The Role of Uncertainty, Awareness, and Trust in Visual Analytics

Dominik Sacha, Hansi Senaratne, Bum Chul Kwon, *Member, IEEE*, Geoffrey Ellis, and Daniel A. Keim, *Member, IEEE*

Fig. 1: Knowledge generation model for visual analytics including uncertainty propagation and human trust building. Uncertainty originates at the data source and propagates through the system components which introduce additional uncertainties. Uncertainty awareness influences human trust building on different knowledge generation levels.

Abstract— Visual analytics supports humans in generating knowledge from large and often complex datasets. Evidence is collected, collated and cross-linked with our existing knowledge. In the process, a myriad of analytical and visualisation techniques are employed to generate a visual representation of the data. These often introduce their own uncertainties, in addition to the ones inherent in the data, and these propagated and compounded uncertainties can result in impaired decision making. The user's confidence or trust in the results depends on the extent of user's awareness of the underlying uncertainties generated on the system side. This paper unpacks the uncertainties that propagate through visual analytics systems, illustrates how human's perceptual and cognitive biases influence the user's awareness of such uncertainties, and how this affects the user's trust building. The knowledge generation model for visual analytics is used to provide a terminology and framework to discuss the consequences of these aspects in knowledge construction and though examples, machine uncertainty is compared to human trust measures with provenance. Furthermore, guidelines for the design of uncertainty-aware systems are presented that can aid the user in better decision making.

Index Terms—Visual Analytics, Knowledge Generation, Uncertainty Measures and Propagation, Trust Building, Human Factors

1 INTRODUCTION

In the visual analytics process, users arrive at new knowledge after performing numerous sensemaking activities. The goal of visual analytics is to foster effective collaboration between human and machine that improves the knowledge generation process. To succeed in this process, end users need to be able to trust their knowledge generated by means of visual analytics. Analysts can often be unaware of uncertainties in their data sources, pre-processing, analysis processes or visualisations that are hidden by a 'black box' approach of visual analytics systems.

In criminal investigation analysis, where analysts use a visual analytics application to analyse a collection of reports and to identify crime suspects, the system may hint at otherwise hidden connections between pieces of evidence using a trained machine learning algorithm. To progress, the analyst needs to trust this outcome. However, if the analyst is not aware of the inherent uncertainties, they may waste their time following wrong leads and may, in the worst case, incriminate innocent people. Likewise, overestimating uncertainties can have a

- *Data Analysis and Visualisation Group, University of Konstanz. E-mail: forename.lastname@uni-konstanz.de*

negative impact upon decision making. It is therefore crucial for users to be provided with an accurate estimation of uncertainties from visual analytics systems so that they can trust acquired knowledge.

The literature describes some parts of uncertainty propagation and trust building in visual analytics processes, however, the interplay of trust and knowledge within the knowledge generation process in visual analytics has not yet been established. Prior studies have investigated sources of uncertainties in subsets of the visualisation process (e.g., [17]). Other studies have looked at human analysts behaviours while building trust in the knowledge generation process, with respect to perception [79], cognitive biases [31], and analytic roadblocks [48]. What is missing is a unified framework that bridges the concepts of uncertainties on the machine side and the trust building process on the human side. Recently, the IEEE VIS2014 Workshop on Provenance for Sensemaking called for research in defining uncertainty, trust, and data quality. MacEachren also highlighted human's decision making and reasoning processes under uncertainty as future research direction [52]. Building such a framework can provide a common language of the concepts that are largely uncharted in the visualisation domain.

Our goal is to investigate uncertainty propagation, trust building, and the interplay between uncertainty and trust during the knowledge generation process within visual analytics. Building on the related work in *Uncertainty Propagation* and *Human Trust Building under Uncertainty*, the paper describes a novel model of uncertainty and trust using the knowledge generation model [64] as a framework and brings

in human cognition and perception issues through the concept of *Awareness*. We choose the most recent and complete knowledge generation model for visual analytics because it "integrates human thinking whilst describing visual analytics components" [64]. Furthermore, this model was the foundation for our initial investigations and discussions about defining and relating its concepts to uncertainty and trust. To extend the usefulness of the model, we provide guidelines on how to improve decision making, avoiding misperceptions and pitfalls generated in the visual analytic processes. Finally, we explore future directions and opportunities for handling uncertainties and trust.

2 RELATED WORK

We group related work into *Uncertainty Propagation* and *Human Trust Building under Uncertainty*. The former covers works on capturing and deriving uncertainty measures within visual analytics pipelines. The latter covers the knowledge generation and the trust building processes in visual analytics.

2.1 Uncertainty Propagation

In this section we give a brief overview on *Types of Uncertainty*, *Uncertainty Propagation*, and *Data Provenance* theories.

Types of Uncertainty: Many works have individually tackled uncertainties that arise through the various components of a system.

Source: Data source uncertainty is inherent in the data. Lush et al. [50] introduced the GEOLabel[1], by which users can rapidly ascertain the source uncertainty of particularly geospatial datasets in terms of metadata (e.g., dataset producer information, compliance to international standards, ratings) through specially designed icons.

Models: Model uncertainty corresponds to the structure of the model and the parameterisation of the model. Chatfield [14] describes how uncertainty is fundamentally propagated in data models that represent real-world phenomena. He also describes the main sources of uncertainty in models, which we discuss in Section 3.1. Cullen and Frey [18] comprehensively address the methods for variability and uncertainty in models, while Lee and Chen [49] comparatively analyse the various uncertainty propagation methods in terms of their performance.

Uncertainty Propagation: Uncertainty is created and passed on from the source to the model and subsequently to the visualisation. Haber and McNabb [33] introduced uncertainty propagation to their visualisation reference model, where the visualisation of uncertainty focuses on the uncertainties that are in the measurement and simulation data (referred to as data source uncertainty, henceforth). They discuss how uncertainty propagates from the filtering stage, mapping stage, to the rendering stage of a traditional pipeline model. They call this uncertainty of visualisation. Uncertainty propagation within the context of visual analytics is the process of quantifying the underlying uncertainties generated throughout their components. In introducing a framework for uncertainty-aware visual analytics, Correa et al. [17] suggest propagating and communicating the uncertainties that arise inherently in the data and its transformations in the information visualisation pipeline. Furthermore, Zuk and Carpendale [83] extend the data uncertainty visualisation pipeline of Pang et al. [61] to include these propagated uncertainties. These workflows facilitate the analyst in identifying the inherent and propagated uncertainties in their data.

Visualisation of Uncertainty: A large body of work has contributed towards visualising the propagated uncertainties (e.g., [35, 55, 72]). Based on their extensive research, MacEachren [51], and Howard and MacEachren [37] introduced several dichotomous categories for uncertainty visualisation, based on the principle that, the way in which data and uncertainties are linked, should be meaningful. These act as guidelines in designing the visualisations appropriately for the data and task at hand. Several researchers have evaluated uncertainty visualisation techniques in various data/task settings (e.g., [68]). However, very little effort has been put into evaluating the effects of visualisations (e.g., visual clutter) that have major effects (e.g., cognitive load) in user perception and problem solving ability.

---

[1] http://www.geolabel.info/

Data Provenance: Data provenance can be described as a way to record all data derivations from its origin until the final data product. Consistent representations of data provenance information derived from workflows and databases can be leveraged within the analysis process. Simmhan et al. [70] survey data provenance techniques and present a taxonomy that covers *Usage*, *Subject*, *Representation*, *Storage* and *Dissemination* of provenance.

2.2 Human Trust Building under Uncertainty

We distinguish the relevant human focused theories in *Knowledge Generation*, *Trust Building* and *Analytic Provenance*.

Knowledge Generation: Tory and Möller [76] give an introduction to human factors and highlight that visualisations serve as cognitive support and address human computer cooperation. They suggest that analysts perceive visualisations and match them to their mental model of the problem. Other human factors include a users' knowledge, expertise and tasks but also factors on perception and cognition. Zuk and Carpendale [83] extend the typology on uncertainties by Thomson et al. [75] for reasoning. Both of the typologies include a category about *Subjectivity* that represents the "amount of interpretation or judgment that is included" [75] or the "amount of private knowledge or heuristics utilised" [83]. Green et al. [31] propose the *Human Cognition Model* that covers various human aspects of knowledge creation with visual analytics. They point out that hypothesis generation is very much influenced by the human tendency to accept confirmatory evidence more than disconfirmatory and that the computer can help to mitigate this cognitive bias. Winters et al. [80] show in their study that humans can have different roles, expertise and knowledge that can be applied during the analysis process. This influences how individuals approach visualisations and also how they reason about their problems. Gahegan [28] summarises different kinds of reasoning and relates them to human activities, visualisation tools or computational tools. MacEachren et al. [54] mention that analysts and decision makers behave differently with and without the usage of uncertainty visualisations, whether they are aware of the uncertainties or not. They differentiate between information uncertainty and an "analysts' or decision makers' uncertainty" and also suggest that to capture, represent and understand these uncertainties are future research challenges.

