears to work against them, which will eventually lead to distrust.

Many interactive visualization tools emphasize their integration with computational software such as MATLAB and R. However, as Mühlbacher et al. pointed out, such computational software is usually used as a black box that runs in isolation, providing no output other than the final result once it is ready and defeating the purpose of a visual-interactive data analysis [48]. More importantly, users have very limited knowledge of what is going on in the algorithms behind the scenes and limited agency over the process. When errors arise, users rarely have the option to probe into the computational processes to inspect the potential causes, therefore being unable to verify what went wrong and calibrate their trust level accordingly. A user trust study in intelligent systems by Holliday et al. found a similar pattern that without explanations of how the systems work, user trust might deteriorate over time, which is why the perceived transparency of the system becomes increasingly important for users to trust it [27]. Based on currently available computation infrastructure, Mühlbacher et al. subsequently proposed four different strategies to achieve user involvement [48], which in turn provide knowledge about the algorithms, insight into how they run, as well as agency to users to calibrate their trust levels.

In addition to user involvement and understanding, Friedman and Nissenbaum pointed out that technical and social constraints

can transfer into issues in computer systems [16]. In the context of computational processes in VA, algorithmic bias is a notable issue. Algorithmic bias touches on systematic errors in the algorithms that might create unfair results. Danks and London gave some good examples on such issues – the training data might be skewed due to moral or legal reasons, or the algorithm could be designed to counter overfitting noisy data but then ending up more biased in other scenarios [10]. If results from such biased algorithms are still consistent with the users’ expectations, they might end up mistrusting an actually untrustworthy computational process. It is therefore important to at least identify and communicate potential computational bias from the algorithms to calibrate trust. To cope with algorithmic bias, Cabrera et al. developed FAIRVIS [6] to aid discovering intersectional bias in machine learning and creating more equitable algorithmic systems. Such tools can be helpful to uncover and communicate algorithmic biases, helping to avoid mistrust.

3.1.3 Should I trust the visualization?

Visualization displays the results from the computational process to make it easier for the human user to gain insights. To do so, most VA systems provide a limited selection of different visual mapping and rendering techniques, and such techniques are very often not an accurate one-to-one mapping from the data space to the view space. While this is only natural in the age of big data where we have many more data points to plot than available pixels on our screens, it still misconstrues the data and is thus a potential cause of distrust.

Many visualizations are visually pleasing, which can help to build initial trust, especially with inexperienced users. However, if such visually pleasing graphs do not communicate the underlying data accurately and provide effective means to discover insights, such initial trust will sooner or later prove to be mistrust and eventually lead to distrust. It is therefore important to calibrate trust through providing some form of guidance that can help to avoid mistrust. Recommendation systems such as Tableau’s “Show Me” [38] and Moritz et al.’s Draco [47] can to some extent avoid “visualization design mirages” [45] by incorporating design knowledge and guidelines in their recommendations. Furthermore, visualization linting can help to uncover improper visual mappings. Similar to code linting, visualization linting searches for common visualization mistakes and automatically highlights them to help users recognize and potentially correct them [44].

The rendering of visualizations can also be an important trust factor. On one hand, technical constraints like low resolution and inadequate contrast might make it hard for users to clearly perceive the visualization, hindering them from gaining accurate insights [5]. On the other hand, some rendering techniques simply struggle to put all information in the available display space, which can lead to important information being hidden at subpixel resolution. To nevertheless point the user towards this information, Luboschik et al. have shown guidance to be a valuable means [37]. As they highlight display regions in which data at subpixel level deviate from the currently shown view, the VA system is transparent about its rendering limitations and users know where to zoom-in to find any deviations. This transparency aligns expectations and thus actively calibrates the trust in the VA system.

3.1.4 Should I trust the interaction?

Usability and user experience of the interaction with a VA system are important for the trustworthiness of it. Coherence is especially crucial for users to understand and trust the VA system, as discrepancies in the interaction might trigger users to scrutinize a digital system further [55]. When users take actions in a VA system, they have conscious or subconscious expectations of the system’s reactions. Discrepancies between these expectations and the provided reaction pose a threat to a VA system’s trustworthiness. A framework that reflects different forms of such mismatch is Tominski and Schumann’s

conceptual separations, spatial separations, and temporal separations regarding interaction costs with a VA system [66].

