

Beware of Validation by Eye: Visual Validation of Linear Trends in Scatterplots

Daniel Braun , Remco Chang , Michael Gleicher , and Tatiana von Landesberger 

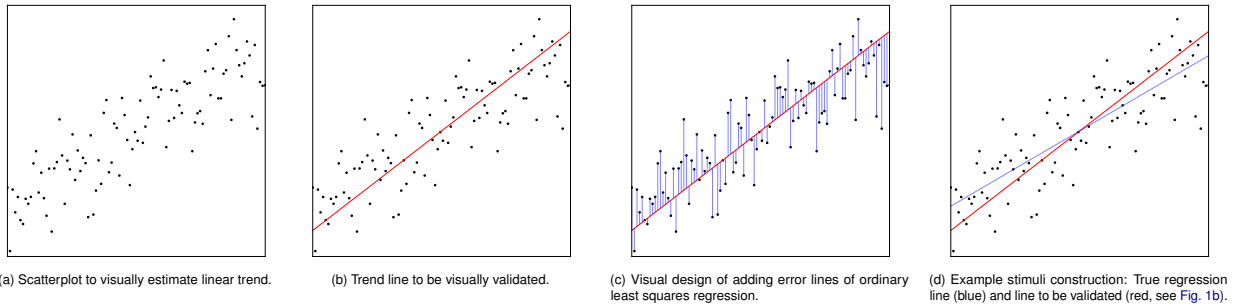


Fig. 1: Example scatterplot shown to participants in our user studies to investigate the visual validation of linear trends. In experiment 1 we tested the difference between visual validation (b) and visual estimation (a), while in experiment 2 (c) we compared different visual designs for regression to support visual validation. In (d), the blue line shows the true OLS regression.

Abstract—Visual validation of regression models in scatterplots is a common practice for assessing model quality, yet its efficacy remains unquantified. We conducted two empirical experiments to investigate individuals’ ability to visually validate linear regression models (linear trends) and to examine the impact of common visualization designs on validation quality. The first experiment showed that the level of accuracy for visual estimation of slope (i.e., fitting a line to data) is higher than for visual validation of slope (i.e., accepting a shown line). Notably, we found bias toward slopes that are “too steep” in both cases. This led to novel insights that participants naturally assessed regression with orthogonal distances between the points and the line (i.e., ODR regression) rather than the common vertical distances (OLS regression). In the second experiment, we investigated whether incorporating common designs for regression visualization (error lines, bounding boxes, and confidence intervals) would improve visual validation. Even though error lines reduced validation bias, results failed to show the desired improvements in accuracy for any design. Overall, our findings suggest caution in using visual model validation for linear trends in scatterplots.

Index Terms—Perception, visual model validation, visual model estimation, user study, information visualization.

1 INTRODUCTION

Visual validation of statistical models serves as an important task in modern statistics and machine learning applications. The complexity of models often demands visual inspection for assessing their correctness and reliability [7, 8]. This is crucial, as the model outcomes have great implications in critical domains [23], e.g., the estimation of pandemic outbreaks [14] and meteorological forecasting [24]. Without visualization, traditional statistical metrics are often insufficient in describing the underlying data and model. For example, datasets with significantly different characteristics can share the same numerical metrics [3, 28, 50]. As a result, visualization researchers have advocated for visual validation of statistical models as a core part of data analysis.

To date, most research on the perception of statistical models has focused on *visual estimation* – individuals’ ability to visually fit a model to data [11, 12, 20–22, 25, 31, 35, 36, 43, 44, 47, 48, 52–54]. While these studies contribute to our understanding of the visual estimation process, there is a lack of research on *visual model validation* – individuals’ ability to assess the fit of a given model to the underlying data. In a recent

study by Braun et al. [5], the authors found significant differences in individuals’ performance in visual validation and estimation of averages in scatterplots. In this paper, we build upon this work and investigate the accuracy and effectiveness of visual model validation for a more complex model – *linear trends* in scatterplots. While the average value offers a fundamental measure in one dimension, the exploration of linear trends provides valuable insights into the relationships between two data dimensions, which can be particularly relevant in machine learning applications where linearity assumptions are common. Scatterplots are the standard approach for showing relationships between two quantitative variables [33].

Linear trends can be interpreted in two different ways [2]: *Regression* and *correlation*. Regression estimates the underlying linear relationship between two variables to find a function that predicts the value of one variable based on the other. Correlation measures the strength of the linear relationship between two variables. The closer points lie to a straight line, the stronger their relationship. Due to the relevance of regression models for decision making in wide-ranging fields [16, 17, 32, 49], visual validation of regression models is important. In this paper, we investigate the perception of linear regression models and answer the following research questions:

- **RQ1:** How does performance in visual validation of linear trends relate to the accuracy of visual estimation?
- **RQ2:** Can common visual designs enhance the performance of visual validation of linear trends?

For RQ1, we examined individuals’ perception of the slope of a linear trend in a between-subject user study. Participants were randomly assigned to either validate the slope of a shown trend line in a scatterplot

- Daniel Braun, and Tatiana von Landesberger are with University of Cologne. E-mail: {braun | landesberger}@cs.uni-koeln.de.
- Remco Chang is with Tufts University. E-mail: remco.chang@tufts.edu.
- Michael Gleicher is with University of Wisconsin - Madison. E-mail: gleicher@cs.wisc.edu.

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxx

(Fig. 1b) or estimate the trend line on their own (Fig. 1a).

Our results confirm the previous findings of Braun et al. [5] that participants are more accurate in model estimation than in model validation. Further, we found that participants systematically overestimated the trends' slope and were particularly inaccurate at recognizing trend lines with a high slope value (i.e., lines that are "too steep"). This implies that the perception of the trend lines is biased – not symmetrical, i.e., lines with slopes that are too high and too low are perceived differently. A post-hoc analysis revealed that participants' responses were more consistent with the non-standard orthogonal regression (ODR: minimizes orthogonal distances between data points and regression line) than with the usual vertical regression (ordinary least squares - OLS: minimizes vertical distances between data points and regression line), since ODR regression has a higher slope per calculation (see Fig. 6). This implies they assumed errors in both x and y variables rather than just one. Our results show similar effects between positive and negative trends.

With respect to visual design, several recent studies have found significant effects of visualization on model estimation [21, 31, 47, 48, 50, 53]. We similarly investigate whether the addition of augmentations to the visual designs improve individuals' performance in the visual validation of models (RQ2). To answer this research question, we modified the visual designs and repeated our first study on the slope validation. Using the same data, we added three visual designs to the shown trend line - OLS error lines, 95% confidence intervals, and bounding boxes (see Fig. 12). The results showed that the addition of error lines reduced the bias in recognizing lines with slopes that are either too high or too low. The addition of bounding boxes slightly increased visual model validation accuracy. However, none of the designs lead to the desired improvements (i.e., higher acceptance rate for correct models and higher rejection rate for incorrect models). In fact, participants reported an increased task difficulty with the addition of visual designs with no benefits to task completion time.

Altogether, the results of our two studies find evidence to caution when using the common practice of "validation by eye" [23] – visually validating statistical models by overlaying a visualization of the model over raw data. Further research is needed on how to support individuals in the visual validation process.

The paper is structured as follows: After an overview of current research studies on visual validation and visual estimation in Sec. 2, we give details on the general data, study, and analysis design in Sec. 3. The two experiments - validation versus estimation and visual designs as well as their analysis and results are presented in Sec. 4 and Sec. 5. Sec. 6 discusses limitations of our work and possible future work.