Trust Building: Muir [58] discusses trust relations between humans and machines and builds on Barber's trust dimensions, which are *Persistence*, *Technical Competence*, and *Fiduciary Responsibility* [3]. Furthermore, Muir gives the following recommendations for improving trust calibration: "(1) improving the user's ability to perceive a decision aid's trustworthiness, (2) modifying the user's criterion for trustworthiness, (3) enhancing the user's ability to allocate functions in the system, (4) identifying and selectively recalibrating the user on the dimension(s) of trust which is (are) poorly calibrated" [58]. Dzindolet et al. [21] investigate how trust develops during the usage of a system. Initially, all participants considered the decision aid as trustworthy and reliable. Observing errors caused the participants to distrust the systems unless an explanation was provided. Understanding the errors helped the users to increase their trust in the decision aid, even under uncertainty. Castelfranchi [11] relates trust to the process of knowledge management and sharing and provides a theory that considers the process to be a decisional act of *passing* and *accepting* knowledge. Trust is related to these activities as a mental attitude, but also a decision (e.g., intention to delegate trust) and a behaviour (e.g., relation between trustor and trustee). Uggirala et al. [77] studied humans using systems that include uncertainties by having the users rate their trust at each level through questionnaires. Their study showed that trust relates to competence and an inverse relation to uncertainty, meaning that an increase in uncertainty decreases trust in the systems. Visser et al. [19] provide a design taxonomy for trust cue calibration that includes all kinds of information that may influence human judgement. Figure 2 (left) illustrates trust calibration and the included problems. A miscalibration between the humans' trust and the systems' trustworthiness leads to over- or distrust that are directly connected to disuse and misuse of automation. Skeels et al. [71] deliver a comprehensive perspective on uncertainties for information visualisation and also briefly discuss the role of awareness

Fig. 2: *Left:* Trust calibration adapted from [19], *Right:* Awareness classification adapted from [71].

(see Figure 2 right). They identify *Unidentified Unknowns* as the worst kind of missing information because in that case humans are building more trust than they should do.

Analytic Provenance: Recent research focuses on tracking interaction in order to investigate human analytic processes. Dou et al. [20] present an approach to capture human reasoning processes that distinguishes between internal (interactions within the system) and external capturing (outside the system). Ragan and Goodall [63] go on to mention that provenance tools support the human in memorising and communicating analytic processes. Nguyen et al. [59] survey analytic provenance and show that typical provenance consists of three stages: *Capturing*, *Visualising* and *Utilising* with capturing on different levels (events, low-level actions, high-level sub-tasks, top-level tasks). They also describe the benefits of analytic provenance that go beyond recalling the analysis process to support "evidence in constructing the reasoning process, and facilitating collaboration [...]"[59]. Examples for leveraging analytic provenance will be given in the guidelines for handling uncertainties in Section 4.

In summary, our review discovers two distinctive groups of literature. One group deals with uncertainty propagation and visualisation on machine aspects; the other investigates human-machine trust relations. We observe important gaps in research within visual analytics frameworks. These are uncertainties in the visualisation itself, uncertainties in the coupling between model and visualisation, and uncertainties in the model building. Only few studies relate uncertainties to these human trust building processes. Furthermore, there is no clear terminology that differentiates between uncertainties from machine and human. In the following, we address these issues and provide a model that integrates uncertainty propagation and human trust building.

3 UNCERTAINTY PROPAGATION AND TRUST BUILDING WITHIN THE KNOWLEDGE GENERATION PROCESS

Within the knowledge generation model for visual analytics, uncertainties are propagated in the different components, causing various issues of trust on the human side. In the following, we give an exposé on how uncertainty is propagated on the machine side and trust is callibrated on the human side. Finally, we discuss the concept of awareness that glues uncertainty and trust together.

3.1 Uncertainty Propagation on the Machine Side

We extend the framework of Correa et al. [17], to include uncertainties in the model building, visualisation, and uncertainties in the coupling between the model and visualisation. A brief description of the role of uncertainty in each of these components in a system, follows.

Uncertainty, generally known as the state of not knowing, is attributed to the discrepancy between a data measurement and its representation. According to Griethe and Schumann [32], concepts regarded as uncertainty are *errors*, *imprecision*, *accuracy*, *lineage*, *subjectivity*, *non-specificity*, *noise* (the authors provide a full description of these terms in [32]). Uncertainties also vary depending on the application domain, for example, *topological consistency* in geospatial data. Throughout this paper we will refer to any one of these concepts as uncertainty. These uncertainties in data can be classified mainly into two categories: 1) source uncertainty, and 2) propagated uncertainty. Accounting for such uncertainties in data is important for thorough data analysis, information derivation and informed decision making.

3.1.1 Source Uncertainty

The source uncertainty (Fig. 3-s2) is inherent in data, and largely depends on the way in which this data is collected (authoritative vs. non-authoritative data). In most cases, non-authoritative data such as social media data contains high uncertainties due to the lack of professional gatekeepers and quality control standards among other reasons [27]. Uncertainties in authoritative data are mainly due to reasons such as erroneous measurements, data entry errors, low resolution, and incomplete data. One way of representing source uncertainties is in terms of qualitative measures such as *purpose*, *usage*, and *lineage* (also known as data provenance in most disciplines) [5].

3.1.2 Propagated Uncertainty

Uncertainty in data is propagated during the data modeling (where data undergo transformations such as interpolation, sampling, quantisation etc., [32, 60]) and visualisation stages, where these propagated uncertainties keep aggregating as data travels through these stages in the system side of the knowledge generation model (Figures 1 and 3).

Data Processing (Fig. 3-s3): Processing techniques transform data for preparation purposes (e.g., data cleansing). They can be broadly grouped into normalisation, sub-sampling (reducing the amount of data) or interpolation (increasing the amount of data). Uncertainty measures related to these processing types can be calculated with statistics [36].

Model Building (Fig. 3-s5): During the model building phase, if users have previous knowledge of the model, they achieve a best approximation by typically fitting a parameterised form of the model to the data. Issues of uncertainty arise due to the complexity of the parameterisation (e.g., how many parameters are suitable?) or the appropriateness of the parameters (are the parameters perfect/ good/ bad?), or even the random variation of the model variables. At this stage model calibration introduces a lot of uncertainties by the process of estimating values of unknown parameters in the model. Other uncertainties arise if the distance functions (e.g., euclidean distance or weightings within the similarity function) do not fit data and tasks. Chatfield [14] classifies these types of uncertainties as arising from i) model misspecification.

Model Usage (Fig. 3-s8): Chatfield [14] states that a lack of previous knowledge of the underlying phenomenon causes inadequacies of the model, which gives rise to structural uncertainties. He introduces ii) specifying a general class of models, where the true model is a special, unknown case, and iii) choosing between several models of different structures, as reasons that give rise to uncertainties in model usage. Additionally, the model carries uncertainties in terms of its suitability to the task at hand. Numerical errors and approximations that occur during the implementation of a model gives rise to algorithmic uncertainty [42]. As stated by Brodlie et al. [7], these uncertainties have not been the focus of uncertainty visualisation research thus far.

Visual Mapping (Fig. 3-s4): During the mapping process, the computation of the geometric model (typically done in the mapping process) may be prone to errors due to approximations. Furthermore, the mapping itself causes errors, if the mapping does not fit the underlying data, e.g., when the chosen visual variables do not correspond to the underlying data types. These issues cause uncertainties in this process, which may hinder the comprehensibility of the underlying data. In general, data should be mapped to proper visualisation techniques using the right visual variables (e.g., glyph vs. colour).

Visualisation (Fig. 3-s7): In addition to the visualisation of the uncertainties of the data and above processes, the visualisation itself may contain uncertainties. This is mainly due to the resolution, clutter, and contrast effects of the output visualisation which may hinder the user in gaining insights of the underlying data. Such effects in visualisations that cause uncertainty in the reasoning process are discussed by Zuk and Carpendale [83] and MacEachren and Ganter [53].