Conceptual separations concern the misalignment between the mental model that users have about the system, the implementation model the system adheres to, and the presented model of its interface. If the users' mental model does not match with the presented model, they might subsequently internalize such mismatch as an error in the system, pushing users to scrutinize and even distrust.

Spatial separations relate to the spatial placements and distances between different interactive elements and system reactions. This is problematic when the user's interaction and visual response from the system are inconsistent. Such inconsistency between users' spatial expectations and the actual spatial separations in the interface would make it harder for users to understand the action-effect causality of the VA system, causing confusion at best and mistrust at worst.

Temporal separations reflect the latency between a user's action and the system's visual response. Users might have some expectations of the duration of certain internal processes, and when the actual latency drastically deviates from their expectations, they will become suspicious of the system and underlying process.

For example, in coordinated multiple views, users' actions in one view are expected to influence several others. Yet, if this influence is not clearly represented across the different views (spatial separation), or the actions take too long to propagate to other views (temporal separation), users might not be able to understand which of their actions impacted in which ways the other views and misinterpret the underlying logic (conceptual separation), leading to mistrust.

A way to counter such trust miscalibration is to communicate the system's response and latency regarding users' possible interactions. For conceptual separations, scented widgets [70] can serve as a preview to align users' expectations of their actions with the reactions from the systems by adding cues to the corresponding interactive elements. For large spatial separation between users' action and systems' reaction, visual links [65], arrows or highlights can enable users to follow the action-effect causality and instill trust. Regarding temporal separations, providing estimates of computationally intensive actions in either textual or visual forms can help to calibrate user expectation and trust. However, inaccurate estimates can also create even more discrepancies and induce distrust.

3.2 Should I Trust Myself?

When users interact with VA systems, variations in their perception, knowledge, judgment, and situational state influence their actions. These factors are essential for trust calibration: On one hand, they may be the reason users misplace their trust or distrust in the first place – e.g., because of their confirmation bias, users trust results more if these align with their beliefs. On the other hand, they can interfere when trying to communicate uncertainties or algorithmic details – e.g., when change blindness makes it hard to follow computational updates. Hence, in this section, we dissect how these human factors are related to trust calibration and which means have been proposed to alleviate the issues they cause.

3.2.1 Should I trust my perception?

Perceptual factors, such as visual expectation, visual memory as well as visual attention, are important in visualization, as it is the foundation of human sensemaking from large and often complex datasets. To calibrate trust, we need to consider if one can perceive visual information true to what the VA system present. To this end, not only should we be aware that there can be a spectrum of perceptual abilities among users, for example, different degrees and types of colorblindness, dyslexia, or autism. We also need to consider that human perception is far from optimal and error-free, as it is evident by the broad range of visual illusions.

Among the perceptual abilities, visual abilities are relatively well-researched. Thus, it is well-known that sensitivity to color dete-

riorates with age and colorblindness can also seriously limit the quantity and quality of information we can extract from visual representations [60]. Many tools have thus put color perception into consideration to make sure one can trust what one perceives. For example, ColorBrewer specifically enables choosing only colorblind-safe color scales [21]. VisCheck and Daltonize show how a visualization or user interface looks like for users with different kinds of colorblindness and provide corrections [15]. Other perceptual differences like synesthesia or dyslexia have not been the focus of dedicated studies in visualization. However, they can be expected to also impact the perceptual process in VA systems, as underlined by recent work on a "synesthetic color palette" [56].

Regardless of individual predisposition, perceptual errors such as change blindness or line width and sine illusions arise for any user perceiving visualizations. To achieve trustful visual communication between the VA system and the user, they thus need to be considered. Change blindness occurs when people do not notice changes in visible elements of a scene. In the context of VA, users might not be aware of how an animated visualization changes or how a static visualization updates. This in turn makes it hard to judge the trustworthiness of the system with up-to-date information. Nowell et al. discussed possible solutions using morphing, crossfading, and wireframes to draw attention to regions of change in the view space [50]. In addition, due to human's tendency to perceive distance between curves as the minimal distance rather than the vertical distance, line width and sine illusions are widely discussed in statistical graphics literature, especially when representing areas between two curves. Hofmann and Vendettuoli proposed *Common Angle Plots* to address line width illusion [26], and VanderPlas and Hofmann demonstrated possible solutions to counter the sine illusion [67].