2 RELATED WORK

Visual Model Validation The research topic of visual model validation has received little attention in the past, but has become more prominent recently as it is recognized as an essential part of exploratory data analysis [23]. Braun et al. [5] were the first to investigate the perceptual differences between visual validation and visual estimation using the example of average values in scatterplots. They found that participants were more accurate in estimation than in validation, that the visual validation of averages is unbiased, and that the critical point between accepting and rejecting a given value is close to the statistical 95% confidence interval. This study motivated this work on the validation of linear trends.

Few other studies relate to our experiment. Majumder et al. [27] examine the visual validation of statistical inference in linear models, in which participants had to visually find the most deviating model (e.g. highest slope) in a small multiple setting. Their study showed that visual tests have higher power than conventional tests when the effect size is large. The findings of Correll et al. [13] on visual validation of data distributions in scatterplots, histograms, and density plots suggest problems with overplotting, which informed the stimuli design in our study.

Most of the time, the possibility for visual model validation is given in an interactive way as part of a visual analytics or machine learning system. Chatzimparmpas et al. [7] give an overview of the current

use of visualization for machine learning model interpretation. Bögl et al. [4] provide a visual analytics process to assist domain experts in selecting suitable models in time-series analysis. While the Visual (dis)Confirmation tool by Choi et al. [10] allows users to perform data analysis by automatically generating appropriate visualizations based on hypotheses framed in natural language, Chegini et al. [9] support users in identifying local patterns in large scatterplot spaces by automatically comparing local regions using a model-based pattern descriptor. Mühlbacher and Piringer [34] introduce a partition-based framework for validating regression models both qualitatively via visualizations and quantitatively via a relevance measure for ranking features. Our work is synergistic with the prior work in that it contributes to a better understanding of visual model validation that can improve the development of visual analytics and machine learning systems in the future.

An additional component of model validation is the viewer's trust in these models and their visualizations, which has been widely studied in the past [15, 19, 29, 40, 45]. The survey by Chatzimparmpas et al. [8] summarizes the importance of this relationship between visual design and trust in machine learning. In our study, we tried to minimize the influence of trust by giving the participants as less information as possible without impairing their understanding of the tasks.

Visual Model Estimation Most studies on the perception of statistical models aimed to understand visual estimation. In addition to research on the estimation of average values [20, 22, 54], there are many papers on the estimation of linear trends.

Correll and Heer [12] examined the basic perceptual process of visual estimation and had participants perform "regression by eye" for linear trends in scatterplots and other visualizations. They find that an individual's ability to estimate the slope of a linear trend with respect to the least squares regression model depends on both visual features and data features, without bias for positive and negative trends. Ciccione and Dehaene [11] were also interested in the accuracy and bias of visual regression estimation in scatterplots. Their results indicate that people consistently overestimate trends. These works played a key role in our study, hypotheses, and stimuli design.

A different aspect of linear trends is the correlation of the two data dimensions. Rensink and Baldridge [44] investigated the influence of statistical properties on the perception of correlation in scatterplots. Besides statistical properties, correlation perception research dealt with the influence of visual designs, features, and ensemble coding [50]. Yang et al. [53] showed that visual features, such as bounding boxes, are used as proxies for estimating correlation in scatterplots. Xiong et al. [52] found that people estimate correlations more accurately in scatterplots with generic axis labels than with semantic labels. Comparisons of different correlation visualizations based on Weber's law ranked scatterplot as the best visualization design and showed that performances of correlation estimations differ for positive and negative correlations [21, 25]. These studies confirm our choice of scatterplots as the visual design for our studies, and their results have implications for our stimuli design.

3 EXPERIMENTAL DESIGN

We conducted two experiments to gain insights into the perceptual process of the visual validation of linear trends and to answer our research questions:

RQ1: How does performance in visual validation of linear trends relate to the accuracy of visual estimation? and

RQ2: Can common visual designs enhance the performance of visual validation of linear trends?

The **first experiment** aimed to answer RQ1. Therefore, it contained two *tasks* to compare the following perceptual processes:

- *Visual validation:* Participants were shown scatterplots with an already drawn trend line (see Fig. 1b). They were asked to indicate whether the line was "too steep", "too flat", or "about the same" (i.e., the shown line represents the actual trend) in relation to the true slope of the linear trend of the data.




Fig. 2: Example of the same amount of positive and negative slope deviation in the validation task.

- *Visual estimation*: Participants were asked to fit trend lines to the data in the scatterplots (see Fig. 1a) by adjusting the slope of the line by moving a slider.

The **second experiment** aimed to answer RQ2 by evaluating three regression visualization designs for validation. In this study, participants only had to perform the validation task described in experiment 1.

3.1 Study procedure

The two experiments addressed different research questions, but used the same study structure and data.

The *study procedure* began with demographic questions about the participants, followed by a short training period for the participants to familiarize themselves with the study interface. To avoid bias in the participants' responses, we did not provide training feedback. To minimize learning effects, the order of trials was randomized. In the study interface, each page displayed one trial (i.e., on plot). For further analysis, response times were recorded. At the end, we asked the participants for their strategy in performing the task and to rate the difficulty of the tasks on a 5-point Likert scale (1 – very difficult, 2 – difficult, 3 – neutral, 4 – easy, 5 – very easy) [51].

Using a *between-subject study design* (i.e. the participants were divided into two groups and had to either estimate or validate trends), we prevented learning effects between the trials and reduced the number of trials per participant [6]. To ensure consistency and comparability, the same data were used in both studies for all of the between-subject groups. The only difference between the two experiments was the visual representation of the data in the scatterplots (see Fig. 1b and Fig. 1c as an example).

3.2 Data Generation and Stimuli Design

Our data generation is inspired by the approaches used by Braun et al. [5] and Correll and Heer [12]. Each trial displayed a scatterplot with a size of 700×700 pixels. Therefore, we recommended a screen size of 13" or larger. The scatterplots contained 100 data points in the range of $[0, 1] \times [0, 1]$ uniformly distributed along the x-axis. We used the standard regression model for the point generation in order to be able to compare the validation and estimation results with the ordinary least squares (OLS) regression: $y = ax + b$. We used this function to generate set of points along particular trends. The trend lines were centered in the scatterplot (as by Correll and Heer [12]), i.e., $f(0.5) = 0.5$, and both positive and negative slopes were used. The y-coordinates of the resulting data points were then permuted using a normal distribution. Investigating the perception of the slope of linear regressions, the slope parameter a was the only variable in our studies. Centering the target trends in the plot ensured that participants could always estimate the true regression line by solely manipulating the slope value.

To keep the difficulty level as similar as possible between trials, the standard deviation of the normal distribution was fixed to 0.1 and the slopes of the trends were set to a range of $[0.35, 0.65]$. Since the permutation of the y-coordinates could lead to a deviation of the resulting regression from the original target trend, we performed rejection sampling to ensure that the slope of the trend of the resulting points was within 10^{-3} of the target slope.

To be able to measure the accuracy of visual validation, we showed the participants trend lines with slopes deviating from the true regression slope. Due to the centering of the underlying trends, a variation of the slope means a rotation of the line around the point $(0.5, 0.5)$. We define the **deviation** of the displayed line as the deviation from the true slope as a proportion of the regression's standard error:

$$\text{shown slope} = \text{true slope} + \text{deviation} \cdot \text{standard error} \quad (1)$$

Fig. 2 shows an example validation trial with the same amount of positive and negative slope deviation (0.5 in this example). The standard error of the regression is data dependent and represents the average distance between the points and the regression line. This allows us to compare and analyze trials with different data sets regardless of their properties, since an absolute change in slope is perceived differently in graphs with high and low point dispersion.

