DOI: 10.1111/cgf.13325
COMPUTER GRAPHICS *forum*
Volume 37 (2018), number 6 pp. 300–312

A Study of the Effect of Doughnut Chart Parameters on Proportion Estimation Accuracy

X. Cai[1], K. Efstathiou[2], X. Xie[1], Y. Wu[1], Y. Shi[3] and L. Yu[2,3]

[1]State Key Lab of CAD&CG, Zhejiang University, China
{xwcai, xxie, ycwu}@zju.edu.cn
[2]University of Groningen, Johann Bernoulli Institute for Mathematics and Computer Science, The Netherlands
K.Efstathiou@rug.nl, mail@yulingyun.com
[3]Hangzhou Dianzi University, China
hzshiyan@gmail.com

Abstract

Pie and doughnut charts nicely convey the part–whole relationship and they have become the most recognizable chart types for representing proportions in business and data statistics. Many experiments have been carried out to study human perception of the pie chart, while the corresponding aspects of the doughnut chart have seldom been tested, even though the doughnut chart and the pie chart share several similarities. In this paper, we report on a series of experiments in which we explored the effect of a few fundamental design parameters of doughnut charts, and additional visual cues, on the accuracy of such charts for proportion estimates. Since mobile devices are becoming the primary devices for casual reading, we performed all our experiments on such device. Moreover, the screen size of mobile devices is limited and it is therefore important to know how such size constraint affects the proportion accuracy. For this reason, in our first experiment we tested the chart size and we found that it has no significant effect on proportion accuracy. In our second experiment, we focused on the effect of the doughnut chart inner radius and we found that the proportion accuracy is insensitive to the inner radius, except the case of the thinnest doughnut chart. In the third experiment, we studied the effect of visual cues and found that marking the centre of the doughnut chart or adding tick marks at 25% intervals improves the proportion accuracy. Based on the results of the three experiments, we discuss the design of doughnut charts and offer suggestions for improving the accuracy of proportion estimates.

Keywords: user studies, information visualization

ACM CCS: Human-centered computing->Visualization; Empirical studies in visualization

1. Introduction

The doughnut chart is a variant of the pie chart, where a centre disk has been removed and the remaining ring is divided into slices (see Figure 1). Both types of charts, doughnut and pie, nicely convey the part–whole relationship, and for this reason they are being extensively used for showing proportions. Despite its prevalence, the pie chart has long been criticized by information visualization experts. The history of the pie chart and the debate around its use has been reviewed in, among others, [Spe05] and [SL91].

Doughnut charts share several similarities with pie charts and one can consider the latter as a special case of the former where the inner radius becomes zero. Compared to pie charts, doughnut charts have the advantage that their structure can be adapted to the presentation of extra information. Some common adaptations are ‘sunbursts’ [SZ00] supporting the representation of hierarchical data by using multiple rings (e.g. [YYQ17, WLL*16, GND*18]), and chord diagrams [KSB*09] where the hole is used for drawing connections between different slices. At the same time, doughnut charts emphasize different visual encodings compared to pie charts. For example, in pie charts, explicit information of angle can be leveraged to estimate proportions while in doughnut charts angle can be only indirectly inferred. Such differences mean that study results for pie charts cannot be directly applied to doughnut charts.

Many experiments (e.g. [SH87, SL91]) have been carried out on human perception of the pie chart, mainly focusing on its accuracy and effectiveness. Studies comparing pie charts to ‘rectangular’ charts (such as bar charts or waffle charts) show that the former are

