

Towards a visual guide for communicating uncertainty in Visual Analytics[☆]Karsten Seipp^{*,1,a}, Francisco Gutiérrez^a, Xavier Ochoa^b, Katrien Verbert^a^a Department of Computer Science, KU Leuven, Celestijnenlaan 200A, Leuven 3001, Belgium^b Department of Administration, Leadership and Technology, New York University - Steinhardt, 82 Washington Square East, 7th Floor, New York, NY 10003, USA

ARTICLE INFO

Keywords:

Uncertainty representation
 Visual Analytics
 Information Visualisation
 Geospatial visualisation
 Algorithm mechanics
 Visual language
 Non-experts
 Laymen

ABSTRACT

This article presents a first step towards the definition of a visual guide for communicating uncertainty which is to fit into existing visualisation frameworks and toolkits. The first entry in our guide is made by a set of visual variables appropriate for representing areal uncertainty in algorithm mechanics. Such visualisations show users how data points are distributed in the classification space and allow them to understand the “goodness-of-fit” of their data to the algorithm. This is important for Visual Analytics applications, which combine Information Visualisation with information mining techniques in an interactive decision-making process. Model uncertainties stemming from widely spread data points need to be visualised so that the user can make adjustments and improve the analysis.

To capitalise on established knowledge and meaning, we explore whether popular visual variables for representing areal uncertainty in the domain of geospatial visualisation may also be effective for representing uncertainty in the visualisation of the mechanics of K-means clustering and Linear Regression algorithms, as both use a spatial distribution of data points. In a study with 500 participants we find that overall the visual means *opacity* performs best, followed by *texture*, but that *grid* and *blur* may be unsuitable for quantifying uncertainty. The performance of *contour lines* appears to depend on the algorithm visualisation. Using this study, we extend the validity of a set of domain-specific findings from geospatial visualisation to the visualisation of algorithm mechanics and use these to form the first building blocks of a cross-disciplinary visual guide for representing uncertainty, laying promising foundations for future work.

1. Introduction

Visual Analytics is increasingly applied in a variety of application domains, including health and engineering [1], financial analysis [2], and Learning Analytics [3]. By allowing expert users to alter parameters to influence algorithms, the visualisation of data changes and provides insight and understanding [4,5]. In recent years, there has been a growing interest by scholars not proficient in data-processing and visualisation to also employ such tools for their research [6,7]. While these users are experts in their own domain, they may be laymen in the field of data science. To help researchers who lack such skills, a variety of frameworks and toolkits exists [8–11]. These either automatically choose the right presentation for a certain type of data, or allow users to build their own visualisations using a library of widgets.

However, datasets may be incomplete, devices inaccurate and predictions based on incorrect assumptions [12]. This introduces various types of uncertainty, the display of which is not adequately considered

in these powerful toolkits used by users not experienced in visual data analysis. Simply giving a numerical quality estimate of a prediction based on such data may not be enough [13]. It is therefore necessary to not only visualise the classification space of an algorithm used for a prediction or visualisation, but also to explore how different types of uncertainty are best represented to such “laymen” and to provide a visual guide, whose recommendations can be implemented into established frameworks to address this omission.

This is especially important, as the capabilities of data-processing experts and laymen to correctly interpret visualisations and their uncertainty representations may diverge [14]. Whereas data-processing experts may be well-versed in engaging with complex visualisations in Visual Analytics applications, users without a background in data science who want to use such applications for their own research (for example in the humanities) are likely to struggle with interpreting the visualisations and thereby may be unable to use these tools effectively in the sense-making process [15].

[☆] This article was originally submitted to the Journal of Visual Languages and Computing, however during the process of review the journal underwent a name change and is now titled Journal of Computer Languages.

* Corresponding author.

¹ Now at Google.

Steering an algorithm and having the results instantly displayed is a key feature of Visual Analytics applications [1]. To increase user trust and to improve accuracy, showing the algorithm's workings is helpful, as this allows users to better understand how and why a prediction or attribution to a certain class is made [16]. By visually analysing the classification space, users can change algorithm parameters or edit the data set to improve the result of the analysis.

Popular algorithms used in Visual Analytics applications for classification and prediction include K-means clustering and Linear Regression. Due to their simplicity, they are easy to deploy and quick to compute and therefore frequently used. Yet, both may be subject to areal uncertainty, meaning that depending on a data point's distance to the centroid (K-means) or trend line (Regression), a class-attribution or compliance with a prediction may be increasingly uncertain. As a result, visualisations of the mechanics of such algorithms, that are to support the user in the decision-making process, may hold areas in which their validity in representing certain data points is reduced. This in turn can explain why the overall prediction quality is low, or indicate that possible groupings in clusters, especially for certain data points, should be interpreted with caution. This may prompt users to adjust their level of trust into the analysis result or to take action and tighten the algorithm's parameters to iteratively improve the outcome.

For example: If the user sees how far some data points are located from the regression line in a Linear Regression, they can decide to either adjust the weighting of these or to remove them to increase the model's prediction accuracy. Similarly, data points far away from a cluster's centre in a K-means clustering algorithm can be identified and the algorithm adjusted to improve classification accuracy and to reduce "fuzzy" labelling. Therefore, a graphical representation of uncertainty inside the classification space that communicates the "goodness of fit" of various data points to the model, may be much more effective than a numeric indicator of the overall prediction quality [13].

Similarly to the visualisations of algorithm mechanics in Visual Analytics applications, geospatial visualisations may show spatially distributed data which in some areas may not be clearly classified due to imperfect measurements or mixed structures, thereby being subject to "attribute" uncertainty. This type of uncertainty describes the confidence with which an area in the visualisation may be attributed to a certain class, ranging from simple land cover classifications [17] over water salinity levels [18] to demographic constellations [19]. This suggests that geospatial visualisation and algorithm visualisation share a common type of uncertainty for which the former seems to have devised an effective representation through a large body of work that is lacking from the latter.

With both domains having to deal with areal uncertainty, the question arises whether the means to represent uncertainty in the domain of geospatial visualisation could be transferred to the domain of algorithm visualisation. In this context we define geospatial visualisation to be the "source" domain from which the representation of uncertainty is to be transferred. The receiver of this representation, the "target" domain, is to be the domain of Visual Analytics and Information Visualisation, where an established, clear definition of uncertainty representation is still missing.

Taking into account the extension of Visual Analytics to users groups not necessarily familiar with data processing and visualisation [6,7,20], a visual guide that is to enrich existing frameworks and toolkits with adequate representations of uncertainty needs to provide visual means and depictions that are effective and easily understood by laymen. To define the first building blocks of such a guide, we propose a set of research questions.

1.1. Research questions

- Can the most popular visual variables for representing areal uncertainty in geospatial visualisation be successfully applied to representing areal uncertainty in the visualisation of algorithm

mechanics for non-experts?

- Which visual variable is most suitable?
- Do our findings comply with those of previous work in the domain of geospatial visualisation regarding efficiency and user preference?

1.2. Contribution

By exploring the above research questions, this article makes the following contributions:

1. We implement the first exploration of the transferability of uncertainty representation from the domain of geospatial visualisation to that of algorithm-mechanics visualisation, focussing on K-means clustering and Linear Regression.
2. We find that in general, *opacity* may be the most suitable visual variable, followed by *texture*. This corresponds to findings of previous work in geospatial visualisation and demonstrates the transferability of these visual means between the two domains. The efficiency of *contour lines* may be dependent on the algorithm-mechanics to be visualised.
3. We report the first quantitative evaluation of users' ability to correctly retrieve uncertainty values using the variable grid method [21]. We find that this visual means, together with *blur*, is not suitable for reliable quantification of uncertainty in the visualisations of algorithm-mechanics.
4. Our study focuses on laymen's ability to assess uncertainty. The findings may thus be applicable to a range of applications that are not solely used by data-processing experts.
5. With the help of the above points, our work lays the foundations for the definition of an interdisciplinary visual guide for representing uncertainty. This guide can plug into existing frameworks and toolkits to enrich their visualisations and may serve as a basis for further research aiming to define empirically validated representations of other types of uncertainty.

1.3. Structure

This article is structured as follows: In a literature survey, we first review work undertaken in the domains of Information Visualisation and Visual Analytics. We then review uncertainty visualisations in the well-established domain of geospatial visualisation and define five popular visual means that might be transferable to the domain of Visual Analytics.

The main part of the article is formed by a user study, where we explore effectiveness and usability of these visual means by analysing quantitative and subjective data derived from 500 users trying to determine varying degrees of uncertainty in two algorithm visualisations.

After this part, the research questions are answered and our findings defined as building blocks of a visual guide for representing uncertainty. We then discuss our work with regards to that of other researchers to describe how it can be implemented to extend their applications and support users in the decision-making process. The article concludes with the identification of avenues for future work and a discussion of the study's limitations.

2. Previous work

This section will review previous work from the fields of Visual Analytics, visualisation toolkits and languages, and geospatial visualisation to help the reader understand the rationales of our research questions. To simplify the structure, the most relevant work is reviewed separately.

2.1. Visual Analytics and uncertainty

Since the term's original definition by Cook and Thomas [22], the

use of Visual Analytics [1] to support decision making has grown continuously. Visual Analytics extends interaction with traditional Information Visualisation techniques with facilities for updating, steering, and improving the analytic processes. The key objective is to incorporate feedback from end-users to improve an automatic analysis. Daniel Keim, one of the pioneers of Visual Analytics, recently presented his thoughts on its use for gaining insights into linguistic data, quoting Albert Einstein:

“Computers are incredibly fast, accurate and stupid, humans are incredibly slow, inaccurate, and brilliant. Together they are powerful beyond imagination.”

Whereas traditionally Visual Analytics applications have been used by experts in domains such as engineering and finance [1,2], trend analysis in large datasets [23], or the control of process compliance with an ideal model [24], there is an increasing interest in scholars from “soft sciences”, such as learning [3] or digital humanities [6,20], to apply Visual Analytics to support their work.

For the latter field, one particularly interesting piece of work is the interactive visual analysis of German poetics by John et al. [7]. In their work, the authors visualise text alignment and focus on improving alignment algorithms in an iterative and integrated process. Examples from the domain of Information Visualisation include timeline visualisation and geographic maps to study the history of philosophical ideas [25], and a node-link diagram that shows relationships between people [26]. Such set-based visualisations focus on representing relationships between clusters or sets, which is a common task in many visual analysis scenarios [27]. This applies to linguists who may group words into semantic categories when analysing a document, and sociologists who group people into communities to study their relationships when analysing social networks.

Especially the work of John et al. [7] illustrates that visualisation is not only used as a means of presenting the end results of an analytic process, but rather as a fundamental part of the analysis itself, allowing users to evaluate data at all stages of the text-mining and sense-making process: From the initial stages of data exploration and hypothesis evaluation, over user-driven feedback for refining text-mining rules and parameters, to generating new hypotheses and questions.

However, despite the advancement of Visual Analytics and Information Visualisation to these “soft sciences”, little work has been done to enable non-experts in data-processing to conduct and steer complex analysis tasks [5]. While these users are experts in their own domains, they usually have little expertise in data-processing and visualisation, and may be called laymen in this regard. This lack of skills leads to a common pattern where the analytical task is shared by two user roles [28]: In a first step, the domain expert but non data-processing expert (e.g. a humanities scholar or teacher) defines criteria for useful data and requirements for the analysis. In a second step, a data-processing expert – usually ignorant of the domain – is responsible to choose, modify, and integrate automated and visual analysis methods. Although such an approach is feasible and has resulted in several interesting prototypes that can be used by non-technical users (for instance to analyse risk and uncertainty in financial data [29]), it is also constrained by several limitations.

The main issue with decoupling the data-processing design from the data-processing use is the loss of vital information necessary to understand the analysis. The most impactful problem in the visual translation is the suppression of data uncertainty and quality information in favour of visual simplicity and clarity. Yet, dealing with uncertainty and trust in Visual Analytics is nontrivial due to the large amount of noise and missing values originating from heterogeneous data sources, as well as bias introduced by automated analysis methods and human perception [30].

To face this problem, the notion of data quality and the confidence of the analysis algorithm need to be appropriately represented. The analysts need to be aware of the uncertainty and be able to read quality

properties at any stage of the process. Whereas several approaches have been researched under the umbrella of “data wrangling” [31], most of this work concentrates on preprocessing data (i.e. data entry, data (re) formatting, data cleaning, etc.). Yet, what constitutes an error is often context-dependent and requires the human judgement of domain experts [32]. As this is vital to the sense-making process, there is a need to research how such variables can be represented alongside with the outcomes of an iterative analysis. So far, little work has been done on this topic [33].

To address this, Sanyal et al. [34] conducted research into specific techniques for uncertainty visualisation, while Skeels et al. [35] explored what uncertainty itself entails. However, existing work tends to concentrate on preprocessing data before such data is used in applications [31], rather than looking at how to communicate quality indicators to support decision-making by domain experts. Therefore, when aiming to define a visual guide for the representation of uncertainty in Visual Analytics applications, it is imperative to focus on forms that can be correctly interpreted by domain-experts and data-processing experts alike.

2.2. Visual languages and visualisation toolkits

Especially in the domain of computer science, a lot of effort has gone into simplifying complex processes into representations that can be understood by non-experts. A prominent example is the visualisation of algorithm mechanics into flow charts for computer programming classes [36]. By replacing programming instructions with a visual representation that summarises long steps of code and highlights important decision components, researchers have reported an increase in confidence as well as a reduction of errors and task completion times for students [37], especially when these visualisations could be manipulated interactively [38].

On a more formal basis, such a transformation of sets of instructions into graphical modules or glyph-based notations is described by the field of visual languages in computing [39]. With the help of various specification formalisms, visual languages translate complex processes into abstract depictions that can be understood by non-computer scientists and employed by these for otherwise hard-to accomplish tasks.

Using such a visual language and combining building blocks of graphical and textual elements, users can intuitively define behavioural requirements for algorithms [40] and have the process of probabilistic reasoning explained to them in an easy-to-understand manner [41]. This way, non-experts may gain insight and understanding into an otherwise “black box” system [42]. Beyond process explanation, visual languages are also used to conceptually design structures that have several layers of complexity and interactions that are difficult to keep in mind with every step of the design process. Here, the transformation into a symbolic representation hides unnecessary convolution from the design phase, allowing the user to focus while the computer automatically ensures adherence to constructional constraints [43].

In a more casual context, visual languages are used to support laymen with the creation of games. By simplifying game mechanics into a set of graphical, interconnected building blocks that form a library of actions and meanings, non-programmers, such as educators outside the field of computing, can build computer games themselves and focus on content creation and story line, rather than on the acquisition of programming skills [44,45]. The power of visual representation for describing complex coherences is even being employed to support managers in understanding organisational structures and optimising workflows [46].

Similar efforts have been made in the field of Visual Analytics and Information Visualisation to empower laymen with specialist skills through simplification and modularity: Ren et al. present a toolkit for the rapid prototyping of Information Visualisation applications by allowing users to select matching visualisations for their data from a library of widgets [8], as do Elias and Bezerianos [9]. Similarly, Howe

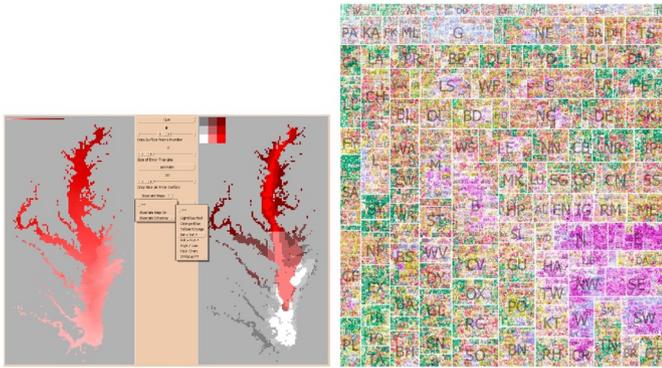


Fig. 1. Left: Figure taken from [47]. Howard and MacEachren map uncertainty levels to colour saturation ranges. **Right:** Figure taken from [52]. Slingsby et al. use colour lightness to represent attribute uncertainty. Permission to reprint has been granted for both figures.

et al. [10] provide a library of visual elements to “assist scientists and researchers in creating interactive visual dashboard applications in seconds with no programming necessary” [10]. More recently Wongsuphasawat et al. [11] presented “Voyager”, a software that automatically suggests matching visualisations for certain types of data to allow laymen to present their findings in an adequate manner, without design knowledge.