3.2.2 Should I trust my knowledge?

Running an analysis with a VA system, users internalize the resulting information they yield from the system through their existing construct of knowledge. Users with different expert knowledge understand and interact with VA systems differently [49]. In particular, a lack of knowledge leads to uninformed actions that potentially cause misunderstanding of the systems and miscalibration of trust.

Domain knowledge about the analyzed dataset is vital for sensemaking in VA systems. The human sensemaking process is based on framing the data presented by the VA system with their existing knowledge construct. As Klein et al. pointed out, "sensemaking is a process of framing and reframing", that fits presented data into the analysts' knowledge construct [34]. On one hand, users are inclined to trust data that fits with their framing and distrust one that does not, which can lead to mistrust and mistrust. On the other hand, trust calibration is also weakened when users do not have enough domain knowledge to judge if they should trust the outputs from the VA system or their own framing. Implementing a form of Analysis of Competing Hypotheses (ACH) can be helpful for mitigating such issues of data-specific domain knowledge. ACH refers to an analytical process to aid decision-making regarding issues with different alternative explanations or conclusions [24]. Enabling users to explore several analytical paths can help to validate different framing of the data and ensure a calibrated level of trust.

In addition to domain knowledge, users' knowledge, or expertise level about VA and the specific VA system has an important impact on the users' ability to take appropriate analytical actions. Lack of knowledge about different computational processes might leave them in a trial-and-error mode when aiming to choose one that is consistent with their intentions; insufficient navigation skills around the VA system might make it increasingly hard for users to discover different options and views that would help them to generate more insights; inexperienced users might not be able to spot errors and understand issues arising in the system and take actions accordingly. These issues hinder a smooth interaction with the VA system, which

impedes the users' perception of the system being truthful and their trust in their own actions. To mitigate the lack of domain and VA knowledge, knowledge-assisted visualization has been proposed to help users navigate through different methods, parameters, and visualization techniques. For example, Jänicke and Scheuermann built a knowledge-assisted visualization for time-dependent multivariate flow datasets, in which users can store process knowledge to aid later analyses [30]. Their user study also shows that knowledge can be extracted and transferred to novice users with this approach.

3.2.3 Should I trust my judgment?

Human decision-making underlies errors and differences in judgment. For example, we tend to seek meaning in things and interpret things within our own experiences, often seeing patterns where there are none. This phenomenon is called *apophenia*, and being aware of it and working actively against it is a skill that is hard to come by [35]. Furthermore, facing the same VA system, different people make different judgments, as they look at the system through their own lenses of reality with different personal traits, habits, and behavioral patterns. Such deviations can interfere with trust calibration.

As part of one's subjective construct of reality, cognitive bias is a systematic deviation from rational judgment caused by the use of heuristics in decision making [28]. A taxonomy of cognitive biases for information visualization by Dimara et al. lists and classifies 154 cognitive biases [14]. For example, *confirmation bias* will have users subconsciously look for evidence that is in line with any prior assumptions, while ignoring findings that contradict their assumptions [51]. Seeing a lot of confirming evidence in a VA system can lead to mistrusting it. This is a clear miscalibration, as the system always shows the full story, but the user only pays attention to one side of it. To address selection bias, Gotz et al. [20] showed the similarity of a selected data subset to the full dataset to ensure the selection is representative. Dimara et al. [13] highlighted optimal choices and altered task framing to mitigate attraction effects. Wall et al. proposed real-time metrics to detect bias [68] and outlined a design space for mitigating bias in VA systems [69].

Moreover, many personal factors, such as culture, gender, and personality can heavily influence how one absorbs information and evaluate trustworthiness. For example, people from cultures with high uncertainty avoidance, such as Greece, Portugal, and Poland, tend to make unnecessarily conservative evaluations, while members of low uncertainty avoidance cultures, such as Singapore, Hong Kong, and Sweden, are more prone to take risky actions [17], which might include trusting something more easily. Regarding personality, locus of control (LOC), which measures the degree to which one feels in control of or controlled by external events [53], is one of the more well-studied personality traits. In the context of VA, Ziemkiewicz et al. found that users with external LOC are able to efficiently complete VA tasks even with unfamiliar visualizations like an inclusion hierarchy, while internal LOC users struggle to do so if not presented with a familiar node-link drawing [71]. To adapt to different personal traits, personalized visualization offers a way to create interfaces that cater to the diversities among users [52], which better align with what different users expect and need, and are thus less likely to lead to miscalibrated trust.