The deviation definition is based on the calculation of the regression confidence interval (CI). In statistical analysis, the 95% CI is a common measure for the uncertainty in models [1]. By construction – in our setting with a fixed number of points and distribution – the slope's 95% CI is set at a deviation of 0.198. We use the confidence interval to compare the user study results with statistical quality measures. In a statistical sense, all lines with a smaller slope deviation should be considered acceptable. However, people may have a smaller confidence interval and reject lines with deviations less than 0.198.

Based on the result of the study for validation and estimation of the average value [5], where deviations greater than 0.7 were consistently rejected, we used the same deviation range of $[-0.7, 0.7]$ (i.e., we showed lines with a maximum slope deviation of 0.7 based on Eq. (1)). The deviations we used for the lines shown in the studies were evenly distributed within this range in 0.05 increments (including deviation 0), resulting in a total of **58 trials** (i.e., 58 different data sets) per participant (29 trends with a positive slope, 29 with a negative slope).

In line with the literature [5, 12, 52], we kept the displayed scatterplots as clean as possible to minimize visual distraction by omitting axes marks and labels (see example stimuli in Fig. 1). The used colors and marker sizes were the same as in the study by Braun et al. [5].

3.3 Analysis Procedure

In order to assess our results in relation to literature, our analysis is similar to that of Braun et al. [5]. For comparability of the estimation and validation results, we transformed a participant's validation responses to binary results: 1 for *accepting* (i.e., "about the same") the shown line, 0 for *rejecting* (i.e., "too steep" or "too flat") the shown line. Similar to previous studies, logistic regression was then applied to the acceptance rates [5, 53] to assess validation accuracy.

The estimation errors (i.e., the deviation in slope of the self-adjusted trend lines) were calculated using the same deviation definition as for validation (Eq. (1)). This allows us to compare the logistic regression of validation acceptance rates with the cumulative distribution (CDF) of estimation errors.

For statistical testing, we used a multi-stage approach with the standard significant level $\alpha = 0.05$ in all tests. First, we performed a

Shapiro-Wilk test on the given responses and response times to test the data for normality, with the result that none of the data fulfilled this property. Based on this, we then used the non-parametric Kolmogorov-Smirnov (*KS*) test for the difference in validation and estimation results (i.e., comparing the logistic regressions of the validation acceptance rates and the CDF of the estimation errors). A Wilcoxon test was used to test the response times and estimation errors on difference in means. For comparing the Likert responses, we used chi-squared test. For experiment 2 we first performed a Kruskal-Wallis test on the response times and a chi-squared test on the categorical responses to test the visual designs for significant differences. As post-hoc pairwise analysis of the respective tests to compare the different design combinations.

4 EXPERIMENT 1: VISUAL VALIDATION VERSUS VISUAL ESTIMATION

Experiment 1 analyzed the performance in visual validation of linear trends in relation to visual estimation. Based on previous research on model perception [5, 11, 12, 41] and our own assessments when generating the data, we propose the following hypotheses for experiment 1:

- **H1:** *The accuracy of visual validation is lower than the accuracy of visual estimation when perceiving the slope of a linear trend in a scatterplot.*
The study by Braun et al. [5] for the average value showed that visual estimation provides a higher level of accuracy compared to visual validation. We expect this result to hold for linear trends.
- **H2:** *People’s critical point between accepting and rejecting a given trend line when validating is close to the boundary of the 95% CI.*
Also based on the results of Braun et al. [5] we expect people’s “visual confidence” to match the statistical CI, as is true for the visual validation of average values.
- **H3:** *For visual validation, the results don’t differ between positive and negative trends.*
Correll and Heer [12] found no bias in visual estimation of linear trends. We expect visual validation and estimation to be similar in this regard.
- **H4:** *For visual estimation, people overestimate the slope of linear trends.*
We expect this behavior found in previous research [11, 12, 41] to be the same in our study.
- **H5:** *Perceived task difficulty and task completion time are lower for visual validation than for visual estimation.*
The validation task is simply a matter of acceptance, whereas the estimation task requires participants to fit a line to the data. Therefore, we expect participants to perceive the validation task as easier and complete it faster.

4.1 Experimental Setting and Participants

We conducted an online study on Limesurvey [26] and recruited participants from the crowdsourcing platform Prolific [42]. A total of 122 participants took part in the study. They had to speak English fluently. No restrictions were made for country of residence. We removed the data from 12 participants because of their incorrect answers to the attention question. Out of the 110 remaining participants, 46 completed the validation and 64 answered the estimation task. Most of them were between 20 and 40 years old (86%) and the gender distribution was close to even (F: 48%, M: 49%, other: 3%). Participants’ educational levels ranged from high school diplomas to doctorate degrees, with the majority having a bachelor’s degree (40%). The overall self-reported expertise in statistical model estimation was relatively low, with 88% of the participants indicating an expertise between 1 and 3 on a 5-point Likert scale. The average time to complete the study was 20 minutes. Participants were compensated with £4.45.

4.2 Results

Given our data generation approach based on the ordinary least squares (OLS) regression (see Sec. 3.2), we analyze the responses of the study participants based on this model.




Fig. 3: Comparison of validation and estimation accuracy (absolute deviation) with respect to OLS regression. Blue line: Cumulative distribution (CDF) for the estimation errors. Orange line: Logistic regression for the validation acceptance. Green line: Statistical 95% CI. Notice that more statistically valid lines were estimated than accepted by validation and more invalid lines were accepted than estimated.

4.2.1 Accuracy of Visual Validation vs. Visual Estimation

We evaluate the accuracy of visual validation and estimation in relation to the “statistical accuracy” of OLS regression. In statistical terms, all trend lines within the 95% confidence interval are considered to be valid. This means that, in a perfect world, participants should accept all trend lines with a slope deviation of less than 0.198 (i.e., the CI slope deviation) and reject all trend lines with a higher slope deviation in the validation task. In the estimation task, participants should ideally only estimate lines with slope deviations less than 0.198.

To be able to compare validation and estimation results, we summarize the acceptance rates and estimation errors. We combine the results for positive and negative trends and positive and negative slope deviations as absolute acceptance rates and errors. Fig. 3 shows the resulting logistic regression for the validation acceptance rates and the cumulative distribution for the estimation errors.

As shown, when the slope deviation is low (e.g., less than the 0.198 – the green line), participants were more accurate in estimating than validating trend lines. For example, about 70% of the participants were able to estimate (i.e., drew) trend lines with slope values below the deviation of 0.198. In contrast, only 60% of the participants correctly validated (i.e., accepted) shown trend lines with the same slope values. When the slope deviation is high, again we see that participants were more accurate in estimating than validating trend lines. For example, only about 20% of the participants estimated (i.e., drew) trend lines with slope values above the deviation of 0.6. In contrast, about 35% of the participants falsely validated (i.e., accepted) shown trend lines with the same slope values. Only in the deviation range adjacent to the confidence interval (0.2 to 0.3) validation is more accurate than estimation. In this statistically invalid range, a slightly higher percentage of lines were estimated than accepted. Notably, the validation acceptance of linear trends has an almost linear relationship with the slope deviation. Ideally, it should have high values at low slope deviations with sharp drop at the 95% CI border (i.e., 0.198).

The KS-test showed the two curves to be significantly different ($p \ll 0.01$, $D = 0.296$). Therefore, **H1** is supported for OLS regression.

The critical points of validation and estimation are precisely these slope deviations, with a 50/50 chance that a line will be accepted or rejected, or estimated with a lower or higher slope. This critical point is slightly lower for estimation ($crit_{val} \approx 0.36 > crit_{est} \approx 0.32$). Moreover, the critical point of validation is much greater than the deviation of the 95% CI (0.198), which does not support our hypothesis **H2** for the OLS regression. The critical points describe an experimental human threshold for the two tasks. People estimated lines that would be correct with the statistical 99.8% CI. For visual validation, they even accepted lines with a deviation greater than the statistical 99.9% CI.