However, while allowing laymen to create visualisations, none of the discussed frameworks and toolkits appear to provide visual means for representing uncertainty in their output.

2.3. The representation of uncertainty in geospatial visualisations

A lot of work has been undertaken to represent uncertainty in cartography and geovisualisation, using intrinsic approaches (manipulation of the visual properties of map content) and extrinsic approaches (addition of graphical elements). A typical example of intrinsic representation may be the manipulation of colour: Howard and MacEachren [47] use saturation to depict different levels of uncertainty of nitrogen distribution in a bay (Fig. 1, left). Similarly, Osorio and Brodlie [48] use varying colour hues and noise for representing different degrees of uncertainty relating to oceanic topography. In a

geographic context, Davis et al. [17] use lightness of colour to indicate uncertain classification of land cover types, as do Bastin et al. [49] and Hengl [50], who employ colour “whiteness” for this purpose. Pang [51] discusses uncertainty visualisation for natural hazards and presents work that utilises colour intensity to differentiate between regions with different probabilities of an earthquake occurring. Lightness of colour is also utilised by Ban and Ahlqvist [19] and Slingsby et al. [52], however not for highlighting the uncertainty of land cover type classification, but that of demographic structures (Fig. 1, right). Finally, Berjawi et al. [53] explore the combined use of transparency, a blurred glyph, and a thermometer glyph to express different types of uncertainty on a map visualisation. Although only providing a prototype implementation, they report that opacity works well as a general indicator of uncertainty for a point of interest, whereas detail provided on-demand can benefit from an additional blur graphic as well as a thermometer glyph to express location and attribute uncertainty.

To indicate uncertainty of positional data, Alesheikh et al. [54] employ “probability contours” that show intervals between which an uncertain data point may be placed. Similarly, Dutton [55] uses varying contour widths. Building on these approaches is the work of Edwards and Nelson [56], who explore the use of fill lightness as well as contour lightness for representing areal uncertainty. Varying contours (together with colour) are also used by more contemporary work, such as that of Brodlie et al. [57], who present the use of a “contour band” and a “spaghetti plot” to show alternative dimensions of oceanographic topography, and that of Bloch et al. [58] and Stamen Design [59], who use varying contours and colour to represent possible variations of strength and direction of tropical storms. Lastly, Spiegelhalter [60] employs changing contours to separate different confidence intervals in a visualisation of health institution performance.

Another means to display uncertain geographic boundaries is blur, as used by Burrough [61] for the visualisation of “fuzzy geographical objects”. This is also discussed by MacEachren [62] (Fig. 2) who refers to this as “contour crispness”, which is similar to the display of a “contour band” or the “spaghetti plots” used by Brodlie et al. [57], but with an applied gradient for transparency. Gershon [63] also employs blur to highlight uncertain areas in the distribution of sea-surface temperature data, as do Djurcilov et al. [18] for indicating areas in which the degree of water salinity may be uncertain.

The discussed intrinsic representations of uncertainty have the benefit of integrating data quality into the data’s actual presentation,

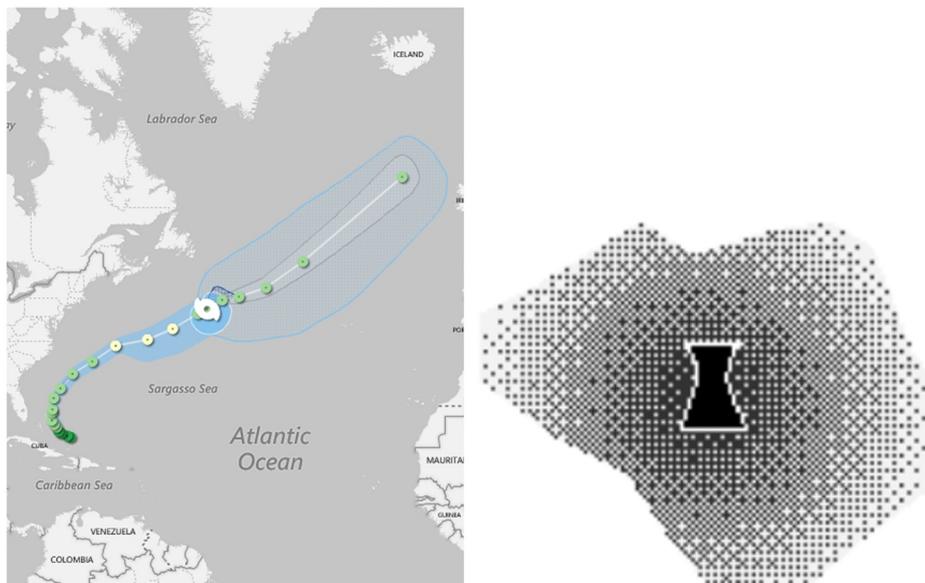


Fig. 2. Figure taken from [62]. MacEachren uses a dithered blur effect to represent the uncertain area of a risk zone. Image is licensed under CC BY 4.0. Permission to reprint has been granted.

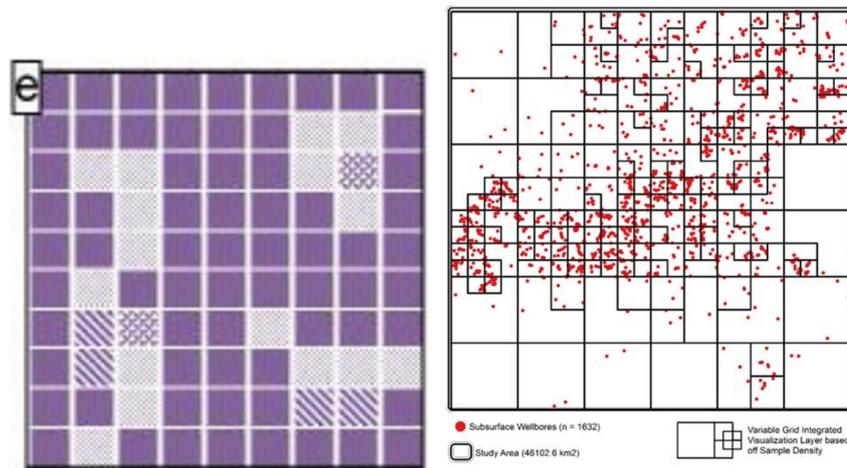


Fig. 3. Left: Figure taken from [64]. Buttenfield shows how texture is used to express attribute uncertainty on an ecological map for bird sightings. Right: Figure taken from [21]. Bauer and Rose use varying grid sizes to represent uncertainty in spatial data. Permission to reprint has been granted for both figures.

reducing clutter by reducing the amount of visual elements. However, extrinsic representations of uncertainty have also been shown to have their merits, and distinction between these two types does not always appear straightforward. For example, Buttenfield [64] discusses the work of Andrlé and Carrol, who use texture for visualising attribute uncertainty of bird sightings (Fig. 3, left). Kinkeldey et al. [65, p. 379] argue that although texture is regarded as an overlay element and thus extrinsic, the visual result of integrating it into a map area may lead to it being categorised as intrinsic. Similarly, Clementini and Di Felice [66] combine a contour band with texture to visualise boundary uncertainty of geographic objects, suggesting this to also be an intrinsic representation of uncertainty. A clearer distinction may be made for the use of a secondary display, such as an adjacent map, where one map displays the data, and another the uncertainty [62]. Other implementations of this concept can be found in the work of Lucieer and Kraak [67], who use multiple map displays to separate the display of the classification of different urban areas from its quality. However, this approach may be criticised for requiring extensive eye movement and reorientation, making some researchers doubt its suitability for simple representations of uncertainty when compared to intrinsic approaches [65].

A technique that attempts to address this problem is the use of grids: Kinkeldey et al. [68] and Cedilnik and Rheingans [69] use a grid laid over a vegetation landcover map. Whereas coloured areas on the map describe the classification type, the “noisiness” of the grid lines represents the uncertainty of class attribution in a certain area, reducing the need for eye movement. Also manipulating the contours of a grid are Hunter and Goodchild [70], who, similar to the “spaghetti plots” [57], use multiple alternative versions laid on top of each other to indicate classification uncertainty in a geospatial visualisation application. Rather than altering grid contours, Bauer and Rose [21] represent uncertainty by varying grid size: A smaller grid unit means high certainty, a larger one low certainty (Fig. 3, right).

Another example of extrinsic uncertainty representation is the use of glyphs – additional objects, such as bars or icons, laid over a map. Following this concept, Cliburn et al. [71] employ coloured glyphs to indicate uncertainties in predicted water balance levels to support decision-making. Sanyal et al. [72] use glyphs together with other visual means to present uncertainty in conglomerates of weather variables. The authors report that experts found this type of visualisation useful for finding outliers in the data. A comparison of several intrinsic and extrinsic uncertainty representations in volumetric data is given by Newman and Lee [73]. They found that multi-point glyphs and ball and arrow glyphs performed best and therefore form an argument for the use of extrinsic representations of uncertainty. However, results were

derived from a sample of 21 research scientists and students with a high degree of computer literacy who cannot be considered laymen in terms of data-processing and visualisation. In addition, Pang [74] suggests that glyphs may be “visually overwhelming”, as focus may be on the glyphs and distract users from the actual data visualisation.

As opposed to static representations of uncertainty, some authors use animation [75] to emphasise areas of low data quality or even sound [76]. Although innovative extensions to the visual layer, differentiation between different uncertainty types may be difficult and may require additional cognitive capacity.

Overall, a plethora of approaches exists for representing uncertainty in the domain of geospatial visualisation. Yet, it remains somewhat unclear which of these might be the most effective and usable. To gain an overview, Kinkeldey et al. [65] present a comprehensive review of uncertainty visualisation studies, focussing on research that measures interpretation accuracy and user confidence. Rather than repeating their findings in detail, we summarise their conclusions:

- transparency tends to be more suitable than colour saturation
- texture (on colour) appears to perform well
- changes in greyscale may provoke more accurate responses than changes in blur or line dashing
- glyphs appear to be effective in 3D visualisations
- grids show promise, but studies lack quantitative data
- using a combined view of data and uncertainty rather than adjacent maps might be preferable due the reduced need for eye movement and better visual linking of data and quality, but could lead to clutter
- animated properties may not perform better than static ones
- intrinsic approaches appear most suitable for quantitative information, extrinsic ones for qualitative information

Despite some trends, the results do not point towards a definitive “champion” for expressing uncertainty. Rather, all techniques seem to have their benefits and limitations, depending on the domain the study was conducted in. Yet, there is evidence that glyphs add clutter to the visual display [74] and that they are better understood by experts than by laymen [14]. Further, animation does not appear to be more effective than static representations [65] and may even cause negative side-effects such as motion sickness or epileptic fits, if not used with caution.

A limitation of the discussed studies may be that the majority appears to focus on experts (despite a few exceptions including laymen) and that they use relatively small numbers of participants [77]. Such a limitation may not come as a surprise, as recruiting experts can be difficult. But with the increasing use of Information Visualisation and

Visual Analytics in consumer applications [78], learning analytics [3], and digital humanities [6,20], it appears worthwhile to explore whether the findings of previous work regarding the performance of certain visual means to represent spatial uncertainty to (mostly) experts in the domain of geospatial visualisation, may also be applicable to laymen in the domain of Visual Analytics.

3. Motivation

The large array of frameworks and toolkits for supporting non-data processing experts in visualising large data sets for exploration and analysis underlines the importance of supporting domain-experts with powerful tools for analysing their findings in an easy-to-understand and visually appealing manner [79].

However, the discussed visual languages and visualisation toolkits lack the graphical representation of uncertainty. Whereas data and prediction results are presented and visualised, their structural problems, inaccuracies, and potential problems in the underlying computational model, are not. This highlights the need for a visual language or toolkit that addresses this important aspect.

With so many well-established frameworks available, it seems wasteful to build a complete approach covering all steps of the visualisation process – this has been done to very high standards in the discussed previous work. Rather, it may be more fruitful to devote time and effort to an “add-on” for these existing toolkits. An add-on, that is independent of implementation and specification formalisms and can therefore “plug-in” to these toolkits in the form of a library of visual guidelines that covers the shortcomings they expose with regards to representing uncertainty. While this could be done in the form of graphical widgets, it seems more flexible and future-proof to leave implementation to the respective authors. Instead, we ought to strive towards providing an empirically founded guideline in the form of a rule-book – the advice of which can be implemented into existing and future applications as required.

3.1. Why transfer knowledge?

When defining a visual language – especially one that is to work as an addition to established ones – it is important to respect existing semantics that stem from well-known, established domains. Doing so will facilitate a novice’s apprehension of such a language, as it will allow them to build on previous experience [80,81]. A promising start may therefore be the examination of the transferability of semantics between domains, where a certain meaning is to be established in the “target” domain that, although needed, may not be as clearly defined and researched as in the “source” domain. By adopting suitable elements as the result of such a transfer, one may capitalise on existing knowledge and meaning and thereby shorten the language’s definition process.

An interesting example for such a transfer can be found in the work of Celentano and Pittarello [82]. The authors employ cartographic metaphors as a means for the design of a visual language for knowledge management. They argue their approach’s helpfulness for representing shared knowledge to a wide audience by employing a shared common ground of meaning, known to users from their everyday lives. Following this, Yusoff and Salim [83] underline the importance of choosing a visualisation that is interpreted the same way among all stakeholders, especially in a collaborative environment. According to their findings, a good visualisation should provide shared social, cognitive, and task-solving support. A similar point is also made by Cybulski et al. [84], who describe how the domain of Visual Analytics in particular can build on primary metaphors derived from everyday experiences and use these effectively in visualisations to support the user’s sense-making process. With regards to a metaphor for uncertainty, Cybulski et al. state that “certainty is firmness” [84], referencing the work of Grady [85]. However, whereas most visual elements can

capitalise on well-defined visual analogies, the concept of uncertainty may not, as it describes a state or feeling, rather than an existing item. But how to express something whose existence is based on absence (in the case of uncertainty stemming from incomplete data), whose reason for being is known, but whose impact on its dependants is as uncertain as the very results given by predictions that consider elements subject to this quantifiable, yet somewhat intangible factor?

Formally describing this multi-faceted entity may be very difficult and, once fathomed, may still lack expressiveness and meaning [13]. The answer, then, may lie in the use of a visual language, where multiple layers of meaning can be compressed into a system-independent graphical representation, that may say so much more than the sum of its alpha-numeric counterparts. Capitalising on previous knowledge and established meaning, we aim to define a visual guide for the representation of uncertainty in Visual Analytics applications.

3.2. Why algorithm mechanics?

Uncertainty can have many sources. However, our first step is to determine the representation of classification and prediction model uncertainty in the visualisation of algorithm mechanics. When we use the term “algorithm mechanics”, we refer to the workings of an algorithm. These can include the distribution of data points in a coordinate system, the algorithm’s assumptions, its interpretations, and its results. Simply put, “algorithm mechanics” comprises what the algorithm does, how it interprets a data set, and how the data points relate to its interpretation. An algorithm visualisation is a depiction of this relationship.