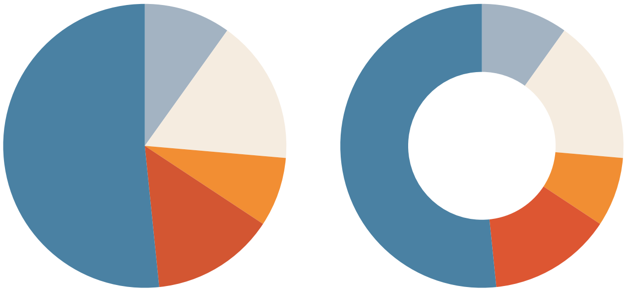

Figure 1: *Pie chart (left) and doughnut chart (right).*

not inferior to the latter for proportion estimation as we describe in detail in the review of related work in Section 2. However, 'round' charts are perceived differently than 'rectangular' charts [ZK10a, ZK10b] and hence their use may be preferable in certain contexts. Moreover, as described in the previous paragraph, doughnut charts have advantages that make them suitable for specific graphical representations.

The aim of the present work is not to compare doughnut charts to other chart types but to find out how to improve the proportion estimation accuracy of doughnut charts for those cases where the use of such charts is preferred. We study this question in two, complementary, ways. First, we determine how the two fundamental design parameters of doughnut charts (outer and inner radius) affect the accuracy of proportion estimates. Secondly, we explore the effect of additional visual cues on the accuracy so that we can make specific suggestions on the use of such cues in the design of doughnut charts. We are not aware of any previous studies on the effect of such additional visual cues for doughnut charts.

Therefore, in this paper, we carry out three experiments to explore the role of the fundamental design parameters and additional visual cues for the accuracy of proportion estimation in doughnut charts. After a review of related work in Section 2, we present the first experiment and its results in Section 3. In this experiment, we test the effect of the overall doughnut size to the accuracy of proportion estimation. In the second experiment, presented in Section 4, we test the effect of the inner radius of the doughnut chart. Then, in Section 5, we present the third experiment which tests the effect of visual cues (marking the centre of the doughnut chart and adding tick marks) on accuracy. Finally, in Section 6 we critically reflect on these results and provide suggestions for the design of doughnut charts.

2. Related Work

Early handbooks on chart design provide suggestions and guidance 'based primarily on authors' intuitions drawn from the wisdom of practice' [FCB01]. More recent work focuses on evidence coming from experiments on graphical perception and empirical studies; see [FCB01, HB10] for reviews of them. This section focuses on the work related to doughnut charts and pie charts. These charts have been extensively used in different applications [WZM*16, SWLL13, LCWL14, MMWC17], such as sport analytics [WLS*18], urban informatics [CYW*16, LWL*17], weather forecast [KTB*18], medical and health data analysis [HPvU*18, CWS*17] and user behaviour analysis [BSBE17, WWL*10, CLS*12].

The pie chart has a history of over 200 years and it is a widely used method for displaying proportion data, especially in popular media. However, its use has not been without its detractors (see [Tuf01, Cle94, Rob05]). For example, Tufte [Tuf01] remarks that pie charts are the worst design 'to show exact numerical numbers' and disapproves of their use 'given their low-data density and failure to order numbers along a visual dimension'. Generally, since pie charts were assumed to be less accurate than other widely used charts (e.g. scatterplots [CCM*14, CZC*15] and bar charts [XCH*14]), they were dismissed by data analysts. For a review of the history of the pie chart and the debate around its use, we refer to [Spe05, SL91].

While such debate provides insightful guidelines for the use of pie charts, in our paper we focus on empirical works, which aim to provide quantitative evidence of the effectiveness of pie charts. Eells [Eel26] appears to be the first to study the effectiveness of pie charts in comparison to bar charts and concluded that pie charts are more accurate than bar charts in presenting component parts. However, von Huhn [VH27] challenged Eells's work and studied different aspects of the question concluding that 'it seems that the only case where the circle may be preferable to the bar is where a single total with rather numerous component parts is to be shown, and where the parts need to be presented not only singly but also in groups'. This issue led to a number of early experimental studies (e.g. [CS32, CS27]) yielding new findings but failing to settle the question [SL91]. More recently, Cleveland [CM84] used the pie chart and the bar chart to study the accuracy of different elementary tasks (elementary visual encodings) and found that position can be more accurately read compared to angle. Simkin and Hastie [SH87] compared the pie chart and two types of bar charts (simple and divided) for different types of judgement. Their results show the interaction of chart type and judgement type. In particular, they show that for comparison judgements pie charts were the least accurate, while for proportion judgements pie charts and simple bar charts were equally accurate and better than divided bar charts. Similar results were obtained in a series of experiments involving pie charts in [Spe90, SL91, HS92, HS98], suggesting that the pie chart is not inferior to the bar chart in proportion judgements and gains an advantage when the number of components increases.