Model-Vis Coupling (Fig. 3-s6): One other aspect that we identified for uncertainty propagation in the system, is the uncertainties caused while coupling the model and the visualisation. These uncertainties mainly impact the users' interaction with the system and the model steering that is coupled to the visualisation interactions. Endert et
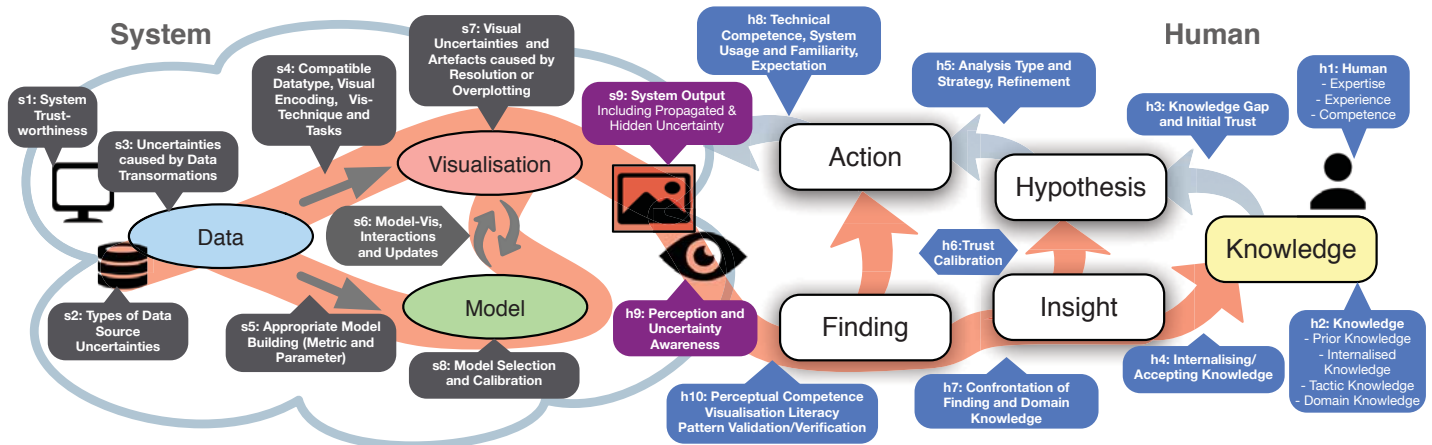
Fig. 3: Each system component may change the data and consequently introduce additional uncertainty. Human trust building within knowledge generation processes is affected by many human factors. The relation between uncertainty and trust is included as the awareness of uncertainties.

al. [23] propose an approach where direct interactions on visualisations are directly translated to model steering interactions (e.g., highlighting an item will increase weighting of the models distance function). If these mappings are not well-designed, these model interactions are translated to model steering interactions that do not fit to the users' intent. Furthermore, the visualisation of the model can be realised in different ways. For example, it is possible to visualise incremental model changes during the training phase (e.g., [26]). However, many visual analytics application just visualise the model result.

All the uncertainties are propagated to the final system output (Fig. 3-s9) which will be observed (Fig. 3-h9) and used by the human for knowledge generation. As important as it is to account for the uncertainties in a system mentioned above, *which* uncertainties to account for is highly dependent on the application scenario of the data.

3.2 Trust Building within Human Analytic Processes

We define *trust* on the human side as a counter part to the machine's *uncertainties* (similar to MacEachran et al.'s [54] distinction between human and machine uncertainties). In the following, we walk through the human concepts of the knowledge generation model [64] and describe them with respect to trust building and highlight influences that are caused by uncertainties. Human trust building can be described as a process of passing and accepting knowledge [11], in our case between human and machine. On the other hand, there are many individual factors that indirectly influence trust building, such as the technical competence [58] and visualisation literacy that is dependent on the users expertise and previous experience with a system.

Trust Calibration (Fig. 3-h6): In each knowledge generation step, users have to calibrate their trust towards their counterpart, the system (or automation as in [19], Fig. 3-s1), and they also need to calibrate their trust between their own previous knowledge (Fig. 3-h1), hypotheses (Fig. 3-h2) and the information that is presented by the system (Fig. 3-s9, h10, h7). Trust calibration is influenced by all the dimensions mentioned by Muir [58]: The "expectation of the persistence of natural physical laws" allows for the creation and usage of mental models (or rule bases). Further, Muir distinguishes three types of technical competence (expert knowledge, technical facility and everyday routine performance) that are essential for trust building. Finally, fiduciary responsibility "is invoked as a basis for trust when the trustor's own technical competence is exceeded by the referent's [(in our case the VA-system)], or is not known to him" [58]. In visual analytics this is often the case when complex processing or data mining algorithms are applied but hidden behind the final visual output.

3.2.1 Knowledge Generation Loop

The knowledge generation loop steers the whole analysis process and consists of knowledge externalisation (there is a knowledge gap) and internalisation (when sufficient trust has been developed). We start our description with the knowledge generation loop because the analysts'

initial trust and hypotheses are based on the prior knowledge and are the foundation for all trust building activities [64].

Knowledge: In general, knowledge can be split into many subparts such as domain-, tactic-, data-, system- or experience-based knowledge (Fig. 3-h2) and has consequently a very individual nature [76, 80]. However, we can distinguish prior knowledge from the knowledge that is gained during analysis and has to be internalised, synthesised and related to the prior knowledge. Within this process, trust develops and pieces of evidence that match or contradict the mental model of the problem are collected and increase or decrease human trust levels. Therefore, trust building depends heavily on the trustworthiness of the machine counterparts (system or data). At the beginning the prior knowledge is assumed to be valid or verified until the analysis reveals evidence that strengthen or weaken it. Through evidence collection supplemental trust emerges and finally transfers gained information to internalised knowledge (Fig. 3-h4). Within this process humans utilize their "private knowledge" in order to judge or interpret pieces of evidences [75, 83]. It is also possible to gain knowledge with analytics, even though the knowledge is based on high uncertainty (if the uncertainty is known and understood as described in [21]). At this stage, we also consider the type of user as an important factor (Fig. 3-h1). A domain expert will behave differently from a machine learning-expert or a novice user. The relation and former experiences with data analysis systems also play a crucial role in trust building. The claim by Muir that "the trust in a machine, once betrayed, may be hard to recover" [58] is backed up by Manzey's study that revealed a similar relationship between error and subjective trust [56]. Furthermore, knowledge includes many sub-components that influence trust building (e.g., the technical competence or subjective attitudes [58]).

3.2.2 Verification Loop

This loop describes higher-level trust building and covers confrontation (of information) and human reasoning. If the trust in the combination of all insights related to the hypothesis exceeds a certain amount and integrates with prior knowledge, we leave the verification loop and arrive at new accepted knowledge (by induction).

Hypothesis: Hypotheses are derived from prior knowledge (because there is a gap) and are the foundation for each verification and exploration cycle (abduction). Initially the trust in a hypothesis is derived from prior knowledge and develops during the analysis process by revealing pieces of evidence (insights) that support or contradict them (Fig. 3-h3). With that respect, humans callibrate their trust and refine their hypothesis in order to come up with an explanation [28]. Also the type and the initial trust in this hypothesis more or less defines the following analysis type (Fig. 3-h5) as the verification loop steers the exploration loop [64]. A very vague and open hypothesis that is weakly trusted will originate analysis with an exploratory fashion that solidifies the analysis step by step, whereas a very defined and a highly trusted hypothesis generates a confirmatory analysis. In reality there are often multiple, conflicting or dependent hypotheses that can be resolved with

the detection of a single expected or unexpected insight.

Insight: Insights are directly related to hypotheses and can be seen as the pieces of evidence for or against them. The trust in insights relates to the amount of similar findings that were produced by the system that support the insight and also on their match to the domain knowledge or the user's mental model of the problem (i.e. what is expected, Fig. 3-h7). If there is a mismatch between the user's mental model and the gained insight one of them has to be adjusted. In other words, the user has to decide whether he trusts himself or the information that was obtained using the machine. As described by Green and Maciejewski [30], human higher level ("System 2") analytical reasoning is able to modify the mental model. Another aspect is that there might be more than one possible interpretation (*insight candidates*) for a finding that can be tested. The analyst develops trust in the alternative interpretations and may (or may not) be able to verify them. However, insights are more likely to occur if they contribute to a plausible narrative (because confirmatory evidence is more likely accepted by humans than disconfirmatory [31]) and therefore calibrate their trust towards evidences that they are comfortable with.

3.2.3 Exploration Loop

The exploration loop covers lower-level trust building and evidence search processes through analysis (deductive). Humans stay in the exploration loop until they develop enough trust in their findings and gain insights by appling their domain knowledge.