When analyzing the individual performance of the participants in visual validation, it is noticeable that 72% of the participants had an individual critical point greater than the 95% CI and 13% accepted incorrect models more often than correct ones.

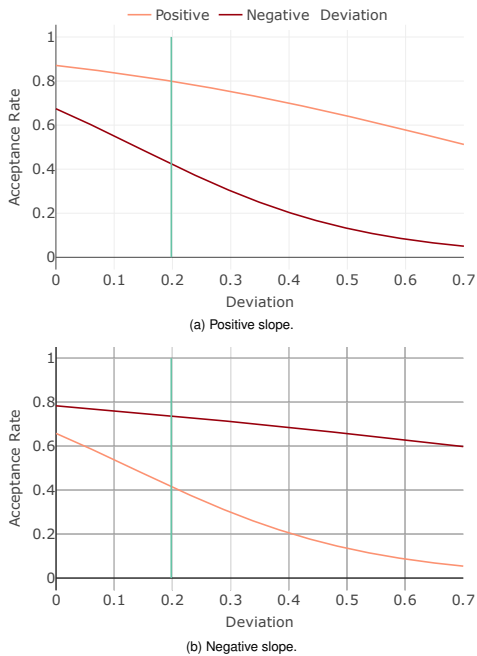


Fig. 4: Comparison of the *validation* acceptance rates of positive and negative deviations for positive and negative trends with respect to OLS regression. Green line: statistical 95% CI.

4.2.2 Bias in Positive and Negative Slope and Deviation

Fig. 4 differentiates the validation acceptance rates by positive and negative trends as well as positive and negative slope deviation. For positive trends, lines with positive slope deviations (i.e., trend lines that were “too steep”) were accepted significantly more often than lines with negative deviations (see Fig. 4a) (KS: $p \ll 0.01$, $D = 0.814$). For negative trends, the results are mirrored, i.e., lines with a negative slope deviation were accepted significantly more often (Fig. 4b) (KS: $p \ll 0.01$, $D = 0.926$). For both trend directions, lines that were “too steep” were still accepted more than 50% of the time, even with the largest slope deviation. Analysis of individual participant acceptance rates showed that the difference between positive and negative slope deviations was significant for each individual participant. Comparing the logistic regressions of the acceptance rates between positive and negative trends, there is no significant difference for “too flat” lines (KS: $p > 0.99$, $D = 0.022$), but a significant difference for “too steep” lines (KS: $p \ll 0.01$, $D = 0.336$), indicating a slightly more accurate validation of positive trends. Thus, hypothesis **H3** cannot be rejected.

A similar pattern can be observed for the estimation errors (see Fig. 5), supporting the results of previous studies [11, 12, 41] and our hypothesis **H4**: People overestimate the slope of trend lines. Without over-estimation, both distributions should be centered at 0 slope deviation. With an average slope deviation of 0.357 for positive trends and -0.359 for negative trends, participants consistently drew the trend lines too steeply (Wilcoxon: $p_{\text{positive}} \ll 0.01$, $V_{\text{positive}} = 1660685$; $p_{\text{negative}} \ll 0.01$, $V_{\text{negative}} = 78910$). No significant differences in estimation errors between positive and negative trends could be found (Wilcoxon: $p > 0.52$, $W = 1743165$).

In summary, both perceptual processes – visual validation as well as visual estimation – are biased toward “too steep” slopes for positive as well as negative trends.

4.2.3 Post-Hoc Analysis: OLS vs. ODR

The results opened a new question: What is the reason why people perform poorly at perceiving trend lines that are “too steep”?

A clue to this question was found in a study by Ciccione and Dehaene [11]. Inspired by their results, we hypothesize that individuals perceive orthogonal distance (ODR) instead of orthogonal least squares




Fig. 5: Histogram of the deviations of the *estimated* lines for positive and negative trends with respect to OLS regression.




Fig. 6: Illustration of the difference between ordinary least squares (OLS, blue line) and orthogonal distance (ODR, yellow line) regression, adapted from [11].

(OLS) regression, since ODR regression has a higher slope per calculation. Fig. 6 illustrates the difference between the two regression models. OLS regression minimizes the sum of the squares of the vertical distances of the points to the line and assumes noise only in the dependent variable (y-axis). ODR regression minimizes the sum of the squares of the orthogonal distances of the points to the line and assumes noise in both variables. The slopes of the ODR regression lines differ from those of the OLS regression. Thus, taking the ODR regression model into account changes the true slope value, which influences the measured accuracy of visual validation and estimation. We analyzed and compared the results of both tasks with respect to both regression models – OLS and ODR – as a post-hoc analysis to see if it explained the bias in slope perception.

Visual Model Validation: The acceptance rates for the shown trend lines in the validation task compared for ODR and OLS are visualized in Fig. 7. They differ significantly for both models (KS: $p \ll 0.01$, $D = 0.342$). For ODR regression, more lines with small slope deviation were accepted and more lines with large deviation were rejected. The




Fig. 7: Comparison of *validation* accuracy (absolute deviation) with respect to OLS and ODR regression.




Fig. 8: Histogram of the deviations of the *estimated* lines for positive and negative trends with respect to ODR regression.




Fig. 9: Comparison of validation and estimation accuracy (absolute deviation) with respect to ODR regression. Blue line: Cumulative distribution (CDF) for the estimation errors. Orange line: Logistic regression for the validation acceptance. Green line: Statistical 95% CI.

logistic regressions intersect closely at the 50% acceptance rate, i.e. the critical point between acceptance and rejection of a line remained almost identical for both regression models.

Visual Model Estimation: There is also evidence of an improvement in visual estimation (see Fig. 8). The estimation errors have improved for both positive and negative trends and the slope is less strongly overestimated compared to OLS ($\mu_{pos} = 0.210$, $\mu_{neg} = -0.216$).

Comparing the two tasks, the differences between estimation and validation were greater for ODR than for OLS (see Fig. 9). The critical point of estimation moved closer to the border of the 95% CI for ODR ($crit_{est} = 0.208$). As a result, the estimation errors were significantly lower than the acceptance threshold for validation (KS: $p \ll 0.01$, $D = 0.234$).

Fig. 10 illustrates the accuracy results of experiment 1. It visualizes an example stimuli with the true OLS and ODR regression lines together with the average responses for validation and estimation. For estimation, this means a line (shown in blue) with the average deviation of the estimated lines. For validation, this is represented as the range of lines (shown in orange) that would be accepted at least 50% of the time.

In sum, we found that participants’ trend perception was more consistent with the ODR regression model than with the OLS model for both estimation and validation tasks.

4.2.4 Response Time and Difficulty

Fig. 11 shows the response times and the distribution of the Likert scale responses for the task’s difficulty. As shown, the response times for visual validation were significantly faster than for visual estimation (Wilcoxon: $p \ll 0.01$, cohensD = 0.12, $\mu_{val} = 10.85sec$, $\mu_{est} = 13.73sec$). In contrast, participants’ self-reported task difficulty was significantly lower for the estimation than for the validation task (Chi-squared: $p < 0.05$, $\chi^2 = 10.08$; $\mu_{val} = 3.28$, $\mu_{est} = 3.84$). Therefore, Hypothesis **H5** cannot be rejected.




Fig. 10: “Visual summary” of the results for experiment 1 in an example stimulus. The figures show the true regression line (green) for OLS (a) and ODR (b) together with participants’ average response for estimation (blue) and the range of lines with an acceptance rate of 50% or higher for validation (orange). Notice that in figure (b), the blue and the green lines are closer to each other than in figure (a) and the orange range better encapsulate the green line. These suggest that the ODR model better fits participants’ perception of trend lines.




Fig. 11: Response times and self-reported difficulty for experiment 1.