As described in the introduction (Section 1), popular algorithms in Visual Analytics comprise K-means clustering and Linear Regression. Data points classified by these may be subject to attribute uncertainty, meaning that their distance to a centroid or regression line in a visualised spatial distribution impacts the adequacy of their attribution to the respective cluster or regression model. This in turn directly influences the accuracy of the analysis. Yet, simply indicating the overall model accuracy of such an algorithm in a prediction may often not be enough for the user to make an informed decision [13]. Instead, they should be able to inspect the workings of the employed algorithm to learn how well their data fits the classification space.

Being aware of data points whose classification or explanation by a model is less certain, allows users to gain a better understanding of the underlying processes that lead to a particular result and enables them to make adjustments to the algorithm or the data set to gradually improve the analysis. This would be much more difficult with only a numerical representation of the prediction accuracy, for users would be left in the dark as to how this number has been generated and how data points are attributed to a cluster or model [16]. Here, the visualisation of uncertainty in the classification space can directly address this issue and improve the sense-making process. As interactively steering algorithms is a key element of gaining insight [1,5], we decided to focus on this step of the analytical process for our research.

With the domain of geospatial visualisation having a rich history in representing attribute uncertainty in spatial distributions, we aim to define the first building block of our visual guide with the help of an empirical study of this domain’s most popular visual means for this purpose, derived from our review: *Opacity, texture, contour lines, blur, and grid* (using variable grid sizes to display uncertainty [21]). As a result, we formulated a set of research questions as defined in Section 1.1 to whose pursuit the remainder of this article is dedicated.

4. Materials and methods

To evaluate the suitability of the five visual means to represent uncertainty in the visualisation of the mechanics of popular algorithms (Linear Regression and K-means clustering), we conducted an online study with 500 users via the Amazon Mechanical Turk service (AMT).

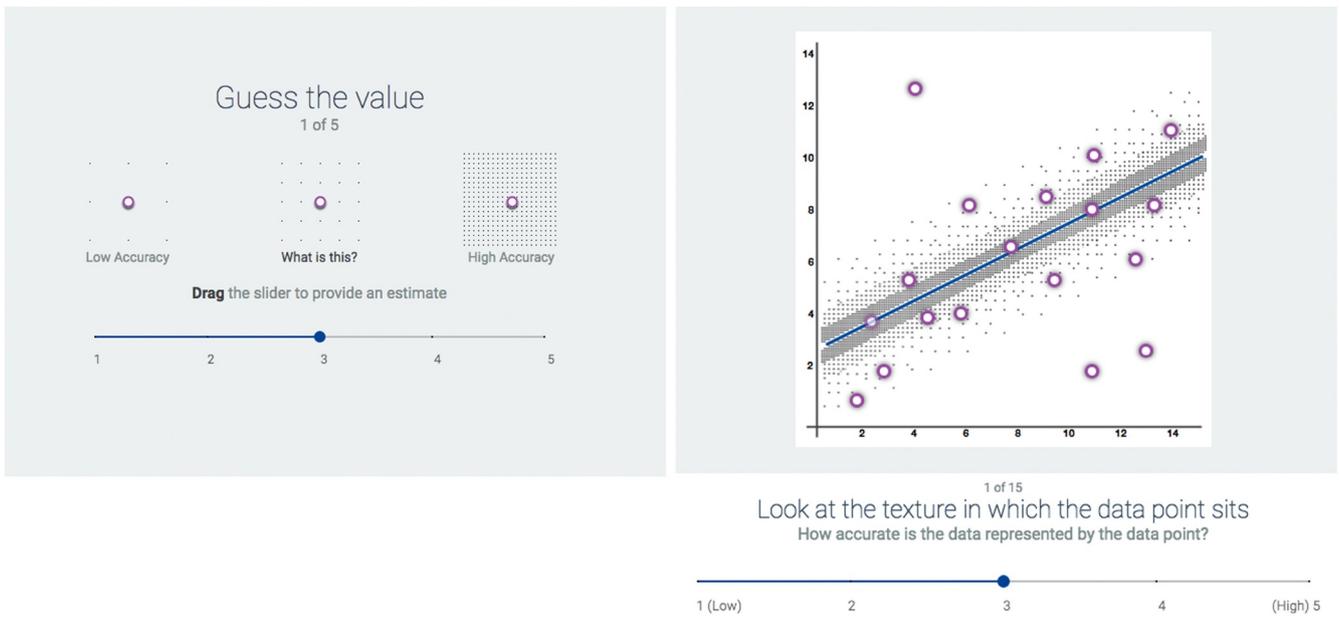


Fig. 4. Left: Tutorial screen for quantifying uncertainty in the *texture* condition. Users had to identify five levels of uncertainty, each corresponding to an increase of 20%. After being shown how to read the degree of uncertainty, users had to complete five test questions (shown) in random order, one per level. **Right:** Task screen of the regression visualisation where uncertainty is represented using changes in *texture* resolution. A user clicks on a blinking dot to then estimate the degree of uncertainty using a slider at the bottom of the screen.

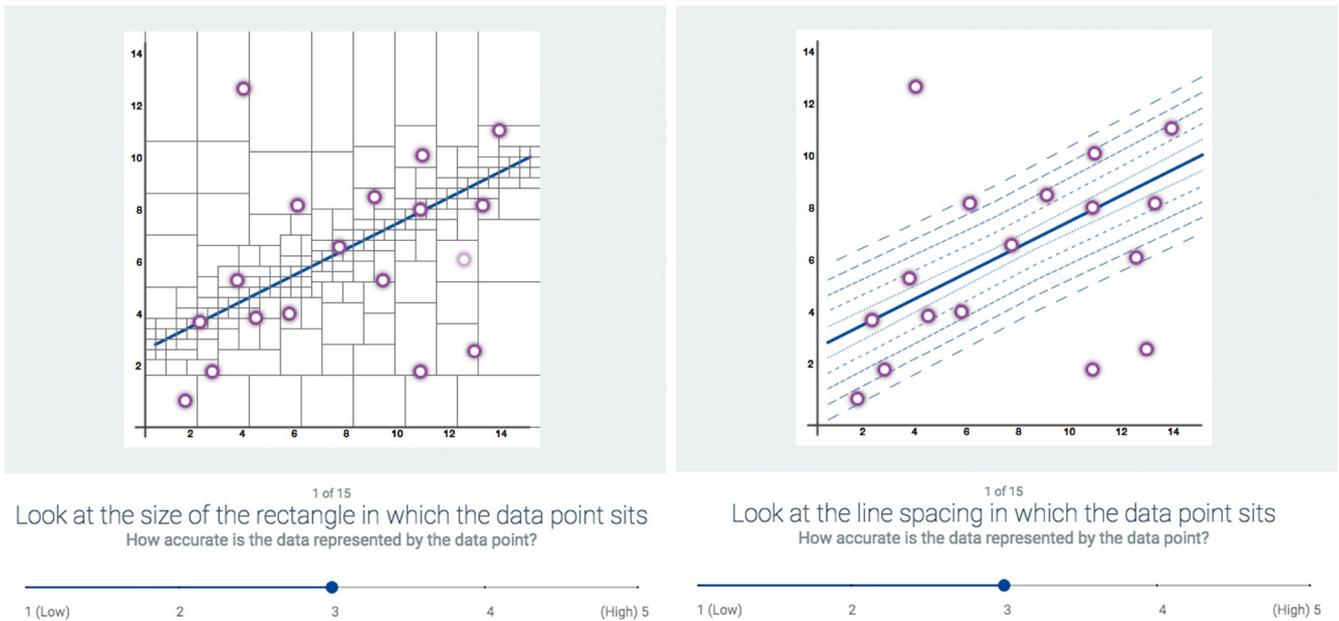


Fig. 5. Left: Task screen of the regression visualisation where uncertainty is represented using changes in *grid* size. A user clicks on a blinking dot to then estimate the degree of uncertainty using a slider at the bottom of the screen. **Right:** Uncertainty visualisation in the regression in the *contour lines* condition.

Respecting AMT’s Terms of Service for worker anonymity, declaration of personal data was optional, providing the following demographics: F: 23%, M: 25%, Unknown: 52%. To qualify for the study, users had to indicate that they had no professional experience in data analysis and that they were reasonably proficient in the English language. In the study, users were shown a visualisation of either a Linear Regression or a K-means clustering algorithm with a distribution of data points (Figs. 4–9). Uncertainty was visualised using the five different visual means *texture*, *opacity*, *blur*, *grid*, and *contour lines*. For the K-means clustering, these indicated the certainty with which a data point may be attributed to a certain cluster. For the Linear Regression, these represented the “goodness of fit” of the regression line for the respective data point.

The points were highlighted at random and users asked to assess their accuracy. Users had to click on the highlighted point to acknowledge the task and then use a slider to enter its accuracy on a level from one (low, 20%) to five (high, 100%). We chose to divide the degree of certainty into five levels of 20%, as previous work suggests that users may struggle with more than six levels of uncertainty [68]. As a result, uncertainty was visualised as follows:

- *Opacity:* Each level of uncertainty was represented by a loss of opacity over five adjacent zones, away from the centroid or regression line (Fig. 6, left, and Fig. 9, left). High levels of opacity meant high accuracy (low uncertainty), low levels meant low accuracy (high uncertainty).

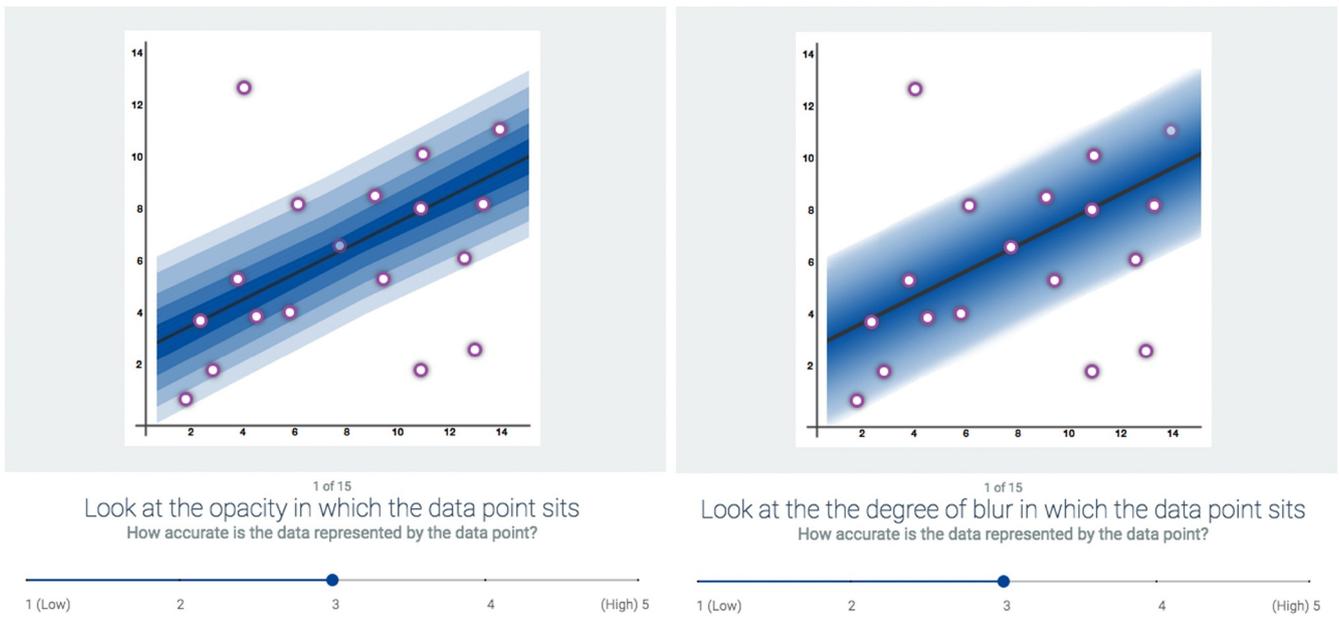


Fig. 6. Left: Task screen of the regression visualisation where uncertainty is represented using changes in *opacity*. A user clicks on a blinking dot to then estimate the degree of uncertainty using a slider at the bottom of the screen. **Right:** Uncertainty representation in the regression in the *blur* condition. The blur was created by calculating a Gaussian blur with an increasing factor for each of the five zones separately and then combining the visual output.

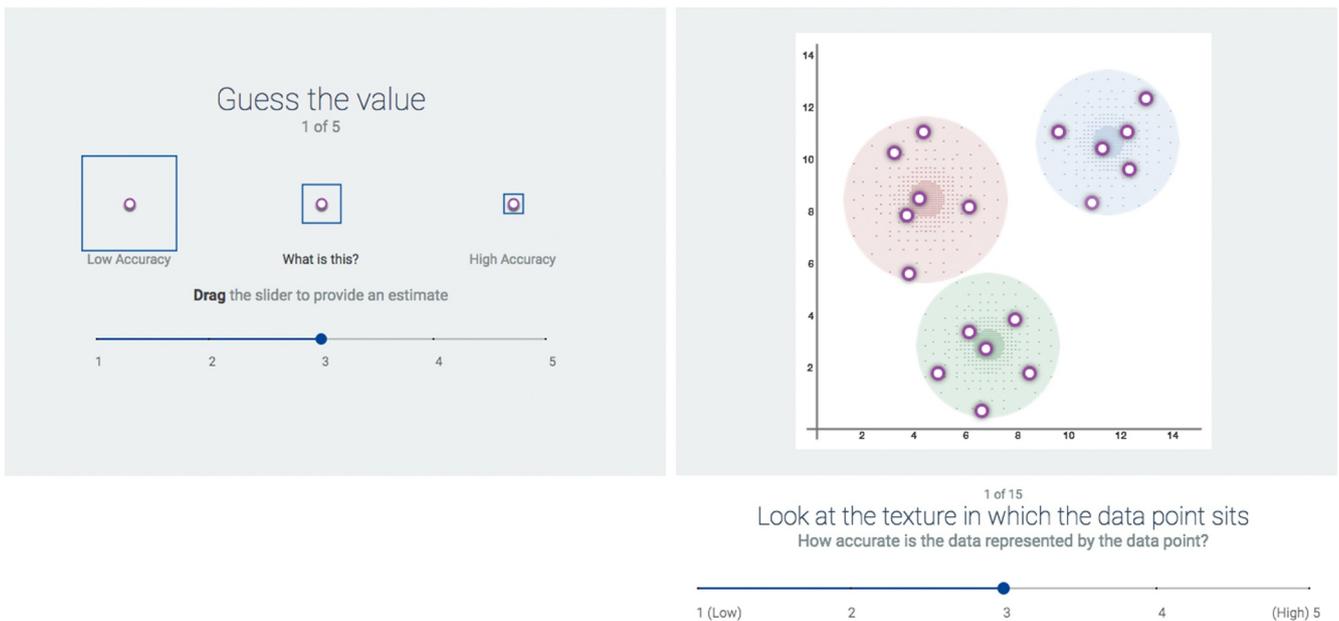


Fig. 7. Left: Tutorial screen for quantifying uncertainty in the *grid* condition. Users had to identify five levels of uncertainty, each corresponding to an increase of 20%. After being shown how to read the degree of uncertainty, users had to complete five test questions (shown) in random order, one per level. **Right:** Task screen of the cluster visualisation where uncertainty is represented using changes in *texture* resolution. A user clicks on a blinking dot to then estimate the degree of uncertainty using a slider at the bottom of the screen.