Action: Actions reflect many aspects in trust building (Fig. 3-h8) as they are the direct interface between human and machine. In general, actions can help the human to develop more trust in the system itself if the user feels in control of the system. On the other hand, hard operable systems and unexpected behaviour (or errors) may result in a general decrease of trust in the system [21]. This trust depends on the humans technical competence and familiarity with the system and its subcomponents (as described by Muir [58]). Actions further help the human to understand and decrease uncertainties. The ability to explore different kinds of uncertainties enables the user to develop an understanding of these uncertainties, where they arise and how they impact the whole system pipeline (or the sensitivity of data items [17]). Another approach is to actively reduce uncertainties by changing the pipeline or the data (e.g., by choosing more suitable processing methods, mappings or models that introduce less uncertainties). Data changes can be done manually by correction or enrichment. In this case an expert is enabled to bring in his user truth (expert knowledge) in order to change data. However, at this stage users adjusting the data according to their needs are at risk of introducing new human created uncertainties.

Finding: As immediate results of an action, users observe a system reaction (Fig. 3-h9, h10). This reaction (either expected or unexpected) contributes to human trust building [21]. If users develop enough trust in their observations they make them to findings. Or in other words, they stay in the exploration loop until they are able to trust what they see. On a perceptual level, users have to consider being misled by their interpretation of visual elements. For example, a user might spot a visual artefact that is not there in reality (e.g. Muller-Lyon illusion). Furthermore, a finding may include many known and visualised, but also hidden, uncertainties that have been propagated through the system. With this respect, human trust building differentiates when uncertainty information of a finding is communicated and considered [28, 77]. Findings are directly related to insights, which are themselves directly related to gained knowledge. Consequently uncertainty propagates from its root, the data source, through the system and human reasoning until knowledge. This is illustrated by the red flow lines in the knowledge generation model in Figures 1 and 3.

3.3 Awareness

We have considered the relationship between trust and uncertainties within the system. Thus far, we assumed that the user is aware (Fig. 3-h9) of the uncertainties. We now consider the possible effect on trust when the user is unaware of uncertainties and how this might manifest

Table 1: Awareness classification

		system	
		no uncertainties	**uncertainties**
human	**aware**	trust = high chance of human error = none 😄	trust = med-low[1] chance of human error = low 🙂
	mistaken[2]	trust = med-low[1] chance of human error = med-low 😐	(*over*) trust = high chance of human error = high 🙁
	unaware	(*under*) trust = medium chance of human error = none 😄	trust = medium chance of human error = high 😐

[1] depends on degree of understanding

[2] 'mistaken awareness' is when the user wrongly believes the opposite, e.g. no uncertainties when in fact there are uncertainties in the system

itself in subsequent errors (similar to [71]). In addition, we illustrate the situation that the user mistakenly believes there are no uncertainties when in fact there are, and vice versa. A proposed classification of the different states of awareness and uncertainties is shown in Table 1.

We can see that trust it is highest when the user is either aware of no uncertainties or mistakenly believes there are no uncertainties. The latter is a case of over-trust leading to a high chance of errors. The lowest trust is when the user is either aware of uncertainties or mistakenly believes there are uncertainties. These are given a value of medium to low as it depends on the users understanding of the uncertainties, the higher the understanding (or mistaken understanding) the higher the trust. The situation where the user is unaware that there are no uncertainties, is a case of under-trust. Making the user aware would increase their confidence and hence trust. In terms of the chance of errors occurring, this is highest when the user is either unaware of uncertainties or wrongly believes that there are no uncertainties, and lowest when the user is aware or unaware that there are no uncertainties.

Whether or not a user becomes aware of uncertainties and indeed takes note of information presented to them, can be influenced by cognitive biases. These, so called, cognitive biases, first introduced by Kahneman and Tversky in the 1970s [40], are deviations in judgment from what rational decision models would predict that occur in particular situations. Importantly, they are involuntary, affect most people to some degree and generally have a negative impact on decision making. For instance, most people have a poor understanding of statistics and instead apply simplifying heuristics to cope with the uncertainty, which leads to irrational decisions. Arnott [2] lists such statistical biases which highlight the inability to comprehend the practical implications of randomness, base rate, sample size, correlations, regression to the mean and probability in many guises.

Visual analytic systems allow the user to explore datasets but this relies on the user wanting to seek further information. Unfortunately, confirmation bias is the tendency to ignore information that does not agree with the user preconception or hypothesis [41]. In a recent study, Phillips et al. [62] demonstrate that users of information systems, tend to reduce the perceived usefulness of information that does not reinforce their current premise, which in turn reduces their likelihood to explore the data. Over-confidence and perceived expertise has a similar effect. Completeness bias, where the user perceives that the data is logical and correct, without uncertainties, may also reduce information seeking.

We need to be aware of other perceptual and behavioural traits when utilising visualisation and automated systems. For instance, our visual perceptual system is subject to errors due to effects such as contrast, colour, clutter and pre-attentive processing. In addition, automation bias can lead the user to overtrust and relying on wrong information that is produced by an automation, overriding their own ability to judge the situation ("looking-but-not-seeing effect"[56]).

As suggested at the start of this section, awareness of uncertainties can reduce errors and increase the users trust in the data. However, cognitive biases may impede the users awareness and additionally may lead to poor decisions, especially when the user is in a state of uncertainty. Principally due to the involuntary nature of cognitive

biases, reducing their negative effects has proved difficult, even when the user is informed of the possible impact of particular cognitive biases. In the next section, we will enumerate some methods to reduce the impact of cognitive biases, perception effects and the automation bias.

4 GUIDELINES, EXAMPLES & CHALLENGES FOR HANDLING UNCERTAINTIES

In this section we formulate guidelines for handling uncertainties and illustrate them with examples from literature. G1 and G2 are the foundation for uncertainty communication by tracking, quantifying and combining uncertainties. G3, G4 and G6 aim to improve the perception of a systems' trustworthiness through the communication of uncertainty information. G5, G7 and G8 take human issues into account in order to enhance, identify and recalibrate poorly calibrated trust dimensions. Our guidelines have been influenced by Muirs recommendations for improving trust calibration [58] (see Section 2.2). Additionally, we put forward some extensions and challenges that suggest future research directions.

4.1 Uncertainties in a System

G1: Quantify Uncertainties in Each Component: We recommend to quantify uncertainties at every stage of the visual analytics pipeline. In the following we give examples for each component starting from left to right side of the model in Figure 3.

Data Source: Data source uncertainty can be quantitatively measured by, for example, standard deviation to measure the precision of the instrument used to collect data, or counting the number of omissions or commissions in a database to measure the completeness of the data. Furthermore, several qualitative measures can be used to get an overview of the source uncertainty of a dataset. These are the lineage, purpose, and usage (described in Section 3.1). These measures are typically documented in the metadata of a dataset by the data producer. These qualitative measures are subjective measures as they pertain to the views of the producer who documents this metadata according to their use cases. In authoritative data, this metadata will be adequately documented due to the professional gate-keeping of this data. In non-authoritative data (such as social media data) we will see little or none such documentation of measures, due to the lack of gate-keepers. Measure such as credibility [12], reputation [57], or trustworthiness [6] are used to measure the uncertainty of such non-authoritative data.

Data Processing: System inputs that go through transformations such as interpolation, extrapolation, normalisation etc., propagate uncertainty at the system output. Choosing a suitable uncertainty propagation method depends on the confidence level, the extent to which uncertainty quantification is needed, and the computational cost that one can endure [49]. Probabilistic approaches (e.g., Monte Carlo Simulation methods) are known to be most robust in quantifying such uncertainties. Lee and Chen [49] describe in detail five types of probabilistic approaches in their comparative study of uncertainty propagation methods. Statistics such as standard deviation, variance and range are further used to propagate data processing uncertainties. Additionally, Cedilnk and Mendoza [13] use distance based functions to measure the similarity of values and point out that interpolated values can also be used.

Model Building: According to Chatfield [14] uncertainties that arise at the model building stage can be lessened by expert background knowledge (e.g., to know which variables to include), and previous experience/information from similar datasets. However, such expert knowledge may not prevent the user in mistakenly excluding an important variable or adding excess variables. The author points out that one way of avoiding model building uncertainty is to use nonparametric procedures that are based on fewer assumptions. One approach to quantify the uncertainties is to use distance functions to measure the distance of parameterisation from the true value.

Model Usage: Chatfield [14] gives a broad discussion on how uncertainties arise in many aspects of a model (as briefly discussed in Section 3.1). To propagate the uncertainties in the model selection bias, he suggests using the Bayesian averaging approach, and points out the non-triviality of biases. He recommends replicating the study to check if the new data fits the model, although makes the point that

replicating studies are not all that simple to conduct. Works of Fernandez et al. [25], and Kennedy and OHagan [42] demonstrate the use of Bayesian approaches to dealing with model uncertainty.