4.2.5 Self-Reported Strategies

Participants’ self-reported strategies were provided in free-text form and subsequently summarized by us. For *validation*, the most often strategy was the comparison of the shown line shown with a self-estimated line (responded by 8 participants). Alternative strategies include a counting strategy ($n = 7$) where the participants counted the dots on both sides of the line, and a strategy that references the overall visual image of the visualization ($n = 6$) where the participants checked whether the trend line passed through the center of the area of dots. The latter strategy is similar to the strategy of a “bounding box” around the scatter dots, which was investigated by Yang et al. [53] for the perception of correlation (see Sec. 5.1).

For *estimation*, one strategy was used very often. 14 participants reported using the counting strategy and balancing the number of points on each side of the line. Only two other strategies were mentioned more than once: drawing a line through the middle of the dots ($n = 3$) and mentally connecting the dots to a line chart ($n = 2$).

For both tasks, the comparison of the results per strategy did not show any significant differences.

4.3 Discussion

Comparing the perceptual tasks of visual validation and estimation, individuals are more accurate at estimating regression models themselves than validating existing ones. One possible reason for this is that, as noted in the self-reported strategies, when validating a trend line, many people compare the line shown to a self-estimated line. Due to their uncertainty in their own estimation, they then may give themselves a margin of tolerance, which is reflected in the larger accepted deviations for the validated lines. This may indicate that people go through a two-step process during visual validation, with estimation as the first step. However, the response times and the self-reported task difficulty provide contradictory supporting evidence. Although participants reported estimation to be easier than validation, they also needed more time to




Fig. 12: Visual designs evaluated for visual validation in experiment 2.

complete the task. One possible reason for the longer response time for estimation is the interactivity associated with the task. The interaction time might be reduced by providing the participants a draggable line instead of a slider. In future work, it would also be interesting to capture the number of readjustments of the estimated lines to compare it with the response time and the accuracy results.

In both estimation and validation, the critical points are well above the 95% confidence interval, meaning that people are not able to perceive regression models with an acceptable level of accuracy by statistical standards. Therefore, individuals cannot rely on their visual validation ability. Given the relatively low self-reported expertise of the participants, the results might be different for people familiar with regression concepts, and an additional experiment solely with domain experts would be interesting for future work.

Our study design allows us to compare our results for the perception of linear trends with those of Braun et al. for the perception of the average value [5]. In general, both visual validation and estimation are less accurate for linear trends than for average values. This is probably due to the increased difficulty of the task, as the average value only examines the characteristic of one variable, while the linear trend examines the relationship between two variables. Nevertheless, the relation between validation and estimation remains similar, but only with respect to ODR regression. The reason for this is most likely that the vertical regression is both statistically and visually decisive when calculating and perceiving the average value.

The results of the experiment showed that the participants in our study were biased towards trend lines that were “too steep”. In the estimation task, they consistently drew lines with slopes that are too high, while in the validation task they accepted significantly more lines that were too steep than lines that were too flat with respect to the true trend line. Possible reasons, such as confirmation bias (i.e., given a trend line, participants may identify points that support the trend line as a valid one) [18] or the overestimation of large values and underestimation of small values [30] could not be verified by our study data and therefore require further investigation. However, the bias is reduced if the results are considered in relation to the ODR instead of the OLS regression. This indicates that, without any context or other assistance, people intuitively estimate and validate an orthogonal regression (ODR) instead of a vertical one (OLS) in scatterplots. That means, people naturally perceive errors in both variables even though they are not present in the data. As OLS is the more commonly used model for regression, we conducted a second experiment that uses additional visual augmentations to the scatterplots to improve the validation quality for OLS regression (see next section).

5 EXPERIMENT 2: VISUAL VALIDATION WITH VISUAL DESIGNS

The results from experiment 1 showed that visual validation is less accurate than visual estimation. Moreover, people more likely assess orthogonal distances between data points and trend line (i.e. ODR) instead of vertical distances (i.e., OLS). However, since OLS is the more commonly-used regression model, in this section, we investigate whether the addition of visualization designs may help improve the participants’ accuracy when performing visual validation with OLS

regression.

Experiment 2 evaluated three common visual augmentations for regression visualization (error lines, bounding boxes, and confidence intervals), as shown in Fig. 12. As described later in Sec. 5.1, we expect error lines to have the most influence on the results. Thus, our hypotheses are:

- **H1:** *Visual designs improve the validation of OLS.* It means that people accept more statistically valid and reject more invalid trend lines with respect to OLS regression.
- **H2:** *The visual design using error lines removes bias towards higher slopes.* This means, there is no perceptual difference between positive and negative slope deviations with respect to OLS regression anymore.
- **H3:** *Error lines reduce the time and difficulty of the task with respect to the other designs including the unaugmented chart.*

5.1 Visual Designs

Visual designs for regression validation mean additional graphical elements that are shown to the user in addition to the trend line in the visualization. We consider common visual designs for showing regression results found in literature (e.g., [37,39,53]) and the strategies used by participants in experiment 1 (see Sec. 4.2.5).

Error Lines The first visual design employs error lines [53]. Error lines show the the vertical distance of each data point from the displayed line, emphasizing the error minimized by OLS regression (see Fig. 12a). This potentially relieves the user of one step of the regression calculation, which should reduce participants’ response times. Moreover, by visually guiding towards OLS regression, the explicit error lines should make people perceive less the ODR regression. Therefore, it should also reduce the bias toward lines that are too steep.

Bounding Box The concept of bounding box covering the data was mentioned by several participants in the self-reported strategies and is inspired by Yang et al. [53]. Our implementation constructs the box by moving two lines parallel to the shown line outward until they reach the outermost points. The surface is then colored with an alpha value of 0.1 for a lower opacity (see Fig. 12b). The resulting box highlights the slope of the line shown. This enlarged area, which includes all data points, is intended to help participants compare the slope of the line shown with the true trend of the data as a whole.

Confidence Interval Visualizing the confidence interval of a regression model is common in several different areas and applications, such as pandemic infection projections or weather forecasts [38,46,55]. It shows the uncertainty in the underlying model [37,39]. All models within the confidence interval should be considered valid in a statistical sense. For our stimuli, we rotate the 95% confidence interval of the true OLS regression model in the same way as the shown line (see Fig. 12c). The goal is to enhance the perception of slope deviation of a line by showing all statistically valid models that accordingly have an even greater deviation.




Fig. 13: Comparison of validation accuracy (absolute deviation) with respect to OLS regression for the visual designs. Green line: Statistical 95% CI.

5.2 Experimental Setting and Participants

We conducted a second user study on Limesurvey [26] (with participants from Prolific [42]) using the identical data and study procedure as in experiment 1. The between-subject groups were now defined by the three designs instead of the two tasks in the first study. All participants performed the visual validation task.

After filtering for attention checks, a total of 108 participants were included in the analysis (32 for *error lines*, 38 for *bounding box*, and 38 for *confidence interval*). The demographic characteristics of the participants were similar to experiment 1: 87% were between 20 and 40 years old with 47% females and 52% males. The education levels were also similar: 43% of the participants had a bachelor's degree and 82% responded with an expertise between 1 and 3 on the Likert scale.

Participants were compensated with £3.54. The average study completion time was 17 minutes.

5.3 Results

We can directly compare the results for the visual designs with the *base line* (i.e., lines without additional augmentations) from experiment 1.

Accuracy With respect to OLS regression (see Fig. 13), the pairwise KS-test indicated significant differences in the acceptance rates between all designs (Tab. 1).

For the *bounding box* and the *confidence interval*, a similar number of valid lines (i.e., trend lines within the CI) were accepted compared to the *base line* condition, while a slightly higher amount of invalid lines were rejected. For the *error lines*, fewer lines were accepted that were within the statistical 95% confidence interval. The true trend lines were accepted only 60% of the time.