- **Contour lines:** Each level of uncertainty was represented by a change in the spacing of the dashes of the demarcation lines of five adjacent zones. The gaps between the dashes grew wider, the further the zone was from the centroid or regression line (Fig. 5, right, and Fig. 8, right). Dense dashes meant high accuracy (low uncertainty), widely spaced dashes meant low accuracy (high uncertainty).
- **Texture:** Each level of uncertainty was represented by a loss of texture resolution by a factor of five. This was done over five adjacent zones, away from the centroid or regression line. After a short user test, small adjustments were made to the factors to make the difference between the five resolutions more evident (Fig. 4, right, and Fig. 7, right). A high resolution (dense) texture meant high accuracy

- (low uncertainty), a low resolution texture (sparse) meant low accuracy (high uncertainty).
- **Grid:** Each level of uncertainty was represented by an increase in grid size over five different sizes, away from the centroid or regression line (Fig. 5, left, and Fig. 8, left). A small grid cell meant high accuracy (low uncertainty), a large grid cell meant low accuracy (high uncertainty).
- **Blur:** Each level of uncertainty was represented by an increase in the blurriness of five zones, away from the centroid or regression line. The blur was created by calculating a Gaussian blur with an increasing factor for each of the five zones separately and then combining the visual output (Fig. 6, right, and Fig. 9, right). A low level

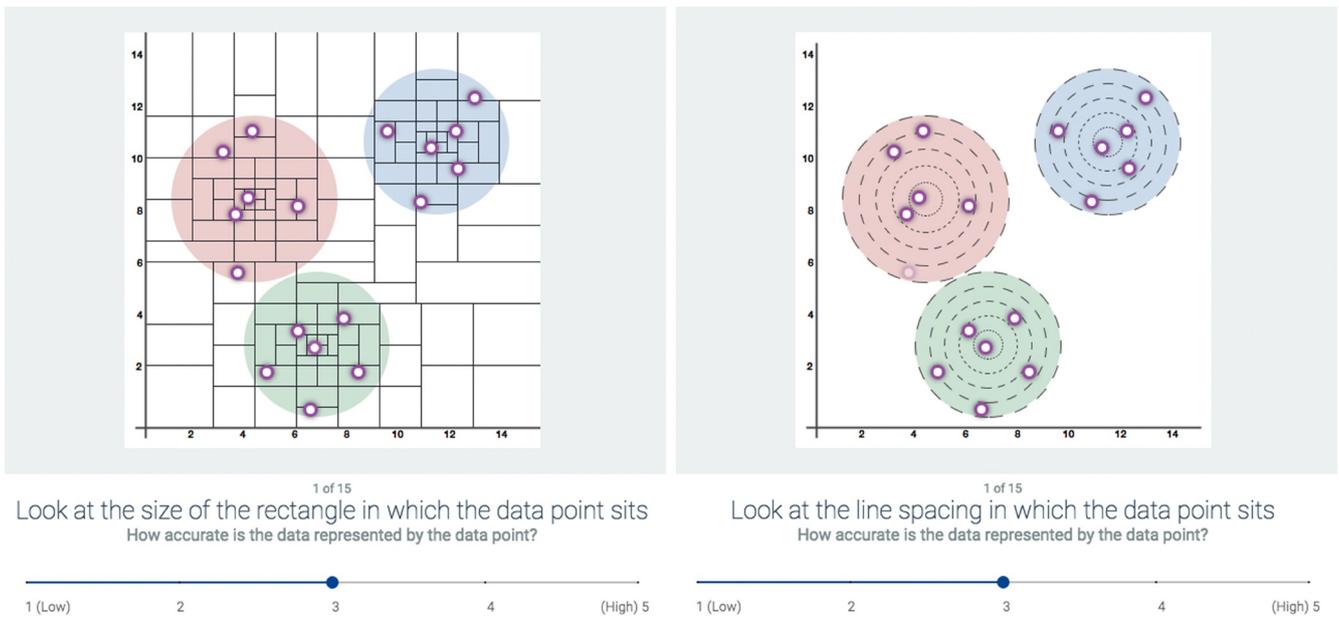


Fig. 8. Left: Task screen of the cluster visualisation where uncertainty is represented using changes in *grid size*. A user clicks on a blinking dot to then estimate the degree of uncertainty using a slider at the bottom of the screen. **Right:** Uncertainty representation in the cluster visualisation in the *contour lines* condition.

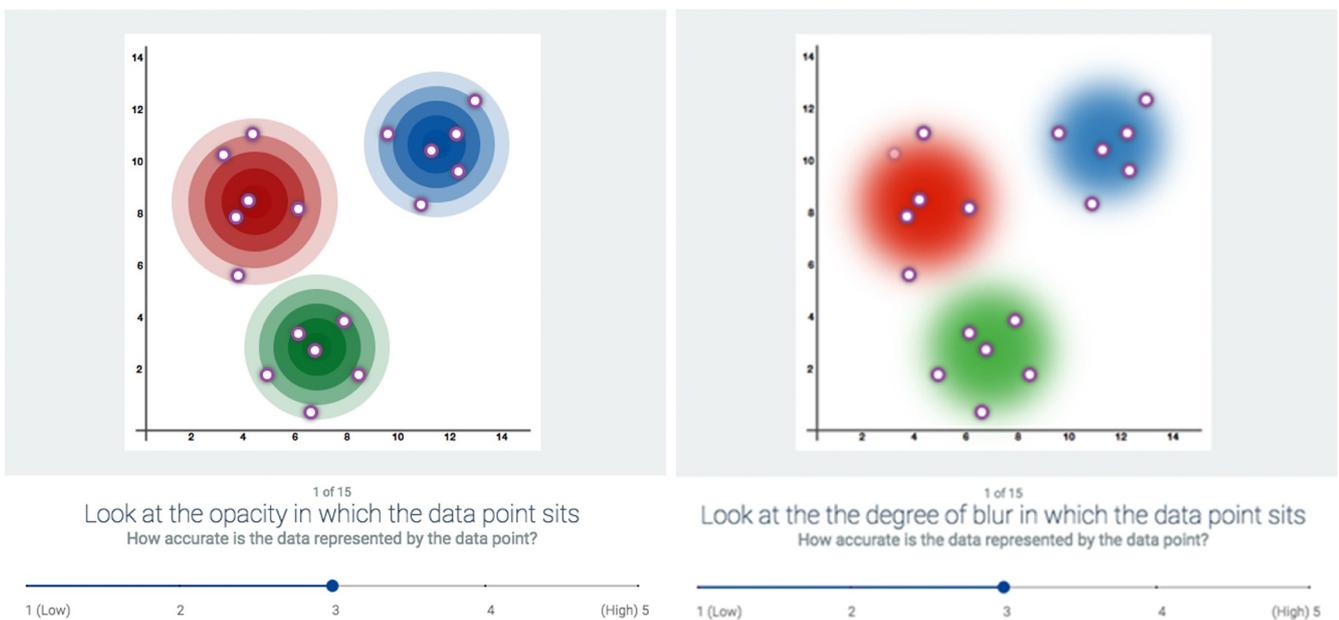


Fig. 9. Left: Task screen of the cluster visualisation where uncertainty is represented using changes in *opacity*. A user clicks on a blinking dot to then estimate the degree of uncertainty using a slider at the bottom of the screen. **Right:** Uncertainty representation in the cluster visualisation in the *blur* condition. The blur was created by calculating a Gaussian blur with an increasing factor for each of the five zones separately, and then combining the visual output.

of blur meant high accuracy (low uncertainty), a high level of blur meant low accuracy (high uncertainty).

The following properties were recorded: The expected response, the given response, and the task completion time. Each level of accuracy had to be assessed three times using three different points in the visualisation, highlighted at random. As we used five levels of uncertainty – each level mapping to an increase of 20% – users had to determine the uncertainty of 15 points in total. Three additional points with the same visual appearance were added to simulate outliers, resulting in a total of 18 visible data points.

The visualisations of the algorithms were created as follows: The Linear Regression used the Ordinary Least Squares Algorithm on a dataset from Anscombe’s quartet (I). The K-means clustering used 15

points from a sample data set from the R statistics package. The visualisations were rendered on a canvas the size of 400×400 px and combined with one of the visual variables. Then we overlaid the data points that users had to assess and added descriptions and controls. The whole application had a size of 630×690 px. Recording of all values started with the highlighting (blinking) of a data point and ended once the user had operated the slider and clicked the “next” button shown below the slider to begin the following task. We conducted a between subjects study, assigning 50 users to each condition. Altogether, the study had 10 conditions: One for each of the five visual variables in either algorithm visualisation (Figs. 4–9).

Before the tasks started, users were shown a tutorial explaining the concept (for example what a large grid size, level of blur, opacity, line style, or texture granularity meant for a certain data point, depending

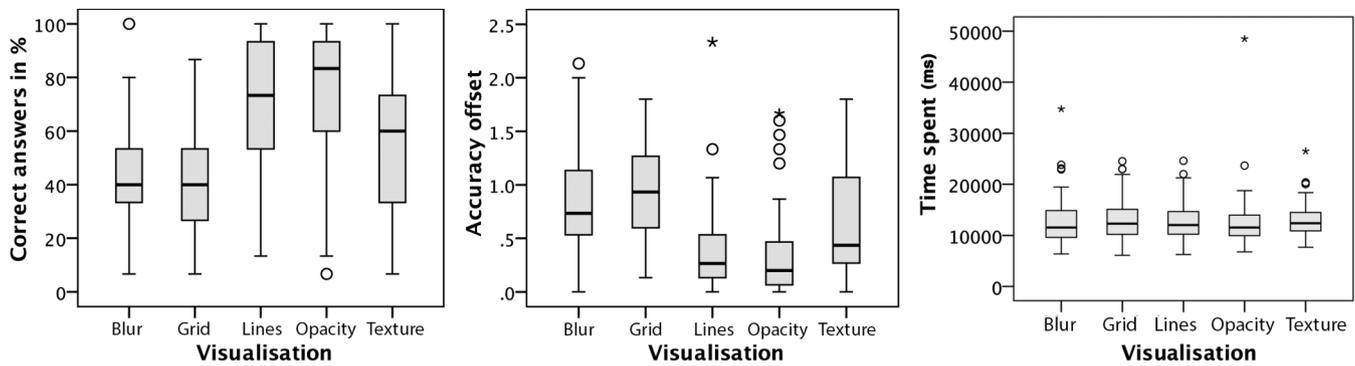


Fig. 10. K-means clustering visualisation: Quantitative results for the five visual variables. **Left:** Median percentage of correctly retrieved uncertainty values. **Middle:** Median accuracy offset. *Opacity* and *contour lines* (lines) are most accurate. **Right:** Median time spent per task per condition in milliseconds.

on the condition) and had to complete five training tasks in random order, one for each level of uncertainty (Fig. 4, left, and Fig. 7, left). The preceding description and tutorial were the only source of training. No key explaining the visual mapping was shown during the main part of the study. To ensure instructions and tasks were clear, we iteratively tested the tutorial and application design with three users from our research group, none of whom spoke English as their first language.

AMT-based studies may be criticised for a reduced validity due to workers “spamming” responses [86]. We therefore limited participation to workers with a minimum HIT approval rate of 90%. Following the suggestion of previous work [86,87], we included five “gold standard questions” randomly into our task list. These were simple mathematical questions, such as “What is 1 + 2?”, which used the same slider as the “normal” tasks for the response to catch inobservant users who simply clicked through the tasks. Users who answered more than one of these “gold standard questions” incorrectly were removed from the data set.

Further, we monitored the answering pattern to validate that spammers did not give the same answer to three consecutive questions. As this procedure lead to the removal of about 15% in three of the conditions, we collected data from an additional 41 participants using the same design to obtain a similar number of valid users per condition, totalling in 482 valid cases. Finally, we monitored the time between task activation and response using two timers. The first timer started as soon as the task started and a data point began to blink. The second timer started once the user had clicked the point and acknowledged the task. We plotted both and compared them to typical task times measured for a control group of five users from our department. No differences were found, suggesting that users did pay attention to completing the tasks considerably.

After users had completed all 20 tasks (15 data point assessments, five “gold standard questions”) they rated the visualisation of uncertainty in a questionnaire using a five-point Likert scale (1: Strongly disagree, 5: Strongly agree) based on the following aspects: Visual appeal, confidence in their decisions, ease of interpretation, and whether they liked that uncertainty in particular was represented in a certain way (preference). We thus followed the example of earlier work investigating the usability of different visual variables in representing uncertainty [88]. The study was running in a browser using JavaScript and HTML. Questionnaires were realised using Google forms. The study took about 25 min to complete and users were rewarded with one Dollar.

5. Results

This section will report the results of the study separately for each algorithm visualisation. In particular, we evaluated **correct interpretation** (how frequently the uncertainty of a data point was correctly interpreted), **accuracy offset** (how “far off” users were with their interpretation), and **time needed** to complete a task. This is followed by the subjective data from the questionnaire, reporting **ease of use**,

confidence, appeal, and preference.

For the vast majority of distributions, either Levene’s test for the equality of variances or the Shapiro–Wilks test turned out to be significant. We therefore chose the Kruskal–Wallis test over the ANOVA for these. By nature, the subjective data collected via a Likert-scale does not meet parametric assumptions. Therefore, a Kruskal–Wallis test was used by default for this.

As a post-hoc test, we chose the non-parametric Mann–Whitney test. Its results were Bonferroni–Holm-corrected, based on the number of comparisons in the respective condition. As the above procedure was applied to all calculations, it will not be repeated in the following text, unless necessary.

5.1. K-means clustering

Correct interpretation: A Kruskal–Wallis test indicated a significant effect of visual means on the correct interpretation (Chi-Square (4) = 61.94, $p < .001$.) A series of post-hoc tests showed that most correct decisions were made when using *opacity* (median 83.33%) to represent uncertainty – more than using *blur* (median 40%), $Z = 5.25$, $p < .001$; more than using *grid* (median 40%), $Z = 5.66$, $p < .001$; more than using *texture* (median 60%), $Z = 3.23$, $p = .001$; but not statistically significantly more than when using *contour lines* (median 73.33%). *Contour lines* elicited more correct interpretations than *blur*, $Z = 5.05$, $p < .001$; more than *grid*, $Z = 5.61$, $p < .001$; more than *texture*, $Z = 2.78$, $p = .006$. The third-highest amount of correct interpretations was found for *texture*, which was interpreted correctly more frequently than *grid*, $Z = 3.34$, $p = .001$, and more frequently than *blur*, $Z = 2.64$, $p = .008$. See Fig. 10, left.

Accuracy offset: The Kruskal–Wallis test showed an effect of visual means on the degree of accuracy offset, Chi-Square(4) = 14.51, $p < .001$. Post-hoc tests revealed *opacity* to have the lowest accuracy offset (median 0.2), meaning users’ interpretation of the degree of uncertainty of a data point was closest to the data point’s actual degree of uncertainty using this visual means. The offset was lower than that of *texture* (median 0.44), $Z = 3.36$, $p = .001$; lower than that of *blur* (median 0.73), $Z = 5.17$, $p < .001$; lower than that of *grid* (median 0.93), $Z = 5.88$, $p < .001$. The visual variable with the second lowest accuracy offset was *contour lines* (median 0.27), which was lower than that of *blur*, $Z = 5.17$, $p < .001$; lower than that of *texture*, $Z = 2.83$, $p = .005$; lower than that of *grid*, $Z = 5.9$, $p < .001$. The means with the third lowest accuracy offset was *texture*, which had a lower offset than *blur*, $Z = 2.63$, $p = .009$, and a lower offset than *grid*, $Z = 3.58$, $p < .001$. See Fig. 10, middle.

Time needed: A Kruskal–Wallis test did not reveal any statistically significant differences in task completion time between the different visual means. See Fig. 10, right.