Visual Mapping: Uncertainties that occur at the visual mapping stage are mainly due to the use of inappropriate visual variables that do not adhere to the data and task at hand. The most sensible approach to assess these uncertainties is through analysing the chosen visual variables and metaphors against existing systematic taxonomies. In his task by data type taxonomy, Schneiderman [69] categorises existing information visualisation techniques according to the type of data (e.g., temporal data) and the task (e.g., zoom or filter). In the case of uncertainty visualisation, we need to consider the added uncertainty dimension to the underlying data. In addition to MacEachrens [51] work on manipulating several visual metaphors to represent uncertainty, Buttenfield and Weibel [9] present a framework for categorising different cartographic visualisation methods according to the uncertainty elements (e.g., positional accuracy or the lineage of the data) and the measurement scale of the data (e.g. discrete or categorical data). Furthermore, in a recent classification, Senaratne and Gerharz [67] categorised popular uncertainty visualisation methods according to the measurement scale of the data (e.g., continuous or categorical), supported data format (e.g., raster or vector), and the type of uncertainty element in the data (e.g., positional or thematic uncertainty).

Visualisation: The works of MacEachren and Howard [37, 51] have developed visual metaphors for representing uncertainty, that fits well with the human cognition model. Examples are the use of blurring effects, transparency, or coarsely structured surfaces to represent uncertainty. Their impact on decision making under uncertainty has been explored in several studies (e.g., [68]). MacEachren and Ganter [53] classify visualisation uncertainties as being developed through two types of errors. *Type 1: seeing what is not really there and Type 2: over-seeing what is really there*. The authors emphasise the need for tools to aid the users in seeing through these Type 1 and Type 2 errors in visualisations. Relating to the Type 2 errors in particular, Brodlie et al. [7] point out uncertainties caused by the lower resolution of the visualisation in contrast to the resolution of the data.

Model-Vis Coupling: We are not aware of existing methods that quantify the uncertainties that arise due to the coupling between visualisation and models. One possibility to quantify differences between model and visualisations is to compare measures of the different spaces (e.g., visual 2D compared to HD as described in [73]) in order to compare model and visualisation characteristics. For example, groups and distances between data items in model space (e.g., between cluster centroids) can be compared to their distances in visual space (e.g., projected distances between cluster centroids). Another approach is to measure how model changes (e.g., via human interaction or data streaming) are propagated to the visualisation. Most of the visualisations take the final model result but there are several cases, and models that deliver incremental results that can be visualised (e.g., [26]).

G2: Propagate and Aggregate Uncertainties: Systems require powerful and sophisticated techniques to support exploration of large datasets. Adding different kinds of uncertainty to this data requires an increase in the level of sophistication of the system. Works of [17] estimate the data source uncertainty and the propagated uncertainty through transformations, via sensitivity analysis and error modeling. To simplify the computations, we require intelligent methods to *aggregate* these propagated uncertainties. Klir and Wierman [44] describe methods to aggregate source uncertainties and propagated uncertainties in the visualisation pipeline. Also, through a remote sensing classification application, Van de Wel et al. [78] describe the use of an entropy measure to build a weighted uncertainty aggregation measure. They map the different kinds of uncertainty to one measure based on a weighted criteria. Learning from this, an alternative would be for the user of the system to weigh each kind of uncertainty stemming from the system, based on its importance to the use case at hand.

G3: Visualise Uncertainty Information: Uncertainty visualisation is known to be a most effective medium to communicate such source and propagated uncertainties. Griethe and Schumann [32] present

a pipeline to show the process of uncertainty visualisation. In their pipeline, they differentiate between four different kinds of data flows. 1) *the basic data transformation process through the visualisation pipeline*: is separated in to data components and their corresponding uncertainties, such that the user sees the underlying uncertainty; 2) *in/output of the acquisition of uncertainty data*: data at every stage of the visualisation will carry uncertainty, and needs to be considered; 3) *dependencies between the visualisation of the raw data and its uncertainty*: while the data is explored, its uncertainty is considered as an integral part of the data. However, decisions in processing the uncertainty is dependent on what raw data is focused on, which rendering techniques and geometric forms for models are chosen. 4) *parameterisation of the pipeline*: uncertainty is not visible by itself at every data component as in 1). Instead, uncertainty will be used to parameterise visualisation of the other data (as done by Schmidt et al. [65]). Furthermore, in visualising the uncertainties in the different stages of the data, one needs to carefully consider the different design principles. Works such as of Pang [60] focus on visualising multi-dimensional uncertainties in data. These can be used as guidelines on how to design visualisations to incorporate different uncertainties propagated through analysis system. Griethe and Schumann [32] further emphasise that the decisions on the amount of user interaction on such an uncertainty visualisation process depends on the users experience and the principles of the visualisation system. Finally, a system should report the uncertainties as cognitive cues about its self-confidence as suggested by Cai and Lin [10]. As a result, users are more comfortable in adjusting their trust appropriately.

G4: Enable Interactive Uncertainty Exploration: Within a visual analytics environment, enabling the user to interact and explore different visualisations for different uncertainties stemming from the different components of the system will enrich the users understanding of the true nature of the data, and additionally, how different propagated uncertainties influence the final output. Also, the ability to use a variety of visualisations may also help with illusion type cognitive biases such as clustering and correlation. It is also important to give the user control to decide which of these uncertainties should influence the final output, or with how much importance it should influence the output. Providing the user with the possibility of giving weighted measures for each uncertainty component would be a realistic approach. Furthermore, Correa at al. [17] present several approaches including uncertainty projections and visualisations that enable the user to explore the uncertainties of individual data items and the impacts of different uncertainties (Figure 4-a).

4.2 Supporting Human Factors

G5: Make the Systems Functions Accessible: Accessible, intuitive and easy to use interaction techniques will increase the technical competence of the analyst and consequently enhance human trust building [58]. In this respect, different user groups have to be considered. Expert and power-users of an analysis tool will need different interaction possibilities and guidance than novice users. For example, with visual analytics tools we often observe users having problems with model steering interactions such as parameter setting [64]. In this case, switching between expert or learning mode might be a first step in that direction. Endert et al. give a nice example of "semantic interaction" [22] where direct manipulation interactions are directly translated to model steering interactions. Furthermore, Chuang et al. [16] provide guidelines for designing trustworthy and understandable visual analytics systems. Their recommendations can verify modeling decisions and provide model interactions during analysis.

G6: Support the Analyst in Uncertainty Aware Sensemaking: Human sensemaking can be supported by offering note-taking or knowledge management components connected to the systems, where humans can externalise and organise their thoughts in order to bridge data and knowledge management [64] (see Figure 4-d, e.g., the Sandbox for Analysis [15, 81]). Our recommendation is to transfer and visualise uncertainty information to the findings that are imported from the analysis part in order to support humans in calibrating their trust in the findings' trustworthiness [19]. We can imagine that a system will automatically

take care of relations between (conflicting) hypotheses, findings, insights and take the role of an unbiased counterpart to the human by including uncertainty information at any stage [31]. A system could, for example, calculate aggregated uncertainty scores for pieces of evidence that have been grouped by the user. In addition, evidence connected to hypotheses and insights may be explicitly marked by the user as *trusted* or *unknown* (or intermediate value). This would be a form of trust annotation that can be matched to uncertainty measures. Furthermore, humans can integrate external evidences from other systems or their own knowledge that might complete the big picture of the analysis. Utilising all the connected information enables a system to offer uncertainty and trust cues e.g., as glyphs connected to the items (such as traffic lights, radar charts or avatars as described in [19]). Figure 4-b illustrates an example view of automation blocks that are enriched with specially designed glyphs that serve as trust cues.

G7: Analyse Human Behaviour in order to Derive Hints on Problems and Biases: Tracking human behaviour can be beneficial in deriving hints on the users of a system. We therefore recommend to leverage analytic provenance approaches suggested by Nguyen et al. [59]. Low level interaction tracking can be used to predict users performance [8] or infer the user frustration [34]. These methods could be enhanced for predicting a users trust level. Closer measures related to uncertainties and trust building can be captured by the rate of overall decision switching. Goddard et al. [29] measured automation bias by noting correct to incorrect prescription switching. Furthermore, Klemmer et al. were able to detect important notes or items based on user tracking [43] (see Figure 4-c). These methods could be leveraged by a system to automatically suggest alternative visualisations or items that have not been utilised. The latter may be useful in mitigating some selection based cognitive biases such as confirmation bias [46]. Another approach to derive human trust measures is to analyse user generated contents. A system could automatically seek for signal words such as "unsure, uncertain, maybe, perhaps ...". Zhou et al. describe 27 language features grouped as: quantity, informality, expressivity, affect, uncertainty, nonimmediacy, diversity, specificity, and complexity [82]. Also, Tenbrink [74] investigated how to derive cognitive analytical processes based on language data. Physical or other human sensors such as eye-tracking can also be used. Kurzhals et al. give an overview on the potential for eye tracking in visual analytics [47]. Furthermore, user analysis may be used during system development and evaluation. Scholtz describes five evaluation areas: Situation awareness, collaboration, interaction, creativity, and utility [66]. We imagine protocol analysis [24] as a useful method to interpret "think aloud" evaluations. User interviews using trust questionnaires could also be conducted [77] in order to investigate the relation between uncertainty and trust for system evaluations. In addition, Bass et al. propose a method to analyse and predict a humans understanding of automation [4].