A more fundamental question is what is the major visual encoding used when decoding numerical information in pie charts. There is consensus that area, arc length and angle information can be extracted from a pie chart, but how perception really works and which one of these encodings is most important remains unclear and controversial. Cleveland and McGill [CM84] assumed that people mainly decode angle information in the pie chart, but suggested that area and arc length may also play a role. The influence of [CM84] is great and many researchers have followed these ideas (e.g. [SH87]), but this has not definitively settled the question as there is not sufficiently strong empirical evidence to support the angle hypothesis to the exclusion of other interpretations. On the same question, Spence and Lewandowsky [SL91] suggest that subjects probably pay no attention to area when presented with a pie chart. Skau and Kosara [SK16] recently tested the effectiveness of individual data encodings (such as the arc, the angle and the area) in pie and doughnut charts. Their results suggest that angle is the least important visual cue for both charts, while both chart types are equally accurate for proportion estimation. We note here that the present work and [SK16] have one similar experiment, namely, our Experiment 2 and their Study 2, which studied the effect of the inner radius on proportion estimation accuracy of doughnut charts. The results of the two experiments are the same although the details of the experimental setup are different. We compare these two experiments in Section 4.3.

In general, doughnut charts have received much less research attention compared to pie charts even though they share similar characteristics. Kosara and Ziemkiewicz [KZ10] compare pie charts, bar charts, doughnut charts and square pie charts in a work related to the design and practicality of online studies. They find that square pie charts are the most accurate while doughnut and pie charts were equally accurate for percentage estimation. Skau and Kosara [SK16] compare pie charts to doughnut charts of different inner radii and they find that the inner radius does not affect the proportion accuracy, except for the case of a very thin doughnut chart. Siirtola [Sii14] compares doughnut charts with bar charts, pie charts and tables in the task of 'perceiving the relative order of the parts of some whole'. The results show that bar charts are superior to doughnut charts and to pie charts, while there is no significant difference between the latter two for this type of task. Although not related to the question of accuracy we would also like to mention here the work by Ziemkiewicz and Kosara [ZK10a, ZK10b] which compares different types of charts (including the doughnut and pie charts) in terms of their semantic aspects.

There are three main choices of judgement tasks for comparing different types of charts: discrimination, comparison and proportion estimation [SH87]. Eells [Eel26] used proportion estimation, and this has been continued in most subsequent empirical studies. However, von Huhn [VH27] criticized Eells's work for the absence of comparison tasks, which are often required in graphical analysis. Spence [SL91] describes magnitude estimation as a 'sensible psychological task for experiments comparing different types of charts', but he also suggests it makes no sense to convey precise data in graphical form rather than in tabular form. He also argues that magnitude estimation 'does not reflect how people use graphs in real life'. Simkin and Hastie [SH87] studied the spontaneous response of 200 undergraduate students to different types of chart. The results showed that most people make comparisons when presented with bar charts and make proportion judgements when presented with pie charts, indicating that people have certain expectations for the use of these charts and the information conveyed by them.

Instead of asking which chart types are better suited for particular tasks, we can consider the complementary question of how we can improve the design of specific chart types. Work in this direction explores how additional visual cues and reference structures can improve perception. Skau *et al.* [SHK15] studied visual embellishments in bar charts and found that they indeed have an impact on the effectiveness of perceiving the chart data. Reference structures such as grids and tick marks are frequently recommended for data charts to aid in relating content to axes [Kos06]. Robbins [Rob05] found that adding grids and tick marks to the design of glyphs could improve performance for the task of reading the exact data values. However, Fuchs *et al.* [FIB*14] found that the star glyph without added reference structures (such as grid lines or tick marks) and without contour lines performs best for similarity search tasks.