Ease of use: The Kruskal–Wallis test indicated an effect of visual means on the ease of interpretation, Chi-Square(4) = 30.07, $p < .001$. The post-hoc tests showed that *opacity* (median 4.5) was rated the

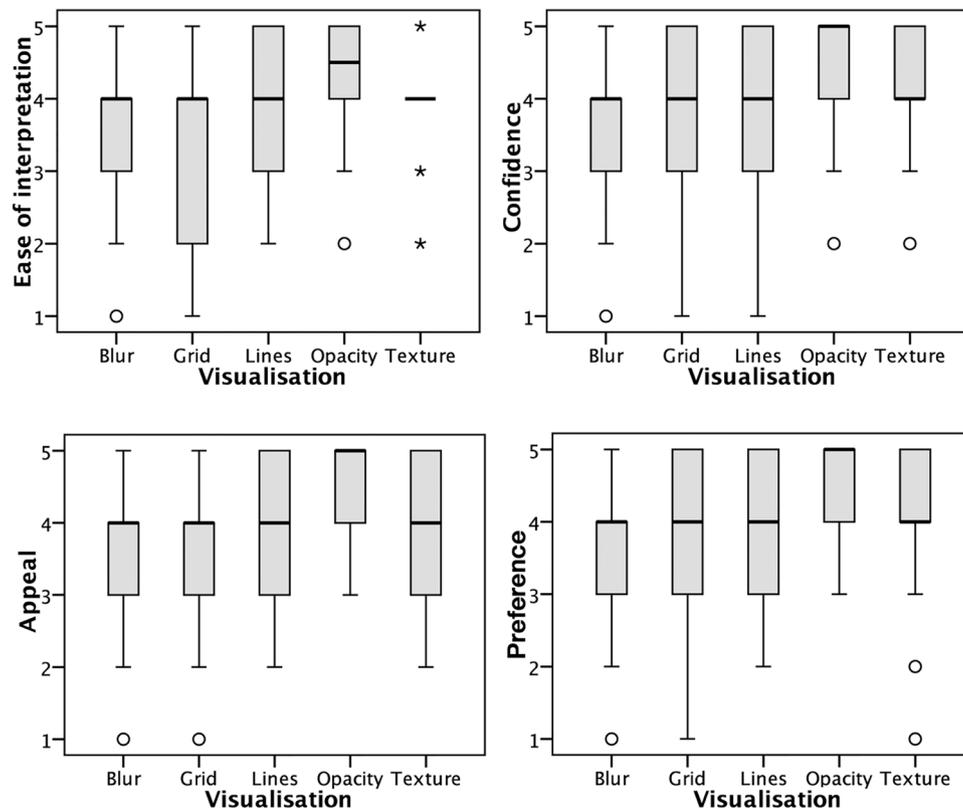


Fig. 11. K-means clustering visualisation: Boxplots of the subjective feedback given on a five-point Likert scale for each of the five visual variables. **Top:** Ease of use and confidence. **Bottom:** Appeal and preference. Overall, *opacity* and *texture* received the most positive feedback.

easiest visual variable to interpret uncertainty – easier than *blur* (median 4), $Z = 4.69$, $p < .001$; easier than *grid* (median 4), $Z = 4.25$, $p < .001$; easier than *texture* (median 4), $Z = 3.5$, $p < .001$; easier than *contour lines* (median 4), $Z = 4$, $p = .007$. The test also indicated differences in the distribution of the responses between *contour lines* and *grid*, *contour lines* and *blur*, and *texture* and *blur*, but these were not statistically significant after the Bonferroni–Holm correction. See Fig. 11.

Confidence: A Kruskal–Wallis test showed an effect of visual means on users’ confidence to have made the right decision, Chi-Square(4) = 26.9, $p < .001$. A series of post-hoc tests revealed that users felt most confident when using *opacity* (median 5) – more confident than when using *blur* (median 4, $Z = 4.77$, $p < .001$), *grid* (median 4, $Z = 3.68$, $p < .001$), and *contour lines* (median 4, $Z = 3.04$, $p = .002$). Differences in the response distribution between *texture* and *blur*, *texture* and *opacity*, and *contour lines* and *blur* were also found, but these were not statistically significant after the Bonferroni–Holm correction. See Fig. 11.

Appeal: The Kruskal–Wallis test indicated that users rated the visual means to have a differing visual appeal, Chi-Square(4) = 18.2, $p = .001$. The post-hoc tests showed a tendency for *opacity* (median 5) to be judged the most appealing – more appealing than *blur* (median 4, $Z = 3.61$, $p < .001$) and more appealing than *grid* (median 4, $Z = 3.64$, $p < .001$). Differences between *opacity* and *contour lines* were also found, but were not statistically significant after the Bonferroni–Holm correction. See Fig. 11.

Preference: A Kruskal–Wallis test indicated an effect of visual means on users stating they liked it as a way to represent uncertainty, Chi-Square(4) = 17.1, $p = .002$. The post-hoc tests hinted at differences between the responses given for the various means, but these were not found to be statistically significant after the Bonferroni–Holm correction. See Fig. 11.

5.2. Linear Regression

Correct interpretation: A Kruskal–Wallis test indicated a significant effect of visual means on the correct interpretation of a data point’s uncertainty (Chi-Square(4) = 58.21, $p < .001$.) A series of post-hoc tests showed that more correct decisions were made when using *opacity* (median 93.3%) – more than when using *blur* (median 46.67%, $Z = 5.49$, $p < .001$), *grid* (median 46.67%, $Z = 6.67$, $p < .001$), *texture* (median 60%, $Z = 5.78$, $p < .001$), and *contour lines* (median 46.67%, $Z = 6.1$, $p < .001$). Differences between the distributions of *texture* and *grid* were also indicated, but turned out not to be statistically significant after the Bonferroni–Holm correction. See Fig. 12, left.

Accuracy offset: The Kruskal–Wallis test revealed an effect of visual means on the degree of accuracy offset, Chi-Square(4) = 58.44, $p < .001$. The post-hoc tests showed *opacity* to have the lowest accuracy offset (median .07), meaning users’ interpretation of the degree of uncertainty of a data point was closest to the data point’s actual degree of uncertainty using *opacity*. It was thus significantly more accurate than *blur* (median .67, $Z = 5.31$, $p < .001$), *grid* (median .6, $Z = 6.56$, $p < .001$), *contour lines* (median 0.8, $Z = 5.92$, $p < .001$), and *texture* (median 0.47, $Z = 5.45$, $p < .001$). *Texture* tended to be more accurate than *grid* and *contour lines*, but differences were not significant after the Bonferroni–Holm correction. See Fig. 12, middle.

Time needed: Levene’s test was not significant and an ANOVA did not show a statistically significant difference in completion time between the visual means. See Fig. 12, right.

Ease of use: The Kruskal–Wallis test indicated an effect of visual means on the ease of interpretation, Chi-Square(4) = 21.3, $p < .001$. Despite all visual means having a median ease-of-use of four, the post-hoc tests revealed a significant difference in the shape of the distribution of the given responses. Fig. 13 shows that *opacity* was more frequently rated to be easy to interpret than *grid* ($Z = 3.38$, $p = .001$),

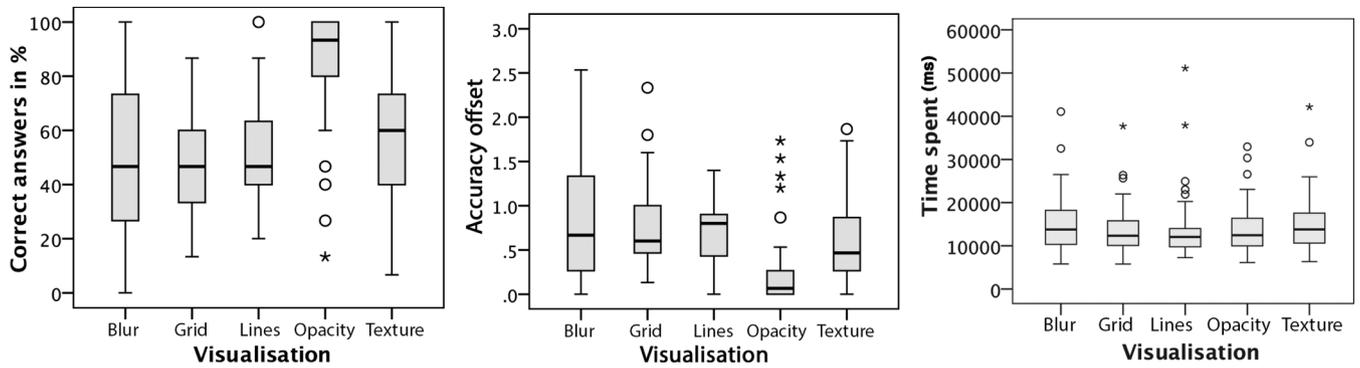


Fig. 12. Regression visualisation: Quantitative results for the five visual variables. **Left:** Median percentage of correctly retrieved uncertainty values. **Middle:** Median accuracy offset. *Opacity* and *texture* are most accurate. **Right:** Median time spent per task and condition in milliseconds.

contour lines ($Z = 3.4, p = .001$), and *blur* ($Z = 2.94, p = .003$). Similarly, *texture* was more frequently rated as easy-to-interpret than *grid* ($Z = 2.7, p = .007$) and *contour lines* ($Z = 2.7, p = .007$). Differences between *texture* and *blur* were not statistically significant after the Bonferroni–Holm correction. See Fig. 13.

Confidence: A Kruskal–Wallis test showed an effect of visual means on users’ confidence to have made the right decision, Chi-Square(4) = 17.43, $p = .002$. The post-hoc tests revealed that despite equal medians (4) for all visual means, users tended to have a higher confidence in their judgement more frequently when using *opacity* than using *blur* ($Z = 3.5, p < .001$) and *contour lines* ($Z = 3.14, p = .002$). Differences in answer distribution between *texture* and *blur*, as well as *texture* and *contour lines*, and *opacity* and *grid* were not significant after the Bonferroni–Holm correction. See Fig. 13.

Appeal: A Kruskal–Wallis test indicated that users rated the visual means to have a different visual appeal, Chi-Square(4) = 23.88, $p < .001$. A series of post-hoc tests showed that users rated *opacity* (median

4) more frequently to be visually appealing than *grid* (median 3), $Z = 4.1, p < .001$, and *contour lines* (median 4), $Z = 3.62, p < .001$. Similarly, *texture* (median 4) was more frequently rated as visually appealing than *grid*, $Z = 2.93, p = .003$. Differences between *texture* and *contour lines* were not statistically significant after the Bonferroni–Holm correction. See Fig. 13.

Preference: A Kruskal–Wallis test indicated an effect of visual means on users stating that they liked it as a way to represent uncertainty, Chi-Square(4) = 17.69, $p = .001$. Despite equal medians for all visual variables (4), post-hoc tests showed that users tended to express a higher degree of preference more frequently for *opacity* than for *blur* ($Z = 3.19, p = .001$) and *contour lines*, $Z = 2.93, p = .003$, and more frequently for *texture* than for *blur*, $Z = 2.84, p = .005$. Similar trends were also revealed for *texture* and *opacity* and the other visual variables, but differences were not significant after the Bonferroni–Holm correction. See Fig. 13.

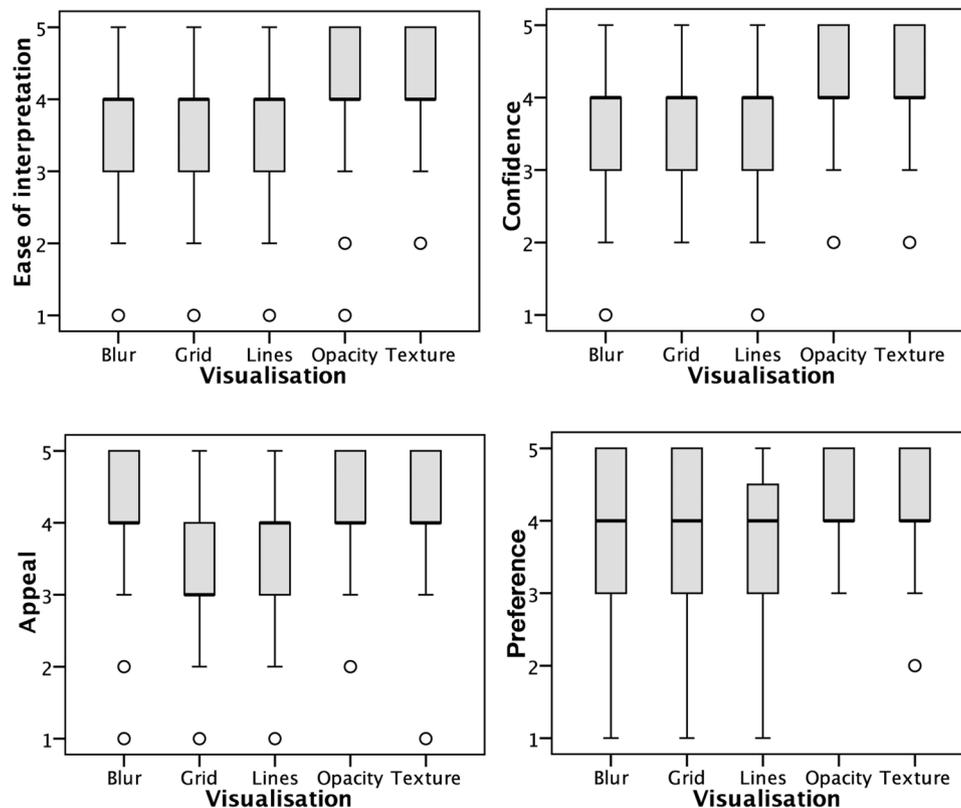


Fig. 13. Regression visualisation: Boxplots of the subjective feedback given on a five-point Likert scale for each of the five visual variables. **Top:** Ease of use and confidence. **Bottom:** Appeal and preference. Overall, *opacity* and *texture* received the most positive feedback.

6. Discussion

Findings are discussed separately for each algorithm visualisation, followed by research questions and limitations.

6.1. K-means clustering

The quantitative results suggest that *opacity* and *contour lines* are the visual means with the highest frequency of correct interpretation. Users estimated the degree of uncertainty correctly using *opacity* in 83.3% of cases and in 73.3% of cases for *contour lines* in comparison to *texture* (60%) and *blur* and *grid* (both 40%).

These results are reflected by those regarding the actual accuracy of users' guesses: The "offset" of these to the actual degree of uncertainty (spanning a total of five levels) was lowest using *opacity* (0.2 levels median offset), followed by *contour lines* (0.26 levels), *texture* (0.44 levels), *blur* (0.73 levels), and *grid* (0.93 levels). However, differences in completion time between the visual means do not appear to exist. The good performance of *opacity* and *contour lines* in the cluster visualisation seems likely to be influenced by their clearly marked regions, offering a simple means of comparison between the different uncertainty zones. In contrast, these are harder to discern in the *texture* and *grid* condition, and hardly at all in the *blur* condition.

The quantitative performance of the visual means for representing uncertainty is partially mirrored by users' subjective feedback for these. Overall, users tended to rate *opacity* as easiest to use, with the greatest appeal, confidence, and preference. However, differences between the other visualisations appear less clear. See Fig. 11.

6.2. Linear Regression

The quantitative results indicate that *opacity* is the visual means with the highest frequency of correct uncertainty estimation. Users interpreted the degree of uncertainty correctly using *opacity* in 93.3% of cases and in 60% of cases for *texture* in comparison to *blur*, *contour lines*, and *grid*, all with a median correct estimation of 46.7%.

The frequency of correct estimation is mirrored by the accuracy offset of the different visual means. It is lowest for *opacity* (0.067), followed by *texture* (0.47), *grid* (0.6), *blur* (0.67), and *lines* (0.8). No differences in task completion time were found. Similar to the cluster visualisation, the good performance of *opacity* seems to be founded in its clear depiction of the different zones. However, this is also given in the *contour lines* and *texture* conditions – yet, accuracy in these is much lower. This is surprising, as *contour lines* performed comparatively well in the cluster visualisation.

Despite similar medians in the subjective feedback, the Kruskal–Wallis tests indicated a different distribution of responses for the various visual means, suggesting that the means which performed well quantitatively, were also perceived positively more frequently. See Fig. 13.