G8: Enable Analysts to Track and Review their Analysis: This guideline points to post-analysis activities as a method to detect and mitigate biases. During analysis, users often focus on searching potential evidences without considering alternatives, errors or uncertainties. In addition, users in their "work flow" should not be interrupted [38]. Therefore, we recommend that the analyst is able (or even encouraged) to look and think about his analysis afterwards, without interruption during the analysis. In our opinion this is a better way than warning users during their analysis (e.g., by popup dialogs) as recent studies show that too often warnings may lead to the opposite [1]. Support to mitigate statistical biases (see Section 3.3) should be provided, such as presenting the user with base rate information (e.g., typical distribution), estimating realistic probabilities or indicating that a particular 'behaviour' is expected rather than a special case, as with regression to the mean. Structured analytic techniques such as a devils advocate may also be ways that help the user to detect problems and lessen the impact of confirmation bias in particular. Furthermore, analysis process visualisation enables involving other users and story telling. Provenance systems such as CzSaw [39], but also systems including story telling [45], are a good starting point in that direction.
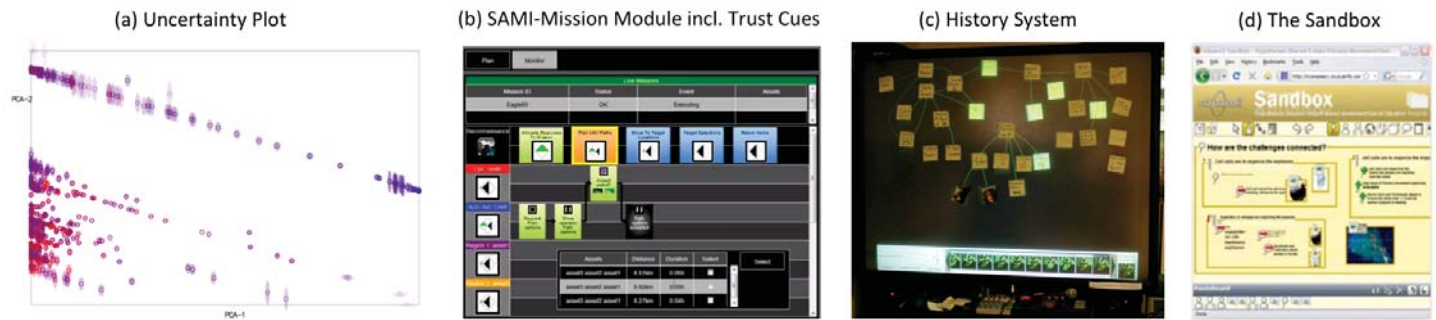
Fig. 4: Examples for uncertainty aware trust building from different domains: (a) An uncertainty projection to explore how data items are affected by uncertainties [17], (b) decision support including trust cue design from [19], (c) deriving important user notes based on user tracking [43], (d) integrating evidences for computer assisted knowledge management [15].

5 DISCUSSION, LIMITATION AND CONCLUSION

In this section, we discuss our findings, provide limitations and open questions of our study, and conclude with some takeaway messages.

5.1 Discussion

Our framework shows that human trust building under uncertainty is an extremely complicated, individual process and is influenced by many factors directly as well as indirectly. Furthermore, users informed with uncertainty information can avoid falling into traps concerning mistaken uncertainties and unaware uncertainties. Readers also need to note that the framework has to be tailored to concrete, individual cases where the scope of uncertainties, users and their tasks are known. The core value of our framework is to provide a balanced view on the role of uncertainty propagation and trust building in visual analytics by considering human and machine aspects together.

The guidelines in this paper will be useful to estimate the dynamics of uncertainties in developing visual analytics applications. The terms and structure we outline in Section 3 provide an overview of uncertainty propagation, both from the source data and from algorithmic manipulation. With this structure, practitioners and researchers can attempt to quantify uncertainties though the process of data transformations. Depending on its use, this quantification will help users determine effective visualisation techniques by thinking of the trade-offs between gaining insights and showing uncertainties.

The framework we provide in Section 3 can be used to educate users of visual analytics applications about uncertainties and their impact, so that they might reduce errors (e.g., cognitive biases) and build trust whilst analysing data. We recommend system developers to provide a simple tutorial of their visual analytics applications using some usage scenarios. In addition to assisting users, the material itself provides a groundwork for education of uncertainties and trust building for designers, practitioners, and researchers. We believe that the impact of uncertainties will decrease as users gain awareness.

There are many implications of this model. As explained, it is necessary for us to capture humans perceived uncertainties and trust levels at a given moment of analysis. One way is to intervene in the analysis process by asking the users to input this directly, which will ensure accurate estimates of their current status. However, to avoid interruptions it would be useful to compute this automatically. This may be possible through tracing and interpreting usage logs to estimate the level of trust. Data provenance may be an effective method to track uncertainty propagation that enables us to increase uncertainty awareness. On the other hand, if analytic provenance methods are used to infer human measures this may give hints on trust building processes. Combining measures/methods from both sides has the potential to identify relations between uncertainty propagation and human trust building.

Our contribution is to categorise types of uncertainties, awareness of them, and human trust building process. However, there are many external factors that can influence individual processes. For example, our model assumes a single analyst perspective, which simplifies the knowledge generation process. In the real world, many interdependent knowledge generation loops run in parallel and often conflict each other, which can result in uncertain outcomes. Furthermore, taking into account collaboration between human analysts would extend the model to explain the dynamics of real world scenarios with a team of analysts.

5.2 Limitations and Open Questions

The scope of this study provides a framework of unpacking uncertainty propagation within visual analytics processes as well as discovering the human trust building process. Here we provide limitations of our approach as well as open questions that future researchers can investigate.

Firstly, uncertainties are difficult to be quantified and categorised into a single process. In visual analytics systems, uncertainties can be propagated and implied through the pipelines, as we discussed. Thus, combination of uncertainties from multiple sources could be larger than the sum. Our model does not provide a quantified model of such intertwined process of uncertainty propagation just yet. As outlined in G1 and G2, some efforts have been made to quantify and aggregate different subsets of uncertainty propagation within visual analytics process. Researchers may need to integrate such efforts using our overarching model and predict such uncertainty propagation in a specific context.

Secondly, another open question is whether the transparency of uncertainty propagation is always good and how much of it is beneficial to users. Our model builds upon an assumption that making the uncertainty propagation transparent will let users be aware of variation in their outcomes. However, providing too much information could confuse, overwhelm, and mislead users, thereby making unwanted human errors. Furthermore, it is also a tradeoff between efficiency and accuracy. For instance, applications for human safety, where uncertainty can result in catastrophic results, may need to consider as much transparency as possible. On the other hand, some business analytics may require fast and reasonable analysis results. Thus, it will be interesting to investigate what are proper amounts and methods to communicate uncertainty information to different groups of visual analytics users.

Thirdly, in line with previous points, it is also an open question whether the awareness of uncertainties leads to increasing or decreasing trust. This question may be from the human's trust building process. To build trust in visual analytics outcomes, users may need to build trust in the visual analytics system first. In this process, the awareness of uncertainties may lead to increasing the awareness of visual analytics process but not to increasing trust in the outcomes. Thus, future research may investigate the sophisticated process of human's trust building steps under uncertainty.

Fourthly, in this regard, we may think of the awareness provenance to verify human's understanding. We introduced the concept of awareness to bridge between machine's uncertainties and human's trust. The awareness again is highly subjective to individuals like the trust level, so it will be difficult to quantify the information. Nonetheless, the awareness indeed affects the entire process, so we call for research into capturing it.

These points above do not capture all limitations and open question from our study, but will be an interesting start for future work.

5.3 Conclusion

In conclusion, we have illustrated how uncertainties arise, propagate and impact human knowledge generation processes by relating the concepts of trust, calibration and awareness. Further, we have given hints on misconfigurations of uncertainty awareness that may cause human errors in data analysis. We provide guidelines that describe various ways to handle uncertainties and to include human factors in order to enhance human trust calibration. Finally, we put forward open research areas that will contribute to more reliable knowledge generation in visual analytics in the future.

REFERENCES

[1] B. B. Anderson, C. B. Kirwan, J. L. Jenkins, D. Eargle, S. Howard, and A. Vance. How polymorphic warnings reduce habituation in the brain—insights from an fmri study. *Proc. of CHI'15*, 2013.