Related to reference structures is the fact that 'people make estimates by starting from an initial value that is adjusted to yield the final answer' as observed by Tversky and Kahneman [TK75] who called this phenomenon *anchoring*. Simkin and Hastie [SH87] found a similar 'anchoring' process in graph perception: the anchor is an initial segmentation basis for the estimations and the accuracy of the proportions that locate close to the natural anchors (for example, 0%, 25%, 50% and 75% of pie charts, 0% and 50% of simple bar charts) would be enhanced. However, pie charts and bar charts provide different natural anchors [SH87, Spe05] and the more accurate anchoring accounts for the superiority of pie charts in proportion estimation tasks [SH87]. It would be interesting to see if a similar 'anchoring' process also exists in reading doughnut charts and whether the accuracy would be improved if a reference structure (such as tick marks) is shown.

In summary, the comparison of different types of charts has been extensively studied in the literature. Nevertheless, research on doughnut charts has not been equally extensive. Even though doughnut charts have similarities with pie charts they also offer advantages (e.g. the ability to present extra information in the centre hole [CGS*11, LCW*15]) and therefore a more thorough study of the effect of their design parameters is necessary. In this work, we focus on the accuracy of doughnut charts for proportion estimation following Cleveland *et al.* [CM84]. We adopt a point of view where we study the doughnut chart as a whole, without trying to separate the effect of different visual encodings, and instead we focus on the effect of the basic design parameters of the chart on its (proportion estimation) accuracy. Our experiments (Experiments 1 and 3) which as far as we are aware have not been performed before, as well as our new findings, have implications for visualization practitioners and the design of more effective and accurate doughnut charts.

3. Experiment 1: The Effect of Size on Accuracy

In our first experiment, we study whether the size of the doughnut chart (its outer radius) affects proportion reading accuracy. Our motivation for this is that mobile devices, having screens of constrained size, are becoming the main reading devices and we wanted to clarify whether the chart size affects the accuracy. In this section, we also explain in detail several choices on the experimental setup and results analysis that are the same in the other two experiments. The three experiments involved different groups of participants. The orderings of the conditions were randomized in each experiment, thus there was no sequence effect (training effect or learning effect) in these experiments.

3.1. Design and procedure

Apparatus: The experiment was conducted on an iPad Air (9.7″ diagonal screen size and resolution 2048 × 1536 pixels). The experiment setup was implemented as a web application accessed through the Safari web browser. Even though the web application could be remotely accessed, all experiments took place in the lab (see Figure 2a). All participants used the same device, set at the

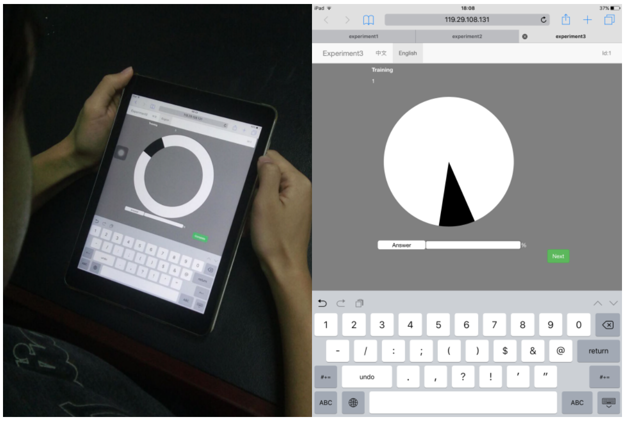

(a) (b)

Figure 2: *(a) A participant is taking our experiment. (b) The interface of the web application for the study.*