The overall decrease in the acceptance rates with visual designs have lowered the critical point between acceptance and rejection of a shown line for all three designs ($crit_{error} = 0.190$, $crit_{box} = 0.251$, $crit_{conf} = 0.298$).

In sum, people rejected slightly more invalid models with visual designs, but did not accept more valid models. Therefore, our hypothesis **H1** cannot be rejected.

Bias Fig. 14 shows the differentiated acceptance rates for the designs. For the *error lines*, the deviation difference between the acceptance' critical points of positive and negative slope deviation decreased ($\Delta_{error}^{pos} = 0.020$ for positive trends and $\Delta_{error}^{neg} = 0.117$ for negative trends). While the acceptance rates for negative trends still differ (KS: $p \ll 0.01$, $D = 0.182$), they are not significantly different for positive

p-value	base	error	box	conf
base	-			
error	$\ll 0.01$	-		
box	$\ll 0.01$	$\ll 0.01$	-	
conf	$\ll 0.01$	$\ll 0.01$	< 0.01	-

Table 1: p-values of the pairwise KS-test for the analysis of the acceptance rates of OLS regression.




Fig. 14: Comparison of the validation acceptance rates for the different visual designs for positive and negative deviations and for positive and negative trends with respect to OLS regression. Green line: statistical 95% CI.

trends anymore (KS: $p > 0.71$, $D = 0.044$). This result suggests that the perception of OLS regression was improved for the participants with *error lines*. It partially supports hypothesis **H2**.

Similar to the *base line*, the *confidence interval* biased the perception of positive and negative slope deviation ($\Delta_{conf}^{pos} = 0.348$, $\Delta_{conf}^{neg} = 0.353$). The *bounding box* showed a slight improvement in the perception of lines that were too steep ($\Delta_{box}^{pos} = 0.186$, $\Delta_{box}^{neg} = 0.159$).

Response Time and Difficulty The response times (measured in seconds) and participants' self-reported task difficulty on a 5-point Likert scale are shown in Fig. 15.

The Kruskal-Wallis test indicated a significant difference in the response times ($p \ll 0.01$, $\chi^2 = 112.42$). A pairwise Wilcoxon test of the visual designs (Tab. 2a) showed that the response times with the *error lines* were significantly higher than for all other designs ($\mu_{error} = 12.64sec$). Among the other designs, the *confidence interval* had the fastest response times ($\mu_{conf} = 9.21sec$), followed by the *bounding box* ($\mu_{box} = 9.96sec$) and the *base line* ($\mu_{base} = 10.85sec$).

A significant difference in the designs was also found in the perceived difficulty (Chi-squared: $p < 0.05$, $\chi^2 = 21.19$). The task was perceived to be the easiest with the *base line* ($\mu_{base} = 3.28$). The only significance could be found when comparing the *base line* to the *error lines* (Tab. 2b) where the *error lines* was perceived as the most difficult ($\mu_{error} = 2.69$). In between the *base line* and the *error lines* were the *bounding box* ($\mu_{box} = 3.21$) and the *confidence interval* ($\mu_{conf} = 2.92$).

Summarizing the results, hypothesis **H3** is not supported.

5.4 Discussion

The results of experiment 2 show that people consistently reject trend lines more often when they are presented with visual designs regardless of whether the trend lines are valid or invalid. There could be several different reasons for this: 1) The visual designs help to identify incorrect regression models. 2) People are generally more skeptical when shown model results with visual designs because they are not used to do




Fig. 15: Response times and self-reported task difficulty for the different visual designs in experiment 2.

p-value	base	error	box	conf
base	-			
error	$\ll 0.01$	-		
box	1.00	$\ll 0.01$	-	
conf	0.92	$\ll 0.01$	0.58	-

(a) Response times.

p-value	base	error	box	conf
base	-			
error	< 0.01	-		
box	0.38	0.06	-	
conf	0.11	0.36	0.18	-

(b) Self-reported difficulties.

Table 2: p-values of the pairwise Wilcoxon-test for the response time analysis and pairwise chi-squared test for analysis of the self-reported difficulties.

not understand this type of presentation. 3) People perceive the type of model to not fit the underlying data. In our study, we intentionally did not provide an explanation for the designs to allow for a purely perceptual study. It is therefore possible that a prior explanation of the visual designs would improve participants’ understanding and thus the results. Similarly, the results may be influenced by the low statistical expertise of our participants.

With respect to OLS regression, the acceptance thresholds improved with all designs compared with the *base line*. This is due to slightly higher rejection rates of invalid trend lines. The validation accuracy of valid lines did not improve with any design.

As with the *base line*, there is no bias in trend direction with the visual designs. Although confidence intervals are commonly used to represent statistical uncertainty, the bias in slope deviation is greatest with the addition of the confidence interval in the visualization. The error lines, on the other hand, provided an unbiased validation in terms of “too steep” and “too flat” trend lines.

The addition of *error lines* in the visualization should in theory reduce the cognitive effort of visual validation, because the cognitive calculation of errors is not needed. However, participants took longer to complete the validation task and found it more difficult than with the *base line*, according to the results of our study. This suggests that people either did not fully understand the concept of error lines without

explanation, that the processing of additional information is cognitively demanding, or that it forces people to intensify their thinking about the shown line and to correct their bias.

The results of experiment 2 showed that the addition of commonly used visual designs for visualizing regression in a scatterplot does not significantly improve people’s ability to validate models visually. Therefore, we are unable to provide design guidelines. As visual estimation remains more accurate, this suggests that guiding people to cognitive estimation as a first step of visual validation might improve accuracy.

6 LIMITATIONS AND FUTURE WORK

We investigated the visual estimation and validation of linear trend lines in scatterplots. The findings and limitations in our experiments may suggest new research questions and future directions.

Model complexity: We found that individuals’ ability to visually validate a linear model is lower than their ability to validate a constant model (i.e., averages [5]). This raises the question whether the ability to visually validate is dependent on the complexity of a model. In order to answer this question, however, it would be necessary to define model complexity in relation to visual perception.

Data characteristics - Outliers: Our study analyzed data with normal distribution as assumed by OLS regression models. However, real world datasets may have special characteristics that impact the regression. For example, Correll and Heer [12] studied the influence of noise and outlier in the data. The addition of outliers would shift the validation question of a correct or incorrect model to a question of including or excluding the outliers in the regression. This question about the correctness of the model itself could be extended to a study on the validation of model types. In this, people would have to be decided whether the type of model fits the data or not. In this paper, we examined the fit of the parameters of a fixed model.

Visual designs for model validation: We tested four designs (unmodified base line and three visual design augmentations) that are commonly used in visualizing regression models. The addition of the three visual designs in experiment 2 failed to improve both the accuracy and bias of visual validation. It would be interesting to see whether the results improve in a separate study where explanation and context to the data and the visual designs are made available to the participant. Additional think-aloud sessions could further provide insights into people’s visual validation process and the use of visual designs. Altogether, the development of novel visual designs for model validation that improves both accuracy and mitigates bias is a future challenge.

7 CONCLUSION

In summary, our research examines the effectiveness of visual validation in assessing linear regression models shown in scatterplots. We conducted two empirical experiments to gain insight into individuals’ abilities to visually validate linear trends and the impact of common visualization designs on validation quality.

The first experiment showed that participants were more accurate at visually estimating linear trends in scatterplots than at visually validating them. In addition, our results revealed a bias in slope deviation (i.e., toward slopes that are “too steep”), but no bias in trend direction. Additional analysis provides evidence that people naturally assess orthogonal regression (ODR) rather than the most commonly used vertical regression (OLS). This indicates that people assume errors in both variables rather than in just the y-coordinate.