[2] D. Arnott. Cognitive biases and decision support systems development: a design science approach. *Inf. Syst. J.*, 16(1):55–78, 2006.

[3] B. Barber. *The logic and limits of trust*, volume 96. Rutgers University Press New Brunswick, NJ, 1983.

[4] E. J. Bass, L. A. Baumgart, and K. K. Shepley. The effect of information analysis automation display content on human judgment performance in noisy environments. *Journal of cognitive engineering and decision making*, 7(1):49–65, 2013.

[5] M. K. Beard, B. P. Buttenfield, and S. B. Clapham. *NCGIA Research Initiative 7: Visualization of Spatial Data Quality: Scientific Report for the Specialist Meeting 8-12 June 1991, Castine, Maine*. National Center for Geographic Information and Analysis, 1991.

[6] M. Bishr and K. Janowicz. Can we trust information?-the case of Volunteered Geographic Information. In *Towards Digital Earth Search Discover and Share Geospatial Data Workshop at Future Internet Symposium, volume*, volume 640, 2010.

[7] K. Brodlie, R. A. Osorio, and A. Lopes. A review of uncertainty in data visualization. In *Expanding the Frontiers of Visual Analytics and Visualization*, pages 81–109. Springer, 2012.

[8] E. T. Brown, A. Ottley, H. Zhao, Q. Lin, R. Souvenir, A. Endert, and R. Chang. Finding waldo: Learning about users from their interactions. *IEEE Trans. Vis. Comput. Graph.*, 20(12):1663–1672, 2014.

[9] B. Buttenfield and R. Weibel. Visualizing the quality of cartographic data. In *Third International Geographic Information Systems Symposium (GIS/LIS 88), San Antonio, Texas*, 1988.

[10] H. Cai and Y. Lin. Tuning trust using cognitive cues for better human-machine collaboration. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, volume 54, pages 2437–2441. SAGE Publications, 2010.

[11] C. Castelfranchi. Trust mediation in knowledge management and sharing. In *Trust Management, Second International Conference, iTrust 2004, Oxford, UK, March 29 - April 1, 2004, Proceedings*, pages 304–318, 2004.

[12] C. Castillo, M. Mendoza, and B. Poblete. Information credibility on twitter. In *Proceedings of the 20th international conference on World wide web*, pages 675–684. ACM, 2011.

[13] A. Cedilnik and P. Rheingans. Procedural annotation of uncertain information. In *Visualization 2000. Proceedings*, pages 77–84. IEEE, 2000.

[14] C. Chatfield. Model uncertainty. *Encyclopedia of Environmetrics*, 2006.

[15] L. Chien, A. Tat, P. Proulx, A. Khamisa, and W. Wright. Grand challenge award 2008: Support for diverse analytic techniques - nspace2 and geotime visual analytics. In *Proceedings of the IEEE Symposium on Visual Analytics Science and Technology, IEEE VAST 2008, Columbus, Ohio, USA, 19-24 October 2008*, pages 199–200, 2008.

[16] J. Chuang, D. Ramage, C. D. Manning, and J. Heer. Interpretation and trust: designing model-driven visualizations for text analysis. In *CHI Conference on Human Factors in Computing Systems, CHI '12, Austin, TX, USA - May 05 - 10, 2012*, pages 443–452, 2012.

[17] C. Correa, Y.-H. Chan, and K.-L. Ma. A framework for uncertainty-aware visual analytics. In *Visual Analytics Science and Technology, 2009. VAST 2009. IEEE Symposium on*, pages 51–58. IEEE, 2009.

[18] A. C. Cullen and H. C. Frey. *Probabilistic techniques in exposure assessment: a handbook for dealing with variability and uncertainty in models and inputs*. Springer Science & Business Media, 1999.

[19] E. J. de Visser, M. S. Cohen, A. Freedy, and R. Parasuraman. A design methodology for trust cue calibration in cognitive agents. In *Virtual, Augmented and Mixed Reality. Designing and Developing Virtual and Augmented Environments - 6th International Conference, VAMR 2014, Held as Part of HCI International 2014, Heraklion, Crete, Greece, June 22-27, 2014, Proceedings, Part I*, pages 251–262, 2014.

[20] W. Dou, W. Ribarsky, and R. Chang. Capturing reasoning process through user interaction. *Proc. IEEE EuroVAST*, 2, 2010.

[21] M. T. Dzindolet, S. A. Peterson, R. A. Pomranky, L. G. Pierce, and H. P. Beck. The role of trust in automation reliance. *Int. J. Hum.-Comput. Stud.*, 58(6):697–718, 2003.

[22] A. Endert, P. Fiaux, and C. North. Semantic interaction for sensemaking: Inferring analytical reasoning for model steering. *IEEE Trans. Vis. Comput. Graph.*, 18(12):2879–2888, 2012.

[23] A. Endert, M. Hossain, N. Ramakrishnan, C. North, P. Fiaux, and C. Andrews. The human is the loop: new directions for visual analytics. *Journal of Intelligent Information Systems*, pages 1–25, 2014.

[24] K. A. Ericsson and H. A. Simon. *Protocol analysis*. MIT-press, 1984.

[25] C. Fernandez, E. Ley, and M. F. Steel. Model uncertainty in cross-country growth regressions. *Journal of applied Econometrics*, 16(5):563–576, 2001.

[26] D. Fisher, S. M. Drucker, and A. C. König. Exploratory visualization involving incremental, approximate database queries and uncertainty. *IEEE Computer Graphics and Applications*, 32(4):55–62, 2012.

[27] A. Flanagin and M. Metzger. The credibility of volunteered geographic information. *GeoJournal*, 72(3):137–148, 2008.

[28] M. Gahegan. Beyond tools: Visual support for the entire process of giscience. *Exploring geovisualization*, (4):83–99, 2005.

[29] K. Goddard, A. V. Roudsari, and J. C. Wyatt. Automation bias: Empirical results assessing influencing factors. *I. J. Medical Informatics*, 83(5):368–375, 2014.

[30] T. M. Green and R. Maciejewski. A role for reasoning in visual analytics. In *46th Hawaii International Conference on System Sciences, HICSS 2013, Wailea, HI, USA, January 7-10, 2013*, pages 1495–1504, 2013.

[31] T. M. Green, W. Ribarsky, and B. D. Fisher. Building and applying a human cognition model for visual analytics. *Information Visualization*, 8(1):1–13, 2009.

[32] H. Griethe and H. Schumann. The visualization of uncertain data: Methods and problems. In *SimVis*, pages 143–156, 2006.

[33] R. B. Haber and D. A. McNabb. Visualization idioms: A conceptual model for scientific visualization systems. *Visualization in scientific computing*, 74:93, 1990.

[34] L. Harrison, W. Dou, A. Lu, W. Ribarsky, and X. Wang. Analysts aren't machines: Inferring frustration through visualization interaction. In *Visual Analytics Science and Technology (VAST), 2011 IEEE Conference on*, pages 279–280. IEEE, 2011.

[35] T. Hengl. Visualisation of uncertainty using the hsi colour model: computations with colours. In *Proceedings of the 7th International Conference on GeoComputation*, pages 8–17, 2003.

[36] G. B. M. Heuvelink and P. A. Burrough. Developments in statistical approaches to spatial uncertainty and its propagation. *International Journal of Geographical Information Science*, 16(2):111–113, 2002.

[37] D. Howard and A. M. MacEachren. Interface design for geographic visualization: Tools for representing reliability. *Cartography and Geographic Information Systems*, 23(2):59–77, 1996.

[38] S. T. Iqbal and E. Horvitz. Disruption and recovery of computing tasks: field study, analysis, and directions. In *Proceedings of the 2007 Conference on Human Factors in Computing Systems, CHI 2007, San Jose, California, USA, April 28 - May 3, 2007*, pages 677–686, 2007.

[39] N. Kadivar, V. Y. Chen, D. Dunsmuir, E. Lee, C. Z. Qian, J. Dill, C. D. Shaw, and R. F. Woodbury. Capturing and supporting the analysis process. In *Proceedings of the IEEE Symposium on Visual Analytics Science and Technology, IEEE VAST 2009, Atlantic City, New Jersey, USA, 11-16 October 2009, part of VisWeek 2009*, pages 131–138, 2009.

[40] D. Kahneman. *Thinking, fast and slow*. Macmillan, 2011.

[41] D. Kahneman and A. Tversky. Subjective probability: A judgment of representativeness. In *The Concept of Probability in Psychological Exper-*

iments, pages 25–48. Springer, 1974.

[42] M. C. Kennedy and A. O'Hagan. Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(3):425–464, 2001.