6.3. Answering the research questions

Can the most popular visual variables for representing areal uncertainty in geospatial visualisation be successfully applied to representing areal uncertainty in the visualisation of algorithm mechanics for non-experts?

The answer to this questions is two-fold: The results measured for *opacity* suggest that this visual means can easily be transferred between the two domains for representing quantifiable uncertainty, regardless of the algorithm visualisation. *Texture* had the same performance in both (median 60% accuracy) and thereby also seems to be applicable to this domain, albeit with less confidence.

However, other visual means appear less suitable. While *contour lines* performed well in the K-means clustering visualisation, they performed poorly in the Linear Regression. *Blur* and *grid* performed poorly

in either visualisation, suggesting that these two means may not be easily transferable.

We conclude that not all of the examined visual means may be applied to the domain of algorithm visualisation when aiming to quantify uncertainty, but that applicability may depend on the type of algorithm to be visualised.

Which visual variable is the most suitable?

Overall, the most suitable visual variable to correctly interpret uncertainty is *opacity*, likely due to its clear demarcation of uncertainty intervals. For a cluster visualisation, the results suggest the following ranking: *Opacity* (83.33%), *contour lines* (73.33%), *texture* (60%), *blur* (40%), *grid* (40%). For a Linear Regression, the order is *Opacity* (93.3%), *texture* (60%), *blur* (46.67%), *grid* (46.67%), *contour lines* (46.67%).

Do our findings comply with those of previous work in the domain of geospatial visualisation regarding efficiency and user preference?

In general, the results correspond to reports of previous work, as our findings verify the general link between performance and preference reported by Senaratne et al. [88]. Other commonalities and differences are as follows:

The positive results observed for *opacity* appear congruent with findings of researchers in the domain of geospatial visualisation [65]. The results also correspond to the reports on *texture*, which had been shown to perform well [89], but with an accuracy of 60% in both algorithm visualisations not as well as *opacity*, again corresponding to Kinkeldey et al.'s report [65].

The performance of *contour lines* differs from previous findings. In the case of the K-means visualisation, performance was better than *blur*, thereby contradicting the findings of Boukhelifa et al. [90], who found the opposite in terms of user preference when manipulating the "sketchiness" of contour lines. However, in the Linear Regression visualisation, performance of *blur* and *lines* was similar (as was user preference).

Little quantitative data is available on the performance of the variable *grid* [21]. Our findings suggest that this is low in both algorithm visualisations and therefore best avoided if *opacity* or *texture* are available. However, it may be useful as a more general indicator of uncertainty, but future work is required to validate this assumption.

MacEachren et al. [91] report a high degree of intuitiveness for *blur*. While our studies indicate that subjective feedback for this visual means was not negative (median 4 in all aspects), quantitative performance was poor, thus contradicting the reports of Boukhelifa et al. [90]. Under the study conditions it appears that *blur* may best be used as a general indicator of uncertainty, rather than a quantifier. Yet, it has to be taken into account that our study has focussed on laymen. As previous work has predominantly focussed on experts, the results may not be directly comparable.

6.4. Towards a visual guide for representing uncertainty

By answering the research questions, the work undertaken in this article has laid the foundations for our visual guide for representing uncertainty. Through an empirical study, we found that changes in *opacity* over multiple "zones of confidence" are a well-performing and well-perceived visual means for visualising quantifiable uncertainty in the depiction of the mechanics of a K-means clustering and a Linear Regression algorithm. The second-most suitable visual variables for this purpose are *contour lines* and *texture* respectively, albeit not as effective.

Using a visual representation of uncertainty when visualising an algorithm's mechanics, allows users of Visual Analytics applications to inspect how well their data fits a model. Giving just a numerical estimate of the prediction quality may neither be very expressive [13], nor helpful with improving the analysis. By directly visualising the algorithm mechanics and the classification or prediction accuracy for data points in certain zones, users of these applications can make a much better and much more informed decision. Allowing a user to investigate

how and why a classification or model has been constructed will increase their understanding and support them directly in the decision-making process [16]. Using the visualisation, they may either tweak the algorithm's settings to better match the data, or decide to remove outliers from the set and thereby increase the model's accuracy. Therefore, visualising uncertainty, especially in this part of the analysis, can provide the user with vital information and directly contributes to improving the iterative analysis.

As a result, the findings of our study are promising candidates for the first entries of our guidelines. In particular, we recommend *opacity* (and *contour lines* or *texture*, depending on the algorithm) for re-presenting classification and model uncertainty in the visualisation of algorithms that employ a spatial distribution of data points as their classification space.

While effective for interpretation by the user, *opacity* may also be easy to parse by the machine, as Visual Analytics is a bilateral affair [92]: A computer could read a pixel's alpha value (altered by the user) to determine the degree of uncertainty. Similarly, an algorithm using a human visual system model [93] may interpret areas of varying opacity as uncertain and adjust its classification result. This way, *opacity* could work well for both entities, human and machine. It thereby represents a powerful first item in a possible library of visual means that support researchers and practitioners in communicating uncertainty effectively.

However, due to the great number of different types of uncertainty in Visual Analytics applications, the results of our study only represent a first start to a growing catalogue of empirically validated visual variables. Bit by bit we aim to extend our guide to ultimately provide a comprehensive library that can be used as an addition to existing frameworks and toolkits to express all types of uncertainties in visualisations created with these. Fig. 14 shows how our visual guide may be combined with such a toolkit or framework to extend it with well-researched uncertainty representations.

Although only examined for communicating classification and prediction uncertainty in spatial distributions and thus requiring more research to validate its potential for more general application, employing *opacity* in the wider Visual Analytics process to represent uncertainty may work as follows for the three main directions of communication in an application:

- **Computer to human:** In data visualisations, points far away from the distributional mean or with unclear properties could use varying degrees of transparency. Similarly, attribute selectors showing the most important attributes contributing to a classification algorithm could use changes in opacity to represent quality and impact. Further, opacity “stacks” well: In additive process visualisations (such as the depiction of a decision tree or flow chart) it may be useful to visualise diffusion and culmination of uncertainty in certain layers built from aggregated sources.
- **Human to computer:** Using varying degrees of transparency, the user may mark certain parts in the data visualisation that they find less relevant or not representative due to their domain knowledge. The machine may then consider these differently in calculations. Sensing the user's possible goal, the computer could use transparency to present alternative results that may match a certain intent or bias the machine has detected in the user's behaviour.
- **Human to human:** The data-processing expert may mark problems or areas of questionable relevance in the visualisation of data sets using transparency to support the analysis of the domain expert and vice versa.

6.5. Extending previous work

Using concrete examples, this section will discuss how our visual guide may be combined with a range of existing toolkits and frameworks to extend these with a much-needed module for uncertainty representation, thereby providing hitherto lacking functionality for these,

increasing their scope and power. We report this separately for data visualisation toolkits and frameworks, and specific algorithm visualisation approaches.

6.5.1. Application in frameworks and toolkits for Information Visualisation and Visual Analytics

In a JVLC special issue on “Information Visualization in Machine Learning and Applications” Ren et al. [8] presented *DaisyViz*, a toolkit for the rapid prototyping of Information Visualisation applications for laymen. However, representations of uncertainty in data or prediction models powering the visualisations of the toolkit were not included, making it difficult to interpret the output correctly. Following this, regression or cluster-based visualisations created with the tool could directly benefit from our findings and use *opacity* to visualise these uncertainties, as our visual library provides the missing representations.

Other good examples of how existing applications and frameworks for data visualisation can benefit from our visual guide are *VizDeck* [10], *Exploration Views* [9], and *Voyager* [11]. All three allow novices to transform their data into powerful visualisations for presentation and extended analysis, but lack uncertainty depiction for algorithms. As our visual guide is designed to be an addition to existing frameworks that acts as a point of reference for the visualisation of uncertainty, it can directly plug into these and extend their visualisations for the benefit of the user (Fig. 14).

Andrienko et al. [94] employ a Growth Ring Map [95] to visualise the spatio-temporal distribution of a set of photos obtained from Flickr. The results are circular clusters of differing colour, distributed over a map to show where and when photos were taken. However, the visualised data could have inaccurate or missing time stamps or geo-coordinates caused by inaccurate measuring devices or user neglect. Yet, the approach does not cater for the visualisation of such uncertainties. Similar to our visualisation of uncertainty in the K-means clustering algorithm, *opacity* could be applied to parts of the Growth Ring where these factors are an issue and thereby help users to interpret the data correctly using our visual guide.

In addition, the authors extract movement patterns from the data visualised in the Growth Rings and use vectors to represent common trajectories to analyse movement. To be able to compare trajectories, Andrienko et al. generalise locations by transforming them into areas. This transformation can lead to uncertainty in the groupings and resulting trajectories, stemming from locations that lie between two areas and whose attribution to one or the other may be fuzzy. Our visual guide could be helpful to visualise this uncertainty by rendering confidence intervals around the main vector (as exemplified by our visualisation of a Linear Regression) using varying degrees of *opacity* and thereby improve users' sense-making process, allowing them to interpret the data under consideration of potential uncertainties.

Avoiding cognitive overload in Visual Analytics applications is an important issue to consider when making these more approachable to laymen [15]. Bertini and Santucci [96] introduce a framework for automatic clutter reduction in such applications by parsing the visualisation for possibly irrelevant artefacts based on their importance in the data. However, the authors do not include the optimisation of uncertainty representation. As glyph representation of uncertainty is valued by experts, but less so by laymen [72], Bertini and Santucci could extend the functionality of their framework by using our results, which highlight the effectiveness of intrinsic uncertainty representation for certain algorithm mechanics: To improve the display for laymen in particular, their extended methodology would not only clear visual clutter caused by unimportant data, it would also improve perceivability and scanability by non-data processing experts by translating glyph-based, extrinsic uncertainty representations into *opacity*-based, intrinsic representations [74].

6.5.2. Application in algorithm visualisation

Clementini [97] developed a geometric model for representing

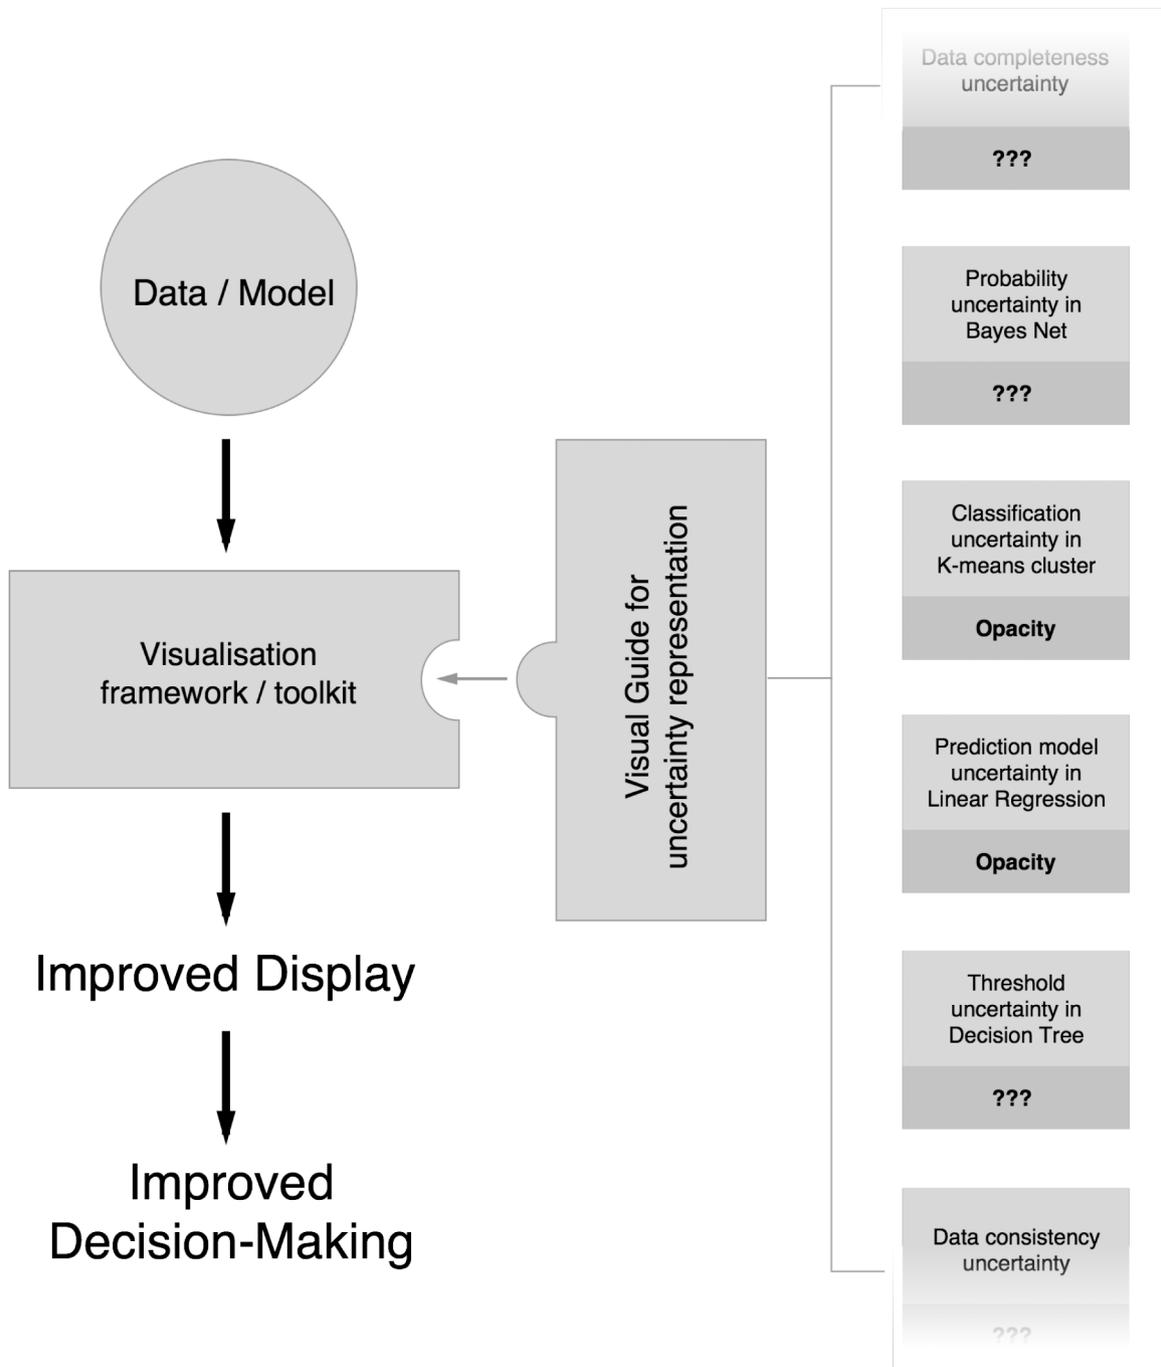


Fig. 14. Depiction of the role of the proposed visual guide for representing uncertainty in the visualisation process. Designed to be an addition to existing frameworks and toolkits, the guide “plugs” into these to extend their visual repository and to provide lacking representations. **Typical flow:** The user provides data which is either directly visualised or used to build a model for an algorithm. The chosen framework applies the most fitting visualisation automatically or manually from its existing repository (as in [9–11]). Depending on the type of data or model visualised, an adequate representation of any uncertainties is automatically applied using the “plugged-in” visual guide, based on empirical evidence. The display is enriched with this additional layer of information and the decision-making thus improved. While the current repository only provides guidelines for representing uncertainty in K-means clustering and Linear Regression visualisation, future work will extend the catalogue to cover additional types of uncertainty, such as that in different types of algorithms or the data itself. These yet-to-be-defined representations are shown as three question marks (???) in the figure.

uncertainty in “spatial objects of linear type” [97]. Such a visualisation might be used for visualising uncertainty in the boundaries of spatial algorithms, such as K-means clustering or Linear Regression, but no user study has been provided. Our work, however, investigated user performance for interpreting uncertainty in such algorithms using *contour lines* and showed limited retrieval accuracy in comparison to *opacity*. It thereby may offer a first evaluation of this means’ usability for this purpose and suggests that while Clementini’s approach may be

powerful, usability seems limited.