3.2.4 Should I trust my situational state?

When users perform a data analysis, their situational state is an underlying factor to their decision-making. Their current internal mood as well as external environment can influence how they perceive, understand, and take actions. Thus, the users' situational state also affects the trustworthiness of their actions and insights.

Regarding mood, studies in decision-making show that negative moods are associated with exploration rather than exploitation behaviors, as well as introducing changes rather than maintaining the status quo [58], among others [54]. This also impacts the use of

VA systems. For example, using VA systems with a negative mood makes the user more likely to introduce unnecessary changes and to explore the visualization for an extended amount of time.

Environmental factors include physical as well as mental ones. Although VA tasks are usually performed in a consistent indoor environment, it not necessarily optimal – especially social factors such as shared offices, interruptions and distractions can make a focused, in-depth analysis session almost impossible. This leads in turn to an increase in perceptual errors, and an inattentiveness to one's own biases and algorithmic biases alike. External mental factors can also have an impact. Risk is an important mental factor among other external situational variables for regulating trust behavior. When users need to make a high-risk decision, they tend to rely on more trustworthy cues and tools [64]. Thus, for a “high-stake analysis” whose outcome will be of great influence – e.g., a trader's investment decision for a fund or a clinician's treatment decision for a patient – analysts are likely to choose methods they have more knowledge about and feel safer with. But this also makes them more likely to misuse methods that do not fit the problem at hand.

Tracking physical parameters like eye gaze and electroencephalogram has been proven useful in gauging users' internal situational factors like mood and intentions [61]. VA systems can also ask questions about users' intentions or external situations before they enter the analysis process. Such information can be used to adapt the VA system to the situational factors.

4 EMERGING FACES OF TRUST

Over the past years, a number of different “flavors of VA” have emerged that introduce new possibilities to the generic VA process that goes back and forth between human and computer. As these emerging approaches have implications on trust building and calibration in VA, we briefly discuss three of them in the following.

Progressive VA (PVA) carries out analytic computations in a step-wise manner on data subsets (so-called *chunks*) in order to visualize and interact with partial results already before the full computation is finished [2]. Researchers list building trust as one of the biggest benefits of adopting PVA, as by communicating the progression of the underlying process, users' gain an understanding of how results are generated. It thus enables building and even calibrating users' trust in the computational process [46]. Yet it also opens the question of how much can one trust the shown intermediate partial result [1]? Since this question is very hard to answer, Jo et al. developed a different approach to put a user's mind at rest: their PVA system ProReveal [31] incorporates safeguards that can be attached to a running progressive computation and formulate a hypothesis about the computation result as a conditional expression. As long as this conditional holds true, the user can move on with the analysis process – but the moment it is no longer true (e.g., because new data has meanwhile been processed that contradicts the hypothesis) the user is notified. Their user study validates that safeguards can alleviate the unsure feelings users have about early and intermediate results. This makes for a very interesting case of temporal separation (cf. Sec. 3.1.4) where one can only fully trust a result once the computation is completed. Yet making use of PVA's inherent ability to continue the analysis already from a good enough partial result, this point of full trust still lies in the future. Safeguard effectively resolve this separation, as they allow moving on with the analysis, even while still not fully trusting the partial result.

Mixed Initiative VA extends VA into a discourse where human and computer are more on par with each other. To that end, a Mixed Initiative VA system infers the users' potential intentions and likely analytic goal from their interactions with the system, so as to proactively support these intentions and goals. This support can range from automatically setting suitable defaults for parameters, to the system offering guidance on how to achieve those analytic goals [9]. An empirical study by Dasgupta et al. found that for

complex sensemaking tasks, Mixed Initiative VA systems can inspire greater trust [11]. In addition to building more trust, Mixed Initiative VA also provides useful tools to calibrate users' trust in themselves and the VA systems. By learning about the intentions of the users, VA systems can adapt to better meet the inferred expectations and needs of the users [39], which Sperrle et al. defined as a co-adaptive guidance process [63]. This is essentially a communication process between VA systems and users to calibrate trust. Relating this dynamic to the trust continuum, Mixed Initiative VA opens up a different direction of trust from VA system to user: the VA system becomes the truster, the user becomes the trustee, and the VA system has to trust the users know what they are doing and behave rationally in order to correctly infer their intentions. Yet as the human user is often irrational, one can already see how miscalibration of trust in the user is a huge challenge in Mixed Initiative VA.