[43] S. R. Klemmer, M. Thomsen, E. Phelps-Goodman, R. Lee, and J. A. Landay. Where do web sites come from?: capturing and interacting with design history. In *Proceedings of the CHI 2002 Conference on Human Factors in Computing Systems: Changing our World, Changing ourselves, Minneapolis, Minnesota, USA, April 20-25, 2002.*, pages 1–8, 2002.

[44] G. J. Klir and M. J. Wierman. *Uncertainty-based information: elements of generalized information theory*, volume 15. Springer Science & Business Media, 1999.

[45] R. Kosara and J. Mackinlay. Storytelling: The next step for visualization. *Computer*, 46(5):44–50, 2013.

[46] D. R. Kretz, B. Simpson, and C. Graham. A game-based experimental protocol for identifying and overcoming judgment biases in forensic decision analysis. In *Homeland Security (HST), 2012 IEEE Conference on Technologies for*, pages 439–444. IEEE, 2012.

[47] K. Kurzhals, B. D. Fisher, M. Burch, and D. Weiskopf. Evaluating visual analytics with eye tracking. In *Proceedings of the Fifth Workshop on Beyond Time and Errors: Novel Evaluation Methods for Visualization, BELIV 2014, Paris, France, November 10, 2014*, pages 61–69, 2014.

[48] B. C. Kwon, B. Fisher, and J. S. Yi. Visual analytic roadblocks for novice investigators. In *2011 IEEE Conference on Visual Analytics Science and Technology (VAST)*, pages 3–11, 2011.

[49] S. H. Lee and W. Chen. A comparative study of uncertainty propagation methods for black-box-type problems. *Structural and Multidisciplinary Optimization*, 37(3):239–253, 2009.

[50] V. Lush, L. Bastin, and J. Lumsden. Developing a geo label: providing the gis community with quality metadata visualisation tools. *Proceedings of the 21st GIS Research UK (GISRUK 3013), Liverpool, UK*, pages 3–5, 2013.

[51] A. M. MacEachren. Visualizing uncertain information. *Cartographic Perspectives*, 13(13):10–19, 1992.

[52] A. M. MacEachren. Visual Analytics and Uncertainty: Its Not About the Data. In *EuroVis Workshop on Visual Analytics (EuroVA)*. The Eurographics Association, 2015.

[53] A. M. MacEachren and J. H. Ganter. A pattern identification approach to cartographic visualization. *Cartographica: The International Journal for Geographic Information and Geovisualization*, 27(2):64–81, 1990.

[54] A. M. MacEachren, A. Robinson, S. Hopper, S. Gardner, R. Murray, M. Gahegan, and E. Hetzler. Visualizing geospatial information uncertainty: What we know and what we need to know. *Cartography and Geographic Information Science*, 32(3):139–160, 2005.

[55] A. M. MacEachren, R. E. Roth, J. O'Brien, B. Li, D. Swingley, and M. Gahegan. Visual semiotics & uncertainty visualization: An empirical study. *Visualization and Computer Graphics, IEEE Transactions on*, 18(12):2496–2505, 2012.

[56] D. Manzey, J. Reichenbach, and L. Onnasch. Human performance consequences of automated decision aids: The impact of degree of automation and system experience. *Journal of Cognitive Engineering and Decision Making*, pages 57–87, 2012.

[57] P. Maué. Reputation as tool to ensure validity of vgi. In *Workshop on volunteered geographic information*, 2007.

[58] 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.

[59] P. H. Nguyen, K. Xu, and B. Wong. A survey of analytic provenance. *Middlesex University*, 2014.

[60] A. Pang. Visualizing uncertainty in geo-spatial data. In *Proceedings of the Workshop on the Intersections between Geospatial Information and Information Technology*, pages 1–14, 2001.

[61] A. T. Pang, C. M. Wittenbrink, and S. K. Lodha. Approaches to uncertainty visualization. *The Visual Computer*, 13(8):370–390, 1997.

[62] B. K. Phillips, V. R. Prybutok, and D. A. Peak. Decision confidence, information usefulness, and information seeking intention in the presence of disconfirming information. *InformingSciJ*, 17:1–24, 2014.

[63] E. D. Ragan and J. R. Goodall. Evaluation methodology for comparing memory and communication of analytic processes in visual analytics. In *Proceedings of the Fifth Workshop on Beyond Time and Errors: Novel Evaluation Methods for Visualization, BELIV 2014, Paris, France, November 10, 2014*, pages 27–34, 2014.

[64] D. Sacha, A. Stoffel, F. Stoffel, B. C. Kwon, G. P. Ellis, and D. A. Keim.

[65] G. S. Schmidt, S.-L. Chen, A. N. Bryden, M. A. Livingston, L. J. Rosenblum, and B. R. Osborn. Multidimensional visual representations for underwater environmental uncertainty. *Computer Graphics and Applications, IEEE*, 24(5):56–65, 2004.

[66] J. Scholtz. Beyond usability: Evaluation aspects of visual analytic environments. In *IEEE Symposium On Visual Analytics Science And Technology, IEEE VAST 2006, October 31-November 2, 2006, Baltimore, Maryland, USA*, pages 145–150, 2006.

[67] H. Senaratne and L. Gerharz. An assessment and categorisation of quantitative uncertainty visualisation methods for geospatial data. In *14th AGILE international conference on geographic information science-advancing geoinformation science for a changing world. AGILE*, 2011.

[68] H. Senaratne, L. Gerharz, E. Pebesma, and A. Schwering. Usability of spatio-temporal uncertainty visualisation methods. In *Bridging the Geographic Information Sciences*, pages 3–23. Springer, 2012.

[69] B. Shneiderman. The eyes have it: A task by data type taxonomy for information visualizations. In *Visual Languages, 1996. Proceedings., IEEE Symposium on*, pages 336–343. IEEE, 1996.

[70] Y. L. Simmhan, B. Plale, and D. Gannon. A survey of data provenance techniques. *Computer Science Department, Indiana University, Bloomington IN*, 47405, 2005.

[71] M. M. Skeels, B. Lee, G. Smith, and G. G. Robertson. Revealing uncertainty for information visualization. *Information Visualization*, 9(1):70–81, 2010.

[72] S. Tak and A. Toet. Color and uncertainty: It is not always black and white. 2014.

[73] A. Tatu, G. Albuquerque, M. Eisemann, P. Bak, H. Theisel, M. A. Magnor, and D. A. Keim. Automated analytical methods to support visual exploration of high-dimensional data. *IEEE Trans. Vis. Comput. Graph.*, 17(5):584–597, 2011.

[74] T. Tenbrink. Cognitive discourse analysis: accessing cognitive representations and processes through language data. *Language and Cognition*, 7:98–137, 3 2015.

[75] J. Thomson, E. Hetzler, A. MacEachren, M. Gahegan, and M. Pavel. A typology for visualizing uncertainty. In *Electronic Imaging 2005*, pages 146–157. International Society for Optics and Photonics, 2005.

[76] M. Tory and T. Möller. Human factors in visualization research. *IEEE Trans. Vis. Comput. Graph.*, 10(1):72–84, 2004.

[77] A. Uggirala, A. K. Gramopadhye, B. J. Melloy, and J. E. Toler. Measurement of trust in complex and dynamic systems using a quantitative approach. *International Journal of Industrial Ergonomics*, 34(3):175–186, 2004.

[78] F. J. Van der Wel, L. C. Van der Gaag, and B. G. Gorte. Visual exploration of uncertainty in remote-sensing classification. *Computers & Geosciences*, 24(4):335–343, 1998.

[79] C. Ware. *Information visualization: perception for design*. Elsevier, 2012.

[80] K. M. Winters, D. Lach, and J. B. Cushing. Considerations for characterizing domain problems. In *Proceedings of the Fifth Workshop on Beyond Time and Errors: Novel Evaluation Methods for Visualization, BELIV 2014, Paris, France, November 10, 2014*, pages 16–22, 2014.

[81] W. Wright, D. Schroh, P. Proulx, A. Skaburskis, and B. Cort. The sandbox for analysis: Concepts and methods. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '06, pages 801–810, New York, NY, USA, 2006. ACM.

[82] L. Zhou, J. K. Burgoon, J. F. Nunamaker, and D. Twitchell. Automating linguistics-based cues for detecting deception in text-based asynchronous computer-mediated communications. *Group decision and negotiation*, 13(1):81–106, 2004.

[83] T. Zuk and M. S. T. Carpendale. Visualization of uncertainty and reasoning. In *Smart Graphics, 7th International Symposium, SG 2007, Kyoto, Japan, June 25-27, 2007, Proceedings*, pages 164–177, 2007.

Knowledge generation model for visual analytics. *IEEE Trans. Vis. Comput. Graph.*, 20(12):1604–1613, 2014.