In another visualisation, Gong et al. [98] employ *blur* to visualise and calculate probability issues of neighbouring locations in Voronoi diagrams. While effective for calculation, the authors state that they still need to validate the cognitive performance of their blur-based visualisation. Our work undertaken in this article for a K-means clustering visualisation may address this by offering such an evaluation through analysing the results obtained from the *blur* condition of our study,

where we visualised the degree of affinity of a point to a reference object. We found that *blur* is not well suited for representing correctly quantifiable uncertainty to users in this type of visualisation. Instead, our findings suggest that the layer for visual presentation should aim to use a “rougher” visualisation that employs varying levels of transparency, which have a more clear-cut demarcation to their neighbouring areas, such as when reducing the resolution of a gradient by lowering the number of available colours. While not necessarily as accurate in representation, our results indicate that such a depiction is easier to interpret by humans. For the internal computer-based analysis, *blur*-based calculation and representation may still be used to satisfy the needs of both entities.

In a more end-user focussed approach, Hansen et al. [38] present an interactive software that supports the visualisation and steering of algorithms for students. Through this, the authors could demonstrate a positive impact on learning and understanding in comparison to traditional teaching methods. However, uncertainty is not visualised in the software. Here our findings could be implemented one-to-one to directly enhance students’ understanding of classification and prediction quality in K-means clustering and Linear Regression, illustrating the practical value of our empirically researched visual guide of visual means for representing various types of uncertainty.

Beyond this, our visual guide could support the representation of uncertainty of movement and trajectories in the work of Wang et al. [23]. Here the authors allow users to explore the impact of changes in certain variables on the position of points in a scatter plot using a vector line. This represents the potential positional trend of the points affected by the change. This vector line is the result of a Linear Regression that predicts the possible “decision flow”. Taking our work as an example, the authors could visualise the confidence intervals of the regressions using changes in *opacity* to represent possible deviations and uncertainties regarding the new variable’s impact, supporting the user in making a more accurate decision. Yet, care should be taken for visual clarity not to suffer.

7. Conclusion and future work

Depicting uncertainty in data and model representations in a visual analysis is important to increase user trust and interpretation accuracy [16]. Indicating to the user with what confidence the data points are classified or how far they are spread out from a model’s ideal distribution line, allows users to understand why and how a decision was made that led to a final result and its potential uncertainty. Further, users can take action and adjust algorithm parameters or remove problematic data points to iteratively improve the end result – a core feature of Visual Analytics applications [1].

While a range of frameworks and toolkits exists to aid laymen in visualising data and processes, [8–11], they often lack an appropriate visualisation of uncertainty, despite this being important for the sense-making process.

To work towards a guide for representing uncertainty in Visual Analytics applications that can be employed as an addition to existing work, we investigated the transferability of popular visual means of uncertainty depiction from the domain of geospatial visualisation to that of algorithm visualisation. We focussed on algorithms whose mechanics and portraiture – similar to elements in the geospatial visualisation domain – use a spatial distribution to indicate class or model affinity. Using a between-subjects study with 500 participants in ten conditions (five per algorithm) we found that *opacity* is the most suitable visual means, followed by *texture*. *Grid* and *blur* do not appear suitable to represent quantifiable uncertainty whereas the usability of *contour lines* appears to depend on the algorithm visualised and is better for clusters than regressions.

By focussing on laymen, our study has illustrated these means’ effectiveness and usability for communicating attribute uncertainty to non-data processing experts and non-visualisation experts, such as

scholars from the field of digital humanities or students. We thereby determined the first two building blocks of our visual guide: Classification and model uncertainty for K-means clustering and Linear Regression are best communicated using changes in *opacity* (Fig. 14). However, for this to be effective, clear points of visual reference have to be given. Simply showing a semi-transparent plane without another visual element to compare it to is unlikely to work well. Therefore, this means may only be effective in communicating uncertainty if it is used to highlight this in at least two adjacent zones in the classification space.

Uncertainty has to be represented in a way that corresponds to users’ expectations and knowledge [84] so that it is easily understood by laymen [15]. In this regard, the good performance of *opacity* may not only be founded in its clear demarcation of different zones of confidence, but also in its possible association with something hazy that lacks clear definition (in terms of its opaqueness). According to Yusoff and Salim [83], an ideal visualisation fosters a common understanding (social support), is perceived and interpreted the same way by all (cognitive support), and allows effective task-solving (task support). Based on this, the use of *opacity* for representing uncertainty in K-means clustering and Linear Regression visualisations appears to be a promising candidate: Users liked it as a means to represent uncertainty (social support), found it visually appealing (social support), and were able to infer uncertainty with high accuracy and low task completion times (cognitive and task-solving support).

In addition to the initial definition of elements for a visual guide to communicate uncertainty in Visual Analytics applications, our results may serve as evidence that findings from the domain of geospatial visualisation may not be confined to the environment they were derived from. Rather, the validity of these findings may be extendible to that of other areas of Information Visualisation, especially when depicting spatial relationships and attributes.

Future work will therefore investigate additional opportunities for transferring knowledge between the two domains and explore how well the examined visual variables may perform in the representation of uncertainty in visualisations of other “spatial” classification algorithms, such as Monte Carlo. Yet, to fill our visual guide with empirically validated representations of various types of uncertainty, we need to trawl multiple domains for other adequate visual means and rigorously test their transferability and effectiveness. Our work may thus only present a first step on a long road towards the definition of such a guide. Yet, as small as this step may be, it is a promising one, inciting us to tread faster and further. Due to the increasing extension of Visual Analytics and Information Visualisation to the “soft sciences”, the extensive research into the transferability of domain-specific knowledge from areas dominated by data-processing experts to those with a majority of data-processing laymen, appears to be a worthwhile endeavour.

8. Limitations

The study conducted in this article faces several limitations, the most obvious of which may be the use of AMT. As reported by previous work [86,87], recruiting users via this service poses challenges concerning the validity of responses. To address this, we applied several measures as discussed in Section 4, but results should still be interpreted with caution. Another aspect to take into account is the uncontrolled environment users completed the study in. Although Web-based studies are common [65], it could not be ensured that users gave the study their full attention or even consulted others. Yet, the reasonable numbers of participants per group ($N = 50$) may help lower the impact of this factor.

Based on Kinkeldey et al.’s [65] extensive review of uncertainty user studies and the potential challenges of glyphs and animations, we limited our primary exploration to a set of the five most popular visual variables for attribute uncertainty. To reduce complexity, we only explored the research questions with regards to the visualisation of K-

means clustering and Linear Regression. However, future work could include other algorithms or methods, whose mechanics also rely on spatial distributions.

Beyond this, it is unclear how well our findings would work in a multidimensional analysis. If this could be represented in a two-dimensional visualisation, then our findings are likely to be of use. However, in a three dimensional environment or where the problem space cannot be translated into two dimensions, further work is necessary to validate our results or to define additional rules regarding their use.

Finally, we only examined a single type of uncertainty (the attribution to a class or model), not multiple types, as is often the case in Visual Analytics applications [52,72]. Our study thus only conceptualises uncertainty as such, and does not include differentiation of specific types. Yet, this was done in line with the suggestions given by Beard and Mackaness [99], who propose three levels of uncertainty indicators: First, indicating that there is uncertainty, second, showing the type of uncertainty, and third, providing tools for the user to investigate the reason of this uncertainty. As our study focussed on probing the transferability of uncertainty representations from the domain of geospatial visualisation to that of Visual Analytics to work towards the definition of a visual guide for representing uncertainty in this domain, we focussed on step one suggested by Beard and Mackaness [99] with the extension of showing the degree of uncertainty. However, we only visualised five different levels of uncertainty, as Kinkeldey et al. [68] reported that users struggle reading this accurately for more than a maximum of six levels. Whether our findings scale to a more fine-grained differentiation is therefore unclear and will be the subject of future work.

Acknowledgements

Part of this work has been supported by the Research Foundation Flanders (FWO, grant agreement no. G0C9515N), the European Union (Horizon 2020 research and innovation programme, grant agreement no. 780751), and the Secretary of Higher Education, Science, Technology and Innovation, Ecuador (SENECYT, Project PIC-15-ESPOL-FWO-001).

References

- [1] D.A. Keim, F. Mansmann, J. Schneidewind, J. Thomas, H. Ziegler, *Visual Analytics: Scope and Challenges*, Springer, 2008.
- [2] T. Schreck, T. Tekušová, J. Kohlhammer, D. Fellner, Trajectory-based visual analysis of large financial time series data, *ACM SIGKDD Explor. Newslett.* 9 (2) (2007) 30–37.
- [3] G. Siemens, P. Long, Penetrating the fog: analytics in learning and education. *EDUCAUSE Rev.* 46 (5) (2011) 30.
- [4] J.S. Yi, Y. ah Kang, J.T. Stasko, J.A. Jacko, Toward a deeper understanding of the role of interaction in information visualization, *Visual. Comput. Graph. IEEE Trans.* 13 (6) (2007) 1224–1231.
- [5] D.A. Keim, J. Kohlhammer, G. Ellis, F. Mansmann, *Mastering the Information Age-solving Problems with Visual Analytics*, Florian Mansmann, 2010.
- [6] M. Jessop, Digital visualization as a scholarly activity, *Lit. Linguist. Comput.* 23 (3) (2008) 281–293.
- [7] M. John, S. Koch, F. Heimerl, A. Miller, T. Ertl, J. Kuhn, Interactive visual analysis of german poetics, *Digital Humanities 2015 Book of Abstracts*, (2015).
- [8] L. Ren, F. Tian, X.L. Zhang, L. Zhang, DaisyViz: a model-based user interface toolkit for interactive information visualization systems, *J. Vis. Lang. Comput.* 21 (4) (2010) 209–229, <https://doi.org/10.1016/j.jvlc.2010.05.003>. Part Special Issue on Graph Visualization
- [9] M. Elias, A. Bezerianos, Exploration views: understanding dashboard creation and customization for visualization novices, in: P. Campos, N. Graham, J. Jorge, N. Nunes, P. Palanque, M. Winckler (Eds.), *Human-Computer Interaction INTERACT 2011: 13th IFIP TC 13 International Conference, Lisbon, Portugal, September 5–9, 2011, Proceedings, Part IV*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2011, pp. 274–291.
- [10] B. Howe, A. Key, D. Perry, C. Aragon, VizDeck: a card game metaphor for fast visual data exploration, *CHI '12 Extended Abstracts on Human Factors in Computing Systems, CHI EA '12, ACM, New York, NY, USA, 2012*, pp. 1667–1672, <https://doi.org/10.1145/2212776.2223690>.
- [11] K. Wongsuphasawat, D. Moritz, A. Anand, J. Mackinlay, B. Howe, J. Heer, Voyager: exploratory analysis via faceted browsing of visualization recommendations, *Visual. Comput. Graph. IEEE Trans.* 22 (1) (2016) 649–658.
- [12] K. Potter, P. Rosen, C.R. Johnson, From quantification to visualization: a taxonomy of uncertainty visualization approaches, *IFIP Adv. Inf. Commun. Technol.* 377 (2012) 226–249, https://doi.org/10.1007/978-3-642-32677-6_15.
- [13] M. Kay, S.N. Patel, J.A. Kientz, How good is 85%?: A survey tool to connect classifier evaluation to acceptability of accuracy, *Proc. CHI'15, ACM, New York, 2015*, pp. 347–356, <https://doi.org/10.1145/2702123.2702603>.
- [14] T.A. Slocum, D.C. Cliburn, J.J. Feddema, J.R. Miller, Evaluating the usability of a tool for visualizing the uncertainty of the future global water balance, *Cartogr. Geogr. Inf. Sci.* 30 (4) (2003) 299–317.
- [15] B.C. Kwon, B. Fisher, J.S. Yi, Visual analytic roadblocks for novice investigators, *Visual Analytics Science and Technology (VAST), 2011 IEEE Conference on, IEEE, 2011*, pp. 3–11.
- [16] D. Sacha, H. Senaratne, B.C. Kwon, G. Ellis, D.A. Keim, The role of uncertainty, awareness, and trust in visual analytics, *IEEE Trans. Visual. Comput. Graph.* 22 (1) (2016) 240–249, <https://doi.org/10.1109/TVCG.2015.2467591>.
- [17] T.J. Davis, C.P. Keller, Modelling and visualizing multiple spatial uncertainties, *Comput. Geosci.* 23 (4) (1997) 397–408, [https://doi.org/10.1016/S0098-3004\(97\)00012-5](https://doi.org/10.1016/S0098-3004(97)00012-5).
- [18] S. Djurcilov, K. Kim, P. Lermusiaux, A. Pang, Visualizing scalar volumetric data with uncertainty, *Comput. Graph.* 26 (2) (2002) 239–248.
- [19] H. Ban, O. Ahlqvist, Representing and negotiating uncertain geospatial concepts: where are the exurban areas? *Comput. Environ. Urban Syst.* 33 (4) (2009) 233–246, <https://doi.org/10.1016/j.compenurbysys.2008.10.001>.
- [20] C.L. Borgman, The digital future is now: a call to action for the humanities, *Digit. Humanit.* 3 (4) (2009).
- [21] J.R. Bauer, K. Rose, Variable grid method: an intuitive approach for simultaneously quantifying and visualizing spatial data and uncertainty, *Trans. GIS* 19 (3) (2015) 377–397, <https://doi.org/10.1111/tgis.12158>.
- [22] K.A. Cook, J.J. Thomas, *Illuminating the Path: The Research and Development Agenda for Visual Analytics*, IEEE Computer Society, Los Alamitos, CA, United States(US)., 2005.
- [23] W.B. Wang, M.L. Huang, Q.V. Nguyen, W. Huang, K. Zhang, T.-H. Huang, Enabling decision trend analysis with interactive scatter plot matrices visualization, *J. Vis. Lang. Comput.* 33 (2016) 13–23, <https://doi.org/10.1016/j.jvlc.2015.11.002>. SI:VIMLA
- [24] I.A. Letia, A. Goron, Model checking as support for inspecting compliance to rules in flexible processes, *J. Vis. Lang. Comput.* 28 (2015) 100–121, <https://doi.org/10.1016/j.jvlc.2014.12.008>.
- [25] H. Van den Berg, G. Parra, A. Jentzsch, A. Drakos, E. Duval, Studying the history of philosophical ideas: supporting research discovery, navigation, and awareness, *Proceedings of the 14th International Conference on Knowledge Technologies and Data-driven Business, ACM, 2014*, p. 12.
- [26] M. Ghoniem, J.-D. Fekete, P. Castagliola, A comparison of the readability of graphs using node-link and matrix-based representations, *Information Visualization, 2004. INFOVIS 2004. IEEE Symposium on, IEEE, 2004*, pp. 17–24.
- [27] N.H. Riche, T. Dwyer, Untangling euler diagrams, *Visual. Comput. Graph. IEEE Trans.* 16 (6) (2010) 1090–1099.
- [28] J. Bernard, T. Ruppert, O. Goroll, T. May, J. Kohlhammer, Visual-interactive pre-processing of time series data, *SIGRAD* 81 (2012) 39–48.
- [29] S. Rudolph, A. Savikhin, D.S. Ebert, Finvis: Applied visual analytics for personal financial planning, *Visual Analytics Science and Technology, 2009. VAST 2009. IEEE Symposium on, IEEE, 2009*, pp. 195–202.
- [30] J. Thomas, J. Kielman, Challenges for visual analytics, *Inf. Visual.* 8 (4) (2009) 309–314.
- [31] S. Kandel, J. Heer, C. Plaisant, J. Kennedy, F. van Ham, N.H. Riche, C. Weaver, B. Lee, D. Brodbeck, P. Buono, Research directions in data wrangling: Visualizations and transformations for usable and credible data, *Inf. Visual.* 10 (4) (2011) 271–288.
- [32] S. Kandel, R. Parikh, A. Paepcke, J.M. Hellerstein, J. Heer, Profiler: integrated statistical analysis and visualization for data quality assessment, *Proceedings of the International Working Conference on Advanced Visual Interfaces, ACM, 2012*, pp. 547–554.
- [33] J. Kennedy, M. Graham, T. Paterson, A. Law, Visual cleaning of genotype data, *Biological Data Visualization (BioVis), 2013 IEEE Symposium on, IEEE, 2013*, pp. 105–112.
- [34] J. Sanyal, S. Zhang, G. Bhattacharya, P. Amburn, R. Moorhead, A user study to compare four uncertainty visualization methods for 1d and 2d datasets, *IEEE Trans. Visual. Comput. Graph.* 15 (6) (2009) 1209–1218.
- [35] M. Skeels, B. Lee, G. Smith, G.G. Robertson, Revealing uncertainty for information visualization, *Inf. Visual.* 9 (1) (2010) 70–81.
- [36] D. Hooshyar, R. Binti Ahmad, M. Fathi, M. Yousefi, M. Hooshyar, Flowchart-based bayesian intelligent tutoring system for computer programming, *Smart Sensors and Application (ICSSA), 2015 International Conference on, IEEE, 2015*, pp. 150–154.
- [37] T. Crews, U. Ziegler, The flowchart interpreter for introductory programming courses, *Frontiers in Education Conference, 1998. FIE'98. 28th Annual, 1 IEEE, 1998*, pp. 307–312.
- [38] S. Hansen, N. Narayanan, M. Hegarty, Designing educationally effective algorithm visualizations, *J. Vis. Lang. Comput.* 13 (3) (2002) 291–317, <https://doi.org/10.1006/jvlc.2002.0236>.
- [39] K. Marriott, B. Meyer, *Visual Language Theory*, Springer Science & Business Media, 2012.
- [40] G. Gleir, L. Bergmans, M. Akit, K. van den Berg, Vibes: a visual language for specifying behavioral requirements of algorithms, *J. Vis. Lang. Comput.* 24 (5) (2013) 350–364, <https://doi.org/10.1016/j.jvlc.2013.08.005>.
- [41] M. Erwig, E. Walkingshaw, A visual language for explaining probabilistic reasoning,