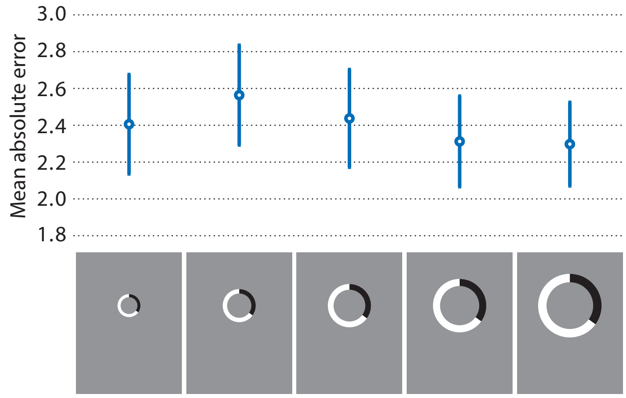

Figure 3: *Mean absolute error and 95% confidence intervals in Experiment 1. The pictures below the chart represent the conditions D_0 to D_4; the dark grey area represents the iPad Air screen and the doughnut charts are drawn to scale.*

same brightness. In each trial, the chart was shown at the top of the screen and a text field accepting numeric input was provided for the answers (see Figure 2b). The on-screen software keyboard was used for input. The text field automatically received focus at the beginning of the test, causing the software keyboard to appear, set for numeric input. After participants had input an answer they could either use the 'Return' key on the software keyboard, or the 'Next' button shown in Figure 2(b), to proceed to the next test.

Participants: A total of 32 participants (17 female) participated in Experiment 1. Their age ranged from 18 to 30 years (mean 22.5). All reported normal or corrected-to-normal vision. Sixteen of them were undergraduate students and the rest were postgraduate students; 12 participants studied Art Design and the rest studied Digital Media Technology. Only one participant reported to be very familiar with quantitative estimation tasks, 16 reported limited familiarity and the rest reported unfamiliarity. A total of 29 of 32 participants reported having no more than 1-year experience with pie or doughnut charts. Each participant received 10 Yuan (Chinese RMB) for participating.

Task: We used five different sizes of doughnut charts (see Figure 3). The values for the outer radius of the doughnut chart were 1.60, 2.31, 3.02, 3.73 and 4.44 cm. The size of the inner radius was fixed to 73.5% of the outer radius.

The five values of the outer radius were determined as follows. We considered five (virtual) screens with 4:3 height–width ratio and diagonal size ranging from a minimum value $d_{\min} = 3.5''$ (roughly corresponding to an iPhone 4) to a maximum value $d_{\max} = 9.7''$ (corresponding to an iPad); the intermediate three diagonal sizes were equally spaced between the two extremes. Then the chart diameter (twice the outer radius) was chosen as 60% of the width of such screen, that is, 36% of the diagonal. In all cases, the corresponding chart was shown on the same device, described above.

We used black colour to show the proportion that should be estimated, and white colour to show the complementary proportion (see Figure 2). Participants were always asked to estimate the proportion of the black part to the whole. The background colour was set to dark grey (50% black).

The proportion values were limited below 50% because we want to avoid the possibility that for proportions over 50% some participants evaluate the complementary sector and subtract their estimate from 100% while other participants try to estimate the proportion directly. The whole set of proportions consisted of 48 values (integer numbers from 1% to 49%, excluding 25%). We partitioned the whole set into 12 subsets, each containing four consecutive numbers, that is, 1–4%, 5–8%, ..., 21–24%, 26–29%, ..., 46–49%.

Overall, our experiment consisted of

$$5 \text{ (sizes)} \times 12 \text{ (proportions)} \times 32 \text{ (participants)} = 1920 \text{ trials.}$$

Every participant took 12 tests for each size condition; there was one test for each one of the 12 proportion subsets and, for each test, a proportion was randomly chosen from the four proportions in the corresponding subset. The order that proportion subsets were tested was randomized. Moreover, since we had 32 participants we arranged that each proportion from each subset was chosen exactly eight times throughout the experiment for all participants. The angular position of the black part was also randomized.