The second experiment aimed to evaluate whether incorporating common visualization designs such as error lines, bounding boxes, and confidence intervals could improve visual validation. Despite the reduction in validation bias observed with error lines, none of the tested designs yielded the desired improvements in accuracy.

Our results emphasize the limitations of relying solely on visual model validation for linear regression models in scatterplots. Further research is needed to investigate the underlying cognitive processes involved in visual validation tasks in order to find appropriate visual solutions for supporting visual model validation.

SUPPLEMENTAL MATERIALS

Supplemental material includes: 1) Python code (Jupyter-notebook) for the data and visual designs generation; 2) User study documentations including the stimuli as PDF files; and 3) Anonymized study results as CSV files.

ACKNOWLEDGMENTS

The authors would like to thank all study participants. This paper is a result of Dagstuhl Seminar 22331 “Visualization and Decision Making Design Under Uncertainty”. This work has been partially supported by the BMBF WarmWorld Project, the SANE Project, and the Risk-Principle Project. This work has also been funded in part by NSF Awards 2007436, 1452977, and 2118201.

REFERENCES

- [1] D. Altman, D. Machin, T. Bryant, and M. Gardner. *Statistics with confidence: confidence intervals and statistical guidelines*. BMJ Books, London, 2000. 3
- [2] V. Bewick, L. Cheek, and J. Ball. Statistics review 7: Correlation and regression. *Critical care*, 7(6):451–459, 2003. doi: 10.1186/cc2401 1
- [3] C. R. Blyth. On simpson’s paradox and the sure-thing principle. *Journal of the American Statistical Association*, 67(338):364–366, 1972. doi: 10.1080/01621459.1972.10482387 1
- [4] M. Bögl, W. Aigner, P. Filzmoser, T. Lammarsch, S. Miksch, and A. Rind. Visual analytics for model selection in time series analysis. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2237–2246, 2013. doi: 10.1109/TVCG.2013.222 2
- [5] D. Braun, A. Suh, R. Chang, M. Gleicher, and T. von Landesberger. Visual validation versus visual estimation: A study on the average value in scatterplots. In *2023 IEEE Visualization and Visual Analytics (VIS)*, pp. 181–185, 2023. doi: 10.1109/VIS54172.2023.00045 1, 2, 3, 4, 7, 9
- [6] G. Charness, U. Gneezy, and M. A. Kuhn. Experimental methods: Between-subject and within-subject design. *Journal of Economic Behavior & Organization*, 81(1):1–8, 2012. doi: 10.1016/j.jebo.2011.08.009 3
- [7] A. Chatzimparmpas, R. M. Martins, I. Jusufi, and A. Kerren. A survey of surveys on the use of visualization for interpreting machine learning models. *Information Visualization*, 19(3):207–233, 2020. doi: 10.1177/1473871620904671 1, 2
- [8] A. Chatzimparmpas, R. M. Martins, I. Jusufi, K. Kucher, F. Rossi, and A. Kerren. The state of the art in enhancing trust in machine learning models with the use of visualizations. *Computer Graphics Forum*, 39(3):713–756, 2020. doi: 10.1111/cgf.14034 1, 2
- [9] M. Chegini, L. Shao, R. Gregor, D. J. Lehmann, K. Andrews, and T. Schreck. Interactive visual exploration of local patterns in large scatterplot spaces. *Computer Graphics Forum*, 37(3):99–109, 2018. doi: 10.1111/cgf.13404 2
- [10] I. K. Choi, N. K. Raveendranath, J. Westerfield, and K. Reda. Visual (dis)confirmation: Validating models and hypotheses with visualizations. In *2019 23rd International Conference in Information Visualization – Part II*, pp. 116–121, 2019. doi: 10.1109/IV-2.2019.00032 2
- [11] L. Ciccione and S. Dehaene. Can humans perform mental regression on a graph? accuracy and bias in the perception of scatterplots. *Cognitive Psychology*, 128, 2021. doi: 10.1016/j.cogpsych.2021.101406 1, 2, 4, 5
- [12] M. Correll and J. Heer. Regression by eye: Estimating trends in bivariate visualizations. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pp. 1387–1396. ACM, New York, 2017. doi: 10.1145/3025453.3025922 1, 2, 3, 4, 5, 9
- [13] M. Correll, M. Li, G. Kindlmann, and C. Scheidegger. Looks good to me: Visualizations as sanity checks. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):830–839, 2019. doi: 10.1109/TVCG.2018.2864907 2
- [14] D. Dansana, R. Kumar, J. Das Adhikari, M. Mohapatra, R. Sharma, I. Priyadarshini, and D.-N. Le. Global forecasting confirmed and fatal cases of covid-19 outbreak using autoregressive integrated moving average model. *Frontiers in Public Health*, 8:580327:1–580327:11, 2020. doi: 10.3389/fpubh.2020.580327 1
- [15] A. Dasgupta, J.-Y. Lee, R. Wilson, R. A. LaFrance, N. Cramer, K. Cook, and S. Payne. Familiarity vs trust: A comparative study of domain scientists’ trust in visual analytics and conventional analysis methods. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):271–280, 2017. doi: 10.1109/TVCG.2016.2598544 2
- [16] R. M. Dawes. The robust beauty of improper linear models in decision making. *American Psychologist*, 34(7):571–582, 1979. doi: 10.1037/0003-066X.34.7.571 1
- [17] R. M. Dawes and B. Corrigan. Linear models in decision making. *Psychological Bulletin*, 81(2):95–106, 1974. doi: 10.1037/h0037613 1
- [18] E. Dimara, S. Franconeri, C. Plaisant, A. Bezerianos, and P. Dragicevic. A task-based taxonomy of cognitive biases for information visualization. *IEEE Transactions on Visualization and Computer Graphics*, 26(2):1413–1432, 2020. doi: 10.1109/TVCG.2018.2872577 7
- [19] H. Elhamedi, A. Gaba, Y.-S. Kim, and C. Xiong. How do we measure trust in visual data communication? In *2022 IEEE Evaluation and Beyond - Methodological Approaches for Visualization (BELIV)*, pp. 85–92. IEEE, New York, 2022. doi: 10.1109/BELIV57783.2022.00014 2
- [20] M. Gleicher, M. Correll, C. Nothelfer, and S. Franconeri. Perception of average value in multiclass scatterplots. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2316–2325, 2013. doi: 10.1109/TVCG.2013.183 1, 2
- [21] L. Harrison, F. Yang, S. Franconeri, and R. Chang. Ranking visualizations of correlation using weber’s law. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1943–1952, 2014. doi: 10.1109/TVCG.2014.2346979 1, 2
- [22] M.-H. Hong, J. K. Witt, and D. A. Szafir. The weighted average illusion: Biases in perceived mean position in scatterplots. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):987–997, 2022. doi: 10.1109/TVCG.2021.3114783 1, 2
- [23] J. Hullman and A. Gelman. Designing for interactive exploratory data analysis requires theories of graphical inference. *Harvard Data Science Review*, 3(3), 2021. doi: 10.1162/99608f92.3ab8a587 1, 2
- [24] A. B. M. Jeuken, P. C. Siegmund, L. C. Heijboer, J. Feichter, and L. Bengtsson. On the potential of assimilating meteorological analyses in a global climate model for the purpose of model validation. *Journal of Geophysical Research: Atmospheres*, 101(D12):16939–16950, 1996. doi: 10.1029/96JD01218 1
- [25] M. Kay and J. Heer. Beyond weber’s law: A second look at ranking visualizations of correlation. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):469–478, 2016. doi: 10.1109/TVCG.2015.2467671 1, 2
- [26] LimeSurvey Project Team / Carsten Schmitz. Limesurvey: An open source survey tool. <http://www.limesurvey.org>, 2024. 4, 8
- [27] H. H. Mahbulul Majumder and D. Cook. Validation of visual statistical inference, applied to linear models. *Journal of the American Statistical Association*, 108(503):942–956, 2013. doi: 10.1080/01621459.2013.808157 2
- [28] J. Matejka and G. Fitzmaurice. Same stats, different graphs: Generating datasets with varied appearance and identical statistics through simulated annealing. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pp. 1290–1294. ACM, New York, 2017. doi: 10.1145/3025453.3025912 1
- [29] E. Mayr, N. Hynek, S. Salisu, and F. Windhager. Trust in information visualization. In *EuroVis Workshop on Trustworthy Visualization (TrustVis)*, pp. 25–29. Eurographics Association, Goslar, 2019. doi: 10.2312/trvis.20191187 2
- [30] C. M. McColeman, L. Harrison, M. Feng, and S. Franconeri. No mark is an island: Precision and category repulsion biases in data reproductions. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1063–1072, 2021. doi: 10.1109/TVCG.2020.3030345 7
- [31] D. Moritz, L. M. Padilla, F. Nguyen, and S. L. Franconeri. Average estimates in line graphs are biased toward areas of higher variability. *IEEE Transactions on Visualization and Computer Graphics*, 30(1):306–315, 2024. doi: 10.1109/TVCG.2023.3326589 1, 2
- [32] V. Morton and D. J. Torgerson. Effect of regression to the mean on decision making in health care. *BMJ (Clinical research ed.)*, 326(7398):1083–1084, 2003. doi: 10.1136/bmj.326.7398.1083 1
- [33] T. Munzner. *Visualization Analysis and Design*. A K Peters/CRC Press, Florida, 2014. 1
- [34] T. Mühlbacher and H. Piringer. A partition-based framework for building and validating regression models. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):1962–1971, 2013. doi: 10.1109/TVCG.2013.125 2
- [35] E. Newburger, M. Correll, and N. Elmqvist. Fitting bell curves to data distributions using visualization. *IEEE Transactions on Visualization and*