- J. Vis. Lang. Comput. 24 (2) (2013) 88–109, <https://doi.org/10.1016/j.jvlc.2013.01.001>.
- [42] R. Sinha, K. Swearingen, The role of transparency in recommender systems, CHI '02 Extended Abstracts on Human Factors in Computing Systems, CHI EA '02, ACM, New York, NY, USA, 2002, pp. 830–831, <https://doi.org/10.1145/506443.506619>.
- [43] G. Iusarczyk, Visual language and graph-based structures in conceptual design, Adv. Eng. Inform. 26 (2) (2012) 267–279, <https://doi.org/10.1016/j.aei.2011.10.005>. Knowledge based engineering to support complex product design
- [44] E.J. Marchiori, del Blanco, J. Torrente, I. Martinez-Ortiz, B. Fernandez-Manjin, A visual language for the creation of narrative educational games, J. Vis. Lang. Comput. 22 (6) (2011) 443–452, <https://doi.org/10.1016/j.jvlc.2011.09.001>.
- [45] E.J. Marchiori, J. Torrente, del Blanco, P. Moreno-Ger, P. Sancho, B. Fernandez-Manjin, A narrative metaphor to facilitate educational game authoring, Comput. Educ. 58 (1) (2012) 590–599, <https://doi.org/10.1016/j.compedu.2011.09.017>.
- [46] K. Zhang, Using visual languages in management, J. Vis. Lang. Comput. 23 (6) (2012) 340–343, <https://doi.org/10.1016/j.jvlc.2012.09.001>.
- [47] D. Howard, A.M. MacEachren, Interface design for geographic visualization: tools for representing reliability, Cartogr. Geogr. Inf. Syst. 23 (2) (1996) 59–77.
- [48] R. Allendes Osorio, K.W. Brodrie, Uncertain flow visualization using LIC, Theory and Practice of Computer Graphics, Eurographics UK Chapter Proceedings, (2009), pp. 215–222.
- [49] L. Bastin, P.F. Fisher, J. Wood, Visualizing uncertainty in multi-spectral remotely sensed imagery, Comput. Geosci. 28 (3) (2002) 337–350, [https://doi.org/10.1016/S0098-3004\(01\)00051-6](https://doi.org/10.1016/S0098-3004(01)00051-6).
- [50] T. Hengl, Visualisation of uncertainty using the HSI colour model: computations with colours, 7th International Conference on GeoComputation, (2003).
- [51] A. Pang, Risk Assessment, Modeling and Decision Support: Strategic Directions, Springer, Berlin, Heidelberg, 2008, pp. 261–294.
- [52] A. Slingsby, J. Dykes, J. Wood, Exploring uncertainty in geodemographics with interactive graphics, Visual. Comput. Graph. IEEE Trans. 17 (12) (2011) 2545–2554.
- [53] B. Berjawi, F. Duchateau, Chesneau, F. Favetta, G. Seccia, C. Cunty, M. Miquel, R. Laurini, Uncertainty visualization of multi-providers cartographic integration, Distrib. Multimedia Syst. DMS2014 Part I 25 (6) (2014) 995–1002, <https://doi.org/10.1016/j.jvlc.2014.10.033>.
- [54] A. Alesheikh, J. Blais, M. Chapman, H. Karimi, Rigorous geospatial data uncertainty models for GISs, Spatial Accuracy Assessment: Land Information Uncertainty in Natural Resources, Ann Arbor Press, Chelsea, Michigan, US, 1999, pp. 195–202.
- [55] G. Dutton, Handling positional uncertainty in spatial databases, Proceedings 5th International Symposium on Spatial Data Handling, (1992), pp. 460–469.
- [56] L.D. Edwards, E.S. Nelson, Visualizing data certainty: a case study using graduated circle maps, Cartogr. Perspect. (38) (2001) 19–36.
- [57] K. Brodrie, R.A. Osorio, A. Lopes, A review of uncertainty in data visualization, Expanding the Frontiers of Visual Analytics and Visualization, Springer, 2012, pp. 81–109.
- [58] M. Bloch, S. Carter, M. Ericson, Hurricane Tracker - NYTimes.com, 2011, Accessed: 2016-03-21.
- [59] Stamen Design, Storm Tracker, 2013.
- [60] D.J. Spiegelhalter, Funnel plots for comparing institutional performance, Stat. Med. 24 (8) (2005) 1185–1202.
- [61] P.A. Burrough, Natural objects with indeterminate boundaries, Geographic Objects with Indeterminate Boundaries, 2 (1996), pp. 3–28.
- [62] A.M. MacEachren, Visualizing uncertain information, Cartogr. Perspect. (13) (1992) 10–19.
- [63] N.D. Gershon, Visualization of fuzzy data using generalized animation, Proc. Visualization'92, IEEE, 1992, pp. 268–273, <https://doi.org/10.1109/VISUAL.1992.235199>.
- [64] B.P. Buttenfield, Mapping Ecological Uncertainty, in: C. Hunsaker, M.F. Goodchild, M. Friedl, T. Case (Eds.), Spatial Uncertainty in Ecology, Springer, New York, 2001, pp. 115–132.
- [65] C. Kinkeldy, A.M. MacEachren, J. Schiewe, How to assess visual communication of uncertainty? A systematic review of geospatial uncertainty visualisation user studies, Cartogr. J. 51 (4) (2014) 372–386, <https://doi.org/10.1179/1743277414Y.0000000099>.
- [66] E. Clementini, P. Di Felice, An algebraic model for spatial objects with indeterminate boundaries, Geographic Objects with Indeterminate Boundaries, 2 (1996), pp. 155–169.
- [67] A. Lucieer, M.-J. Kraak, Interactive and visual fuzzy classification of remotely sensed imagery for exploration of uncertainty, Int. J. Geogr. Inf. Sci. 18 (5) (2004) 491–512.
- [68] C. Kinkeldy, J. Mason, A. Klippel, J. Schiewe, Evaluation of noise annotation lines: using noise to represent thematic uncertainty in maps, Cartogr. Geogr. Inf. Sci. 41 (5) (2014) 430–439.
- [69] A. Cedilnik, P. Rheingans, Procedural annotation of uncertain information, Proc. Visualization 2000, (2000), pp. 77–84, <https://doi.org/10.1109/VISUAL.2000.885679>.
- [70] G. Hunter, M. Goodchild, A new model for handling vector data uncertainty in geographic information systems, URISA J. 8 (1) (1996) 51–57.
- [71] D.C. Cliburn, J.J. Feddema, J.R. Miller, T.A. Slocum, Design and evaluation of a decision support system in a water balance application, Comput. Graph. 26 (6) (2002) 931–949.
- [72] J. Sanyal, S. Zhang, J. Dyer, A. Mercer, P. Amburn, R.J. Moorhead, Noodles: a tool for visualization of numerical weather model ensemble uncertainty, Visual. Comput. Graph. IEEE Trans. 16 (6) (2010) 1421–1430.
- [73] T.S. Newman, W. Lee, On visualizing uncertainty in volumetric data: techniques and their evaluation, J. Vis. Lang. Comput. 15 (6) (2004) 463–491, <https://doi.org/10.1016/j.jvlc.2003.09.001>.
- [74] A. Pang, Visualizing uncertainty in geo-spatial data, Proc. of the Workshop on the Intersections between Geospatial Information and IT, National Research Council Arlington, VA, 2001, pp. 1–14.
- [75] P.F. Fisher, Visualizing uncertainty in soil maps by animation, Cartographica 30 (2–3) (1993) 20–27.
- [76] P. Fisher, Randomization and sound for the visualization of uncertain spatial information, Visualization in Geographic Information Systems, John Wiley & Sons, London, 1994, pp. 181–185.
- [77] A.M. MacEachren, A. Robinson, S. Hopper, S. Gardner, R. Murray, M. Gahegan, E. Hetzler, Visualizing geospatial information uncertainty: what we know and what we need to know, Cartogr. Geogr. Inf. Sci. 32 (3) (2005) 139–160, <https://doi.org/10.1559/1523040054738936>.
- [78] D. Huang, M. Tory, B.A. Aseniero, L. Bartram, S. Bateman, S. Carpendale, A. Tang, R. Woodbury, Personal visualization and personal visual analytics, IEEE Trans. Visual. Comput. Graph. 21 (3) (2015) 420–433.
- [79] M.F. Costabile, D. Fogli, P. Mussio, Visual interactive systems for end-user development: a model-based design methodology, IEEE Trans. Syst. Man Cybern. Part A: Syst. Hum. 37 (6) (2007) 1029–1046.
- [80] A.F. Blackwell, Pictorial representation and metaphor in visual language design, J. Vis. Lang. Comput. 12 (3) (2001) 223–252, <https://doi.org/10.1006/jvlc.2001.0207>.
- [81] A.F. Blackwell, The reification of metaphor as a design tool, ACM Trans. Comput.-Hum. Interact. 13 (4) (2006) 490–530, <https://doi.org/10.1145/1188816.1188820>.
- [82] A. Celentano, F. Pittarello, From real to metaphoric maps: cartography as a visual language for organizing and sharing knowledge, J. Vis. Lang. Comput. 23 (2) (2012) 63–77, <https://doi.org/10.1016/j.jvlc.2011.11.004>. Special issue dedicated to Prof. Piero Mussio
- [83] N.M. Yusoff, S.S. Salim, A systematic review of shared visualisation to achieve common ground, J. Vis. Lang. Comput. 28 (2015) 83–99, <https://doi.org/10.1016/j.jvlc.2014.12.003>.
- [84] J. Cybulski, S. Keller, D. Saundage, Metaphors in interactive visual analytics, Proceedings of the 7th International Symposium on Visual Information Communication and Interaction, VINCI '14, ACM, New York, NY, USA, 2014, pp. 212:212–212:215, <https://doi.org/10.1145/2636240.2636866>.
- [85] J.E. Grady, The Oxford Handbook of Cognitive Linguistics, Oxford University Press, 2007.
- [86] A. Kittur, E.H. Chi, B. Suh, Crowdsourcing user studies with mechanical turk, Proc. CHI'08, ACM, New York, 2008, pp. 453–456, <https://doi.org/10.1145/1357054.1357127>.
- [87] J. Wang, P.G. Ipeirotis, F. Provost, A Framework for Quality Assurance in Crowdsourcing, (2013).
- [88] H. Senaratne, L. Gerharz, E. Pebesma, A. Schwering, Usability of spatio-temporal uncertainty visualisation methods, Bridging the Geographic Information Sciences, Springer, 2012, pp. 3–23.
- [89] M. Leitner, B.P. Buttenfield, Cartographic guidelines on the visualization of attribute accuracy, Proceedings of AUTO-CARTO, 13 (1997).
- [90] N. Boukheila, A. Bezerianos, T. Isenberg, J.-D. Fekete, Evaluating sketchiness as a visual variable for the depiction of qualitative uncertainty, Visual. Comput. Graph. IEEE Trans. 18 (12) (2012) 2769–2778.
- [91] A. MacEachren, R. Roth, J. O'Brien, B. Li, D. Swingley, M. Gahegan, Visual semantics & uncertainty visualization: an empirical study, Visual. Comput. Graph. IEEE Trans. 18 (12) (2012) 2496–2505, <https://doi.org/10.1109/TVCG.2012.279>.
- [92] A. Endert, M.S. Hossain, N. Ramakrishnan, C. North, P. Fiaux, C. Andrews, The human is the loop: new directions for visual analytics, J. Intell. Inf. Syst. 43 (3) (2014) 411–435, <https://doi.org/10.1007/s10844-014-0304-9>.
- [93] J. Yang, F. Li, Z. Sun, S. Jiang, A small target detection method based on human visual system and confidence measurement, J. Inf. Hiding Multimedia Signal Process. 7 (2) (2016) 448–459.
- [94] G. Andrienko, N. Andrienko, P. Bak, D. Keim, S. Kisilevich, S. Wrobel, A conceptual framework and taxonomy of techniques for analyzing movement, J. Vis. Lang. Comput. 22 (3) (2011) 213–232, <https://doi.org/10.1016/j.jvlc.2011.02.003>. Special Issue on Visual Analytics and Visual Semantics
- [95] P. Bak, F. Mansmann, H. Janetzko, D. Keim, Spatiotemporal analysis of sensor logs using growth ring maps, IEEE Trans. Visual. Comput. Graph. 15 (6) (2009) 913–920.
- [96] E. Bertini, G. Santucci, Improving visual analytics environments through a methodological framework for automatic clutter reduction, J. Vis. Lang. Comput. 22 (3) (2011) 194–212, <https://doi.org/10.1016/j.jvlc.2011.02.002>. Special Issue on Visual Analytics and Visual Semantics
- [97] E. Clementini, A model for uncertain lines, J. Vis. Lang. Comput. 16 (4) (2005) 271–288, <https://doi.org/10.1016/j.jvlc.2004.11.003>. Perception and ontologies in visual, virtual and geographic space
- [98] Y. Gong, L. Wu, Y. Lin, Y. Liu, Probability issues in locality descriptions based on voronoi neighbor relationship, J. Vis. Lang. Comput. 23 (4) (2012) 213–222, <https://doi.org/10.1016/j.jvlc.2012.04.002>.
- [99] K. Beard, W. Mackaness, Visual access to data quality in geographic information systems, Cartographica 30 (2–3) (1993) 37–45.