Collaborative VA extends VA from one to multiple analysts, potentially with different backgrounds and expertise, performing the analysis together [22]. Collaboration has been proven to be useful in bringing in diverse perspectives and mitigate individual's limitation of knowledge and cognitive bias. By communicating knowledge, experience and different perspectives between each other, users will be exposed to more new ideas and are therefore more likely to break their habitual behaviors. This can help users to gain a more comprehensive understanding of the information in the VA systems, and therefore calibrate their trust level. Billman et al. conducted a series of empirical studies on collaborative intelligence analysis [3]. They found a reduction in confirmation bias for heterogeneous groups of people with diverse beliefs when using collaborative systems. However, for homogeneous groups with similar beliefs, their initial biases were accentuated. Therefore, to ensure trustful decisions, it is important to promote collaboration with heterogeneous groups of users to make sure that diverse opinions and inputs will be considered.

5 CONCLUSIONS

While research dealing explicitly with trust building has been rare in the field of VA, work that emphasizes trust calibration in VA is even rarer. Inspired by work in related research fields such as automation [12, 25, 36] and intelligent systems [27], we make a clear distinction between trust building and trust calibration, and bring attention to the latter for matching users' perceived trust and the actual trustworthiness of VA systems. Admittedly, trust building is essential to avoid distrust situations where users might abandon the VA systems. However, building trust that is higher than the actual trustworthiness of the VA systems might set user expectations too high, leading users to blindly trust the system, which will result in disappointment sooner or later. This is precisely the point of trust calibration, which aims to find the appropriate trust level for a VA system and dataset at hand. Trust calibration can in most instances be understood as a form of communication between system and human user in which expectations are aligned to avoid disappointments.

In this paper, we established the importance of trust calibration through the conceptual space of a trust continuum and discussed it for VA systems and users. However, much more research needs to be done to gain a more comprehensive understanding of trust calibration in VA. To begin with, trust building and calibration can stand in conflict with each other when the actual trustworthiness of a VA system is low, and thus building perceived trust would actively miscalibrate it. Therefore, it is important to consider and investigate how trust building and calibration should coexist. Furthermore, although there has been some research on evaluating trust levels [11, 27], tracking trust calibration can be a dynamic process that requires continuous monitoring of trust and trustworthiness. How to evaluate trust calibration is therefore an important but complicated question to address. Last but not least, as new VA approaches emerge, trust calibration can become more intricate – PVA brings up additional trust issues when working with incomplete results, Mixed Initiative

VA starts to asks about the trustworthiness of users, and collaborative VA introduces interpersonal trust to the VA process. Both theoretical and empirical research is needed to fully dissect and investigate the trust dynamics in corresponding VA approaches.