- Computer Graphics*, pp. 1–12, 2022. doi: 10.1109/TVCG.2022.3210763 1
- [36] B. D. Ondov, F. Yang, M. Kay, N. Elmqvist, and S. Franconeri. Revealing perceptual proxies with adversarial examples. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1073–1083, 2021. doi: 10.1109/TVCG.2020.3030429 1
- [37] L. Padilla, S. C. Castro, and H. Hosseinpour. Chapter seven - a review of uncertainty visualization errors: Working memory as an explanatory theory. In K. D. Federmeier, ed., *The Psychology of Learning and Motivation*, vol. 74 of *Psychology of Learning and Motivation*, pp. 275–315. Academic Press, 2021. doi: 10.1016/bs.plm.2021.03.001 7
- [38] L. Padilla, R. Fygenson, S. C. Castro, and E. Bertini. Multiple forecast visualizations (mfvs): Trade-offs in trust and performance in multiple covid-19 forecast visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):12–22, 2023. doi: 10.1109/TVCG.2022.3209457 7
- [39] L. Padilla, M. Kay, and J. Hullmann. Uncertainty visualization. In W. Piegorsch, R. Levine, H. Zhang, and T. Lee, eds., *Computational Statistics in Data Science*, pp. 405–421. Wiley, 2022. 7
- [40] D. Peskov, B. Cheng, A. Elgohary, J. Barrow, C. Danescu-Niculescu-Mizil, and J. Boyd-Graber. It takes two to lie: One to lie, and one to listen. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 3811–3854. ACL, Stroudsburg, 2020. doi: 10.18653/v1/2020.acl-main.353 2
- [41] D. R. Proffitt, M. Bhalla, R. Gossweiler, and J. Midgett. Perceiving geographical slant. *Psychonomic Bulletin & Review*, 2:409–428, 1995. doi: 10.3758/BF03210980 4, 5
- [42] Prolific. Prolific - quickly find research participants you can trust. <https://www.prolific.co>. Location: London, UK. Accessed: 2024-03-12. 4, 8
- [43] D. Reimann, C. Blech, N. Ram, and R. Gaschler. Visual model fit estimation in scatterplots: Influence of amount and decentering of noise. *IEEE Transactions on Visualization and Computer Graphics*, 27(9):3834–3838, 2021. doi: 10.1109/TVCG.2021.3051853 1
- [44] R. A. Rensink and G. Baldrige. The perception of correlation in scatterplots. *Computer Graphics Forum*, 29(3):1203–1210, 2010. doi: 10.1111/j.1467-8659.2009.01694.x 1, 2
- [45] B. E. Rogowitz, L. A. Treinish, and S. Bryson. How not to lie with visualization. *Computers in Physics and IEEE Computational Science & Engineering*, 10(3):268–273, 1996. doi: 10.1063/1.4822401 2
- [46] I. T. Ruginski, A. P. Boone, L. M. Padilla, L. Liu, N. Heydari, H. S. Kramer, M. Hegarty, W. B. Thompson, D. H. House, and S. H. Creem-Regehr. Non-expert interpretations of hurricane forecast uncertainty visualizations. *Spatial Cognition & Computation*, 16(2):154–172, 2016. doi: 10.1080/13875868.2015.1137577 7
- [47] G. Strain, A. J. Stewart, P. Warren, and C. Jay. Adjusting point size to facilitate more accurate correlation perception in scatterplots. In *2023 IEEE Vis X Vision*, pp. 1–5, 2023. doi: 10.1109/VisXVision60716.2023.00006 1, 2
- [48] G. Strain, A. J. Stewart, P. Warren, and C. Jay. The effects of contrast on correlation perception in scatterplots. *International Journal of Human-Computer Studies*, 176, 2023. doi: 10.1016/j.ijhcs.2023.103040 1, 2
- [49] A. Sultan, W. Sařabun, S. Faizi, and M. Ismail. Hesitant fuzzy linear regression model for decision making. *Symmetry*, 13(10), 2021. doi: 10.3390/sym13101846 1
- [50] D. A. Szafir, S. Haroz, M. Gleicher, and S. Franconeri. Four types of ensemble coding in data visualizations. *Journal of Vision*, 16(5):11:1–11:19, 2016. doi: 10.1167/16.5.11 1, 2
- [51] W. M. Vagias. Likert-type scale response anchors. *Clemson International Institute for Tourism & Research Development, Department of Parks, Recreation and Tourism Management. Clemson University*, 2006. 3
- [52] C. Xiong, C. Stokes, Y.-S. Kim, and S. Franconeri. Seeing what you believe or believing what you see? belief biases correlation estimation. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):493–503, 2023. doi: 10.1109/TVCG.2022.3209405 1, 2, 3
- [53] F. Yang, L. T. Harrison, R. A. Rensink, S. L. Franconeri, and R. Chang. Correlation judgment and visualization features: A comparative study. *IEEE Transactions on Visualization and Computer Graphics*, 25(3):1474–1488, 2019. doi: 10.1109/TVCG.2018.2810918 1, 2, 3, 6, 7
- [54] L. Yuan, S. Haroz, and S. Franconeri. Perceptual proxies for extracting averages in data visualizations. *Psychonomic Bulletin & Review*, 26(2):669–676, 2019. doi: 10.3758/s13423-018-1525-7 1, 2
- [55] Y. Zhang, Y. Sun, L. Padilla, S. Barua, E. Bertini, and A. G. Parker. Mapping the landscape of covid-19 crisis visualizations. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI '21, article no. 608, 23 pages. ACM, New York, 2021. doi: 10.1145/3411764.3445381 7