Procedure: At the beginning, five training tests were given and the correct answers were shown afterwards to help learning. The actual experiment started after the training session. During the actual experiment, participants were asked to judge the proportion and give an answer, as soon and precisely as possible. However, the correct answer was no longer provided. There were three blocks of tests in the actual experiment and participants could take a break between blocks. For each test, we recorded the participant ID, the size condition, the proportion condition, the estimate given by the participant and the time from the moment the test appeared on screen to the moment the participant confirmed their answer. After participants finished the actual experiment, they were required to complete a questionnaire for background information and preference questions. The typical time for the whole experiment, including the training session and the questionnaire, was about 20 min.

Table 1: *Mean absolute error and 95% confidence intervals in Experiment 1.*

Condition	Radius (cm)	Mean	95% CI
D_0	1.60	2.4058	[2.1359, 2.6756]
D_1	2.31	2.5640	[2.2930, 2.8349]
D_2	3.02	2.4375	[2.1722, 2.7028]
D_3	3.73	2.3123	[2.0665, 2.5582]
D_4	4.44	2.2977	[2.0708, 2.5245]

3.2. Results

The first step in processing the experiment data was to clean up answers that we attributed to mistypes. We removed from the data 7 answers where the absolute error was more than 20%.

In previous studies, two types of error measures have been used to analyse such data. Defining the estimation error as

$$\Delta p = \text{judged proportion} - \text{true proportion},$$

we then have the *absolute error* $|\Delta p|$, and the *log-absolute error*

$$\lambda(\Delta p) = \log_2(|\Delta p| + 1/8),$$

see [CM84]. Using the log-absolute error has the effect of making small errors more prominent while suppressing larger errors. In our setting, where the absolute error is always a non-negative integer, the log-absolute error gives a large difference between exact answers and 1% error. More precisely, an exact answer gives log-absolute error $\lambda(0) = -3$ while errors of $\pm 1\%$ and $\pm 8\%$ give respective log-absolute errors $\lambda(\pm 1) \simeq 0.17$ and $\lambda(\pm 8) \simeq 3.02$. This means that an 1% error is penalized with respect to the exact answer, as much as an 8% error is penalized with respect to 1% error. For this reason, and since we consider small errors to be sufficiently good for the given task, we do not want to have them severely penalized and we prefer to work with the absolute error $|\Delta p|$.

We report the results for the mean absolute error for the conditions D_0 to D_4 in Experiment 1 in Table 1 and Figure 3. We are interested on whether the effect of the chart size on the mean absolute error is significant. First, note in Figure 3 that the mean absolute error in most conditions lies inside the 95% confidence intervals of other conditions, with the exceptions of the pair (D_1, D_4) and (D_1, D_3).

Given the within-subjects design of this experiment, we used repeated-measures ANOVA to test the mean absolute error of each person in each size after verifying that the data obey the hypotheses of normality and homoscedasticity. The outcome of the test was that the main effect of size is not significant ($F(4, 124) = 0.949$, $p = 0.438 > 0.05$).

In the questionnaire, 28 of 32 participants (87.5%) reported that their estimation strategy involved estimating easy to discern proportion sizes, for example 25% or 50%, and making their proportion estimates by comparing to these sizes. A total of 20 of 32 participants (62.5%) reported that they believed that the chart size affected their estimates while the rest reported little or no effect.