REFERENCES

- [1] M. Angelini, T. May, G. Santucci, and H.-J. Schulz. On quality indicators for progressive visual analytics. In *Proc. of EuroVA*, pp. 25–29. Eurographics, 2019.
- [2] M. Angelini, G. Santucci, H. Schumann, and H.-J. Schulz. A review and characterization of progressive visual analytics. *Informatics*, 5(3):31, 2018.
- [3] D. Billman, G. Convertino, J. Shrager, P. Pirolli, and J. Massar. Collaborative intelligence analysis with CACHE and its effects on information gathering and cognitive bias. In *Human Computer Interaction Consortium Workshop*, 2006.
- [4] G.-P. Bonneau, H.-C. Hege, C. R. Johnson, M. M. Oliveira, K. Potter, P. Rheingans, and T. Schultz. Overview and state-of-the-art of uncertainty visualization. In C. D. Hansen, M. Chen, C. R. Johnson, A. E. Kaufman, and H. Hagen, eds., *Scientific Visualization*, pp. 3–27. Springer, 2014.
- [5] K. Brodlić, R. A. Osorio, and A. Lopes. A review of uncertainty in data visualization. In J. Dill, R. Earnshaw, D. Kasik, J. Vince, and P. C. Wong, eds., *Expanding the frontiers of visual analytics and visualization*, pp. 81–109. Springer, 2012.
- [6] Á. A. Cabrera, W. Epperson, F. Hohman, M. Kahng, J. Morgenstern, and D. H. Chau. Fairvis: Visual analytics for discovering intersectional bias in machine learning. In *Proc. of VAST*, pp. 46–56. IEEE, 2019.
- [7] C. Castelfranchi, R. Falcone, and E. Lorini. A non-reductionist approach to trust. In Golbeck [19], pp. 45–72.
- [8] J. H. Cho, K. Chan, and S. Adali. A survey on trust modeling. *ACM Computing Surveys*, 48(2):28, 2015.
- [9] K. Cook, N. Cramer, D. Israel, M. Wolverton, J. Bruce, R. Burtner, and A. Endert. Mixed-initiative visual analytics using task-driven recommendations. In *Proc. of VAST*, pp. 9–16. IEEE, 2015.
- [10] D. Danks and A. J. London. Algorithmic bias in autonomous systems. In *Proc. of IJCAI*, pp. 4691–4697. AAAI Press, 2017.
- [11] A. Dasgupta, J. Lee, R. Wilson, R. A. Lafrance, N. Cramer, K. Cook, and S. Payne. Familiarity vs trust: A comparative study of domain scientists' trust in visual analytics and conventional analysis methods. *IEEE TVCG*, 23(1):271–280, 2017.
- [12] E. de Visser, M. Cohen, A. Freedy, and R. Parasuraman. A design methodology for trust cue calibration. In *Proc. of the Intl. Conference on Virtual, Augmented and Mixed Reality (VAMR)*, pp. 251–262. Springer, 2014.
- [13] E. Dimara, G. Bailly, A. Bezerianos, and S. Franconeri. Mitigating the attraction effect with visualizations. *IEEE TVCG*, 25(1):850–860, 2018.
- [14] E. Dimara, S. Franconeri, C. Plaisant, A. Bezerianos, and P. Dragicevic. A task-based taxonomy of cognitive biases for information visualization. *IEEE TVCG*, 26(2):1413–1432, 2018.
- [15] B. Dougherty and A. Wade. Vischeck simulates colourblind vision, 2006. <http://www.vischeck.com>, Retrieved 23-AUG-2020.
- [16] B. Friedman and H. Nissenbaum. Bias in computer systems. *ACM TOIS*, 14(3):330–347, 1996.
- [17] B. Frijns, A. Gilbert, T. Lehnert, and A. Tourani-Rad. Uncertainty avoidance, risk tolerance and corporate takeover decisions. *Journal of Banking & Finance*, 37(7):2457–2471, 2013.
- [18] L. Gitelman. *Raw data is an oxymoron*. MIT Press, 2013.
- [19] J. Golbeck, ed. *Computing with Social Trust*. Springer, 2009.
- [20] D. Gotz, S. Sun, and N. Cao. Adaptive contextualization: Combating bias during high-dimensional visualization and data selection. In *Proc. of IUI*, pp. 85–95. ACM, 2016.
- [21] M. Harrower and C. A. Brewer. Colorbrewer.org: an online tool for selecting colour schemes for maps. *The Cartographic Journal*, 40(1):27–37, 2003.
- [22] J. Heer and M. Agrawala. Design considerations for collaborative visual analytics. *Information visualization*, 7(1):49–62, 2008.

- [23] M. Herschel, R. Diestelkämper, and H. B. Lahmar. A survey on provenance: What for? what form? what from? *The VLDB Journal*, 26(6):881–906, 2017.
- [24] R. J. Heuer Jr. Analysis of competing hypotheses. In *Psychology of Intelligence Analysis*, pp. 95–110. US Govt. Printing Office, 1999.
- [25] K. A. Hoff and M. Bashir. Trust in automation: Integrating empirical evidence on factors that influence trust. *Human Factors*, 57(3):407–434, 2015.
- [26] H. Hofmann and M. Vendettuoli. Common angle plots as perception-true visualizations of categorical associations. *IEEE TVCG*, 19(12):2297–2305, 2013.
- [27] D. Holliday, S. Wilson, and S. Stumpf. User trust in intelligent systems: A journey over time. In *Proc. of IUI*, pp. 164–168. ACM, 2016.
- [28] K. J. Holyoak and R. G. Morrison. *The Cambridge handbook of thinking and reasoning*. Cambridge University Press, 2005.
- [29] J. Hullman, X. Qiao, M. Correll, A. Kale, and M. Kay. In pursuit of error: A survey of uncertainty visualization evaluation. *IEEE TVCG*, 25(1):903–913, 2018.
- [30] H. Jänicke and G. Scheuermann. Steady visualization of the dynamics in fluids using ϵ -machines. *Computers & Graphics*, 33(5):597–606, 2009.
- [31] J. Jo, S. L’Yi, B. Lee, and J. Seo. Proreveal: Progressive visual analytics with safeguards. *IEEE TVCG*. to appear.
- [32] D. Keim, G. Andrienko, J.-D. Fekete, C. Görg, J. Kohlhammer, and G. Melançon. Visual analytics: Definition, process, and challenges. In A. Kerren, J. T. Stasko, J.-D. Fekete, and C. North, eds., *Information visualization*, pp. 154–175. Springer, 2008.
- [33] K. Kelton, K. R. Fleischmann, and W. A. Wallace. Trust in digital information. *Journal of the American Society for Information Science and Technology*, 59(3):363–374, 2008.
- [34] G. Klein, J. K. Phillips, E. L. Rall, and D. A. Peluso. A data-frame theory of sensemaking. In *Expertise out of context: Proc. of the Conference on Naturalistic Decision Making*, pp. 113–155. Lawrence Erlbaum Assoc., 2007.
- [35] C. Kozyrkov. Your dataset is a giant inkblot test. Towards Data Science, 2017. <https://towardsdatascience.com/your-dataset-is-a-giant-inkblot-test-b9bf4c53ec5>, Retrieved 23-AUG-2020.
- [36] J. D. Lee and K. A. See. Trust in automation: Designing for appropriate reliance. *Human factors*, 46(1):50–80, 2004.
- [37] M. Luboschik, C. Maus, H.-J. Schulz, H. Schumann, and A. M. Uhrmacher. Heterogeneity-based guidance for exploring multiscale data in systems biology. In *Proc. of BioVis*, pp. 33–40. IEEE, 2012.
- [38] J. Mackinlay, P. Hanrahan, and C. Stolte. Show me: Automatic presentation for visual analysis. *IEEE TVCG*, 13(6):1137–1144, 2007.
- [39] S. Makonin, D. McVeigh, W. Stuerzlinger, K. Tran, and F. Popowich. Mixed-initiative for big data: The intersection of human + visual analytics + prediction. In *Proc. of HICSS*, pp. 1427–1436. IEEE, 2016.
- [40] S. Marsh and P. Briggs. Examining trust, forgiveness and regret as computational concepts. In Golbeck [19], pp. 9–43.
- [41] S. Marsh and M. R. Dibben. Trust, untrust, distrust and mistrust – an exploration of the dark(er) side. In *Trust Management: Proc. of iTrust*, pp. 17–33. Springer, 2005.
- [42] E. Mayr, N. Hynek, S. Salisu, and F. Windhager. Trust in information visualization. In *Proc. of TrustVis*, pp. 25–29. Eurographics, 2019.
- [43] B. McGuinness and A. Leggatt. Factors influencing information trust and distrust in a sensemaking task. In *Proc. of the International Command and Control Research and Technology Symposium*, 2006.
- [44] A. McNutt and G. Kindlmann. Linting for visualization: Towards a practical automated visualization guidance system. In *Proc. of the 2nd VisGuides Workshop*, 2018.
- [45] A. McNutt, G. Kindlmann, and M. Correll. Surfacing visualization mirages. In *Proc. of CHI*, pp. 1–16. ACM, 2020.
- [46] L. Micallef, H.-J. Schulz, M. Angelini, M. Aupetit, R. Chang, J. Kohlhammer, A. Perer, and G. Santucci. The human user in progressive visual analytics. In *Short Paper Proc. of EuroVis*, pp. 19–23. Eurographics, 2019.
- [47] D. Moritz, C. Wang, G. L. Nelson, H. Lin, A. M. Smith, B. Howe, and J. Heer. Formalizing visualization design knowledge as constraints: Actionable and extensible models in Draco. *IEEE TVCG*, 25(1):438–448, 2019.