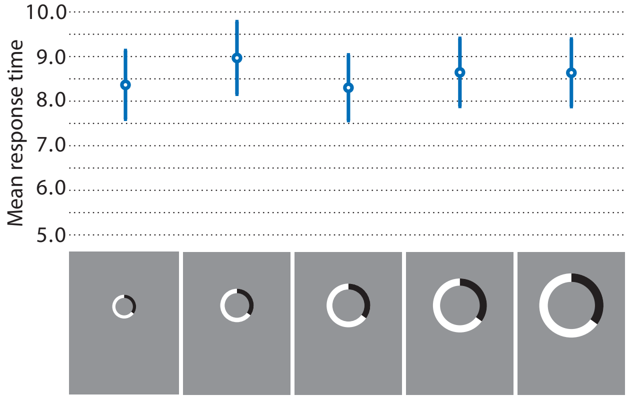

Figure 4: *Mean response time and 95% confidence intervals in Experiment 1.*

Table 2: *Mean response time and 95% confidence intervals in Experiment 1.*

Condition	Radius (cm)	Mean (s)	95% CI
D_0	1.60	8.365	[7.598, 9.132]
D_1	2.31	8.969	[8.203, 9.735]
D_2	3.02	8.302	[7.536, 9.067]
D_3	3.73	8.645	[7.877, 9.067]
D_4	4.44	8.636	[7.870, 9.402]

3.3. Discussion

The main outcome of Experiment 1 is that size does not have any statistically significant effect on the accuracy or speed for proportion estimation tasks for the range of sizes tested. This is in contrast to the subjective opinion of most participants who reported in the questionnaire that the size affected the accuracy of their estimates. The result indicates that for small screens such as those typically found on mobile devices one can safely use chart sizes with diameter $\simeq$3 cm.

In this experiment, as well as in the following ones, we gave unlimited time to the participants to make their estimates since the response time is not our focus in this paper. This occasionally led to very large response times. The response times for each trial varied from 1.11 to 86.15 s, with the mean response time at 8.58 s and 95% confidence interval [8.24 s, 8.92 s]. We report the mean response time for the conditions D_0 to D_4 in Figure 4 and Table 2. Given the within-subjects design of this experiment and the fact that the data obey the hypotheses of normality and homoscedasticity, we used repeated-measures ANOVA to test the mean response time of each person in each size and found that the main effect of size on time is not significant ($F(4, 124) = 0.845$, $p = 0.5 > 0.05$). There does not seem to be a correlation between response times and accuracy: Goodman and Kruskal's gamma test gives $p = 0.1612$ for the null hypothesis H_0 that there is no correlation. Nevertheless, it would be interesting to check how the results would be affected if participants had adopted a more casual attitude or if a time limit was imposed.

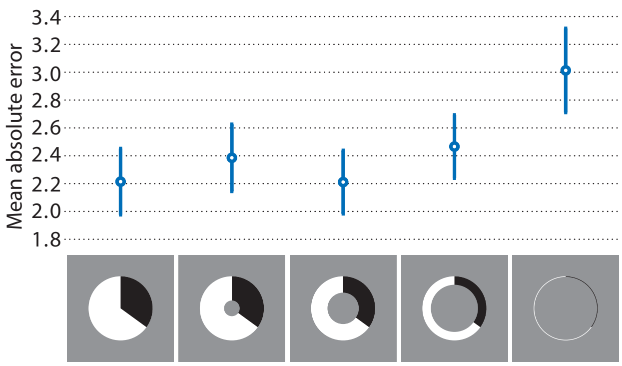

Figure 5: *Mean absolute error and 95% confidence intervals for conditions IR_0 to IR_4 in Experiment 2. Mean absolute error of IR_4 is significantly larger than the other conditions.*

Table 3: *Mean absolute error and 95% confidence intervals (CI) for conditions IR_0 to IR_4 in Experiment 2.*

Condition	Inner radius (%)	Mean	95% CI
IR_0 (pie)	0	2.214	[1.966, 2.462]
IR_1	24.5	2.385	[2.138, 2.633]
IR_2	49	2.211	[1.963, 2.459]
IR_3	73.5	2.466	[2.218, 2.714]
IR_4 (thin ring)	98	3.013	[2.765, 3.261]

Table 4: *p-Values of post hoc Student's* t*-test with Bonferroni correction on the dependence of absolute error on the inner radius. Numbers marked by * and ** represent significance at the 0.05 and 0.01 levels, respectively.*

	IR_0	IR_1	IR_2	IR_3	IR_4
IR_0		1.000	1.000	0.426	0.002**
IR_1	1.000		1.000	1.000	0.019*
IR_2	1.000	1.000		0.649	0.006**
IR_3	0.426	1.000	0.649		0.025*
IR_4	0.002**	0.019*	0.006**	0.025*	

4. Experiment 2: The Effect of Inner Radius on Accuracy

After finding in Experiment 1 that the overall size of the doughnut chart does not affect the chart accuracy, in this experiment we explore another fundamental question: does the size of the inner radius of a doughnut chart affect the accuracy of proportion estimation? We are interested in this question because the inner radius is the only factor to distinguish pie and doughnut charts. If we fix the outer radius of a doughnut chart and increase its inner radius, the doughnut becomes thinner, transforming from a full disk (inner radius 0) to a thin ring.

4.1. Design and procedure

Apparatus: The same setup was used as in Experiment 1.

Participants: A total of 32 participants (eight female) participated in Experiment 2. Their age ranged from 20 to 28 years (mean 23.2) and all reported normal or corrected-to-normal vision. Twenty of them were undergraduate students and the rest were postgraduate students, all of them studied Computer Science. In terms of familiarity with quantitative estimation task 23 participants reported to be familiar or very familiar, 5 moderately familiar and 4 unfamiliar or very unfamiliar. Each of the participants received 10 Yuan (Chinese RMB) as reward.

Task: We used five conditions of inner radius (see Figure 5). More specifically, we considered doughnut charts where the inner radius is 0% (a pie chart), 24.5%, 49%, 73.5% and 98% (a thin ring) of the outer radius. We refer to these conditions as IR_0 to IR_4. For each of these five conditions, we considered 12 proportion conditions, following the same setting as in Experiment 1, see Section 3.1. Overall, Experiment 2 consisted of

$$5 \text{ (conditions)} \times 12 \text{ (proportions)} \times 32 \text{ (participants)} = 1920 \text{ trials.}$$

Procedure: The procedure was the same as in Experiment 1. For each test, we recorded the participant ID, the inner radius condition, the proportion condition, the estimate given by the participant and the time in milliseconds from the moment the test appeared on screen to the moment the participant confirmed their answer. The typical time for the whole experiment, including the training session and the questionnaire, was approximately 15 min.

4.2. Results

Cleaning the data in the same way as in Experiment 1 resulted in removing two trials (of 1920) from the results.

The mean absolute error for each of the conditions IR_0 to IR_4 is shown in Table 3 and Figure 5. We can see that the thinnest doughnut chart (IR_4) leads to the largest mean absolute error and the rest conditions (IR_0, IR_1, IR_2 and IR_3) are at the same level. Given the within-subjects design of this experiment, and the fact that the data obey the hypotheses of normality and homoscedasticity, we used repeated-measures ANOVA to test the mean absolute error of each person in each inner radius condition and found that the inner radius has a statistically significant effect ($F(4, 124) = 8.026$, $p < 0.001$). Since the main effect of inner radius is significant, we used Student's t-test with Bonferroni correction for the *post hoc* test, as shown in Table 4. Expect that the mean absolute error of IR_4 is significantly higher than those of the other conditions, no other significant difference is detected.

In the post-survey questionnaires, 25 of 32 participants reported similar strategies as the participants in Experiment 1. One participant reported to base their estimate on the position of the centre of the circle, two participants reported to base their estimates on comparing the corresponding arc to the entire circle and three participants reported to base their estimates on comparing the corresponding angle to the ring sector. Twenty-four participants chose the pie chart (IR_0), two chose IR_1, three chose IR_2, two chose IR_3 and one chose IR_4 as the most accurate conditions. A total of 19 participants chose IR_0, four chose IR_1, eight chose IR_2 and one chose IR_3 as their most preferred conditions. Therefore, for most participants, the pie chart