- [48] T. Mühlbacher, H. Piringer, S. Gratzl, M. Sedlmair, and M. Streit. Opening the black box: Strategies for increased user involvement in existing algorithm implementations. *IEEE TVCG*, 20(12):1643–1652, 2014.
- [49] B. M. Muir. Trust between humans and machines, and the design of decision aids. *International journal of man-machine studies*, 27(5-6):527–539, 1987.
- [50] L. Nowell, E. Hetzler, and T. Tanasse. Change blindness in information visualization: a case study. In *Proc. of InfoVis*, pp. 15–22. IEEE, 2001.
- [51] A. Nussbaumer, K. Verbert, E.-C. Hillemann, M. A. Bedek, and D. Albert. A framework for cognitive bias detection and feedback in a visual analytics environment. In *Proc. of EISIC*, pp. 148–151. IEEE, 2016.
- [52] A. Ottley. *Adaptive and Personalized Visualization*. Morgan and Claypool, 2020.
- [53] E. J. Phares. *Locus of control in personality*. General Learning Press, 1976.
- [54] R. Raghunathan and M. T. Pham. All negative moods are not equal: Motivational influences of anxiety and sadness on decision making. *Organizational behavior and human decision processes*, 79(1):56–77, 1999.
- [55] M. M. Roghanizad and D. J. Neufeld. Intuition, risk, and the formation of online trust. *Computers in Human Behavior*, 50:489–498, 2015.
- [56] R. Rouw and N. B. Root. Distinct colours in the ‘synaesthetic colour palette’. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 374(1787):20190028, 2019.
- [57] D. Sacha, H. Senaratne, B. C. Kwon, G. Ellis, and D. A. Keim. The role of uncertainty, awareness, and trust in visual analytics. *IEEE TVCG*, 22(1):240–249, 2016.
- [58] B. Scheibehenne, B. Von Helversen, and Y. Shevchenko. Change and status quo in decisions with defaults: The effect of incidental emotions depends on the type of default. *Judgment and Decision making*, 9(3):287–296, 2014.
- [59] H.-J. Schulz, T. Nocke, M. Heitzler, and H. Schumann. A systematic view on data descriptors for the visual analysis of tabular data. *Information Visualization*, 16(3):232–256, 2017.
- [60] S. Silva, B. S. Santos, and J. Madeira. Using color in visualization: A survey. *Computers & Graphics*, 35(2):320–333, 2011.
- [61] G. Slanzi, J. A. Balazs, and J. D. Velásquez. Combining eye tracking, pupil dilation and EEG analysis for predicting web users click intention. *Information Fusion*, 35:51–57, 2017.
- [62] B. Solhaug, D. Elgesem, and K. Stølen. Why trust is not proportional to risk. In *Proc. of ARES*, pp. 11–18. IEEE, 2007.
- [63] F. Sperrle, A. V. Jeitler, J. Bernard, D. A. Keim, and M. El-Assady. Learning and teaching in co-adaptive guidance for mixed-initiative visual analytics. In *Proc. of EuroVA*, pp. 61–65. Eurographics, 2020.
- [64] D. Sprague and M. Tory. Exploring how and why people use visualizations in casual contexts: Modeling user goals and regulated motivations. *Information Visualization*, 11(2):106–123, 2012.
- [65] M. Steinberger, M. Waldner, M. Streit, A. Lex, and D. Schmalstieg. Context-preserving visual links. *IEEE TVCG*, 17(12):2249–2258, 2011.
- [66] C. Tominski and H. Schumann. *Interactive Visual Data Analysis*. AK Peters Visualization Series. CRC Press, 2020.
- [67] S. VanderPlas and H. Hofmann. Signs of the sine illusion—why we need to care. *Journal of Computational and Graphical Statistics*, 24(4):1170–1190, 2015.
- [68] E. Wall, L. M. Blaha, L. Franklin, and A. Endert. Warning, bias may occur: A proposed approach to detecting cognitive bias in interactive visual analytics. In *Proc. of VAST*, pp. 104–115. IEEE, 2017.
- [69] E. Wall, J. Stasko, and A. Endert. Toward a design space for mitigating cognitive bias in vis. In *Short Paper Proc. of VIS*, pp. 111–115. IEEE, 2019.
- [70] W. Willett, J. Heer, and M. Agrawala. Scented widgets: Improving navigation cues with embedded visualizations. *IEEE TVCG*, 13(6):1129–1136, 2007.
- [71] C. Ziemkiewicz, R. J. Crouser, A. R. Yauilla, S. L. Su, W. Ribarsky, and R. Chang. How locus of control influences compatibility with visualization style. In *Proc. of VAST*, pp. 81–90. IEEE, 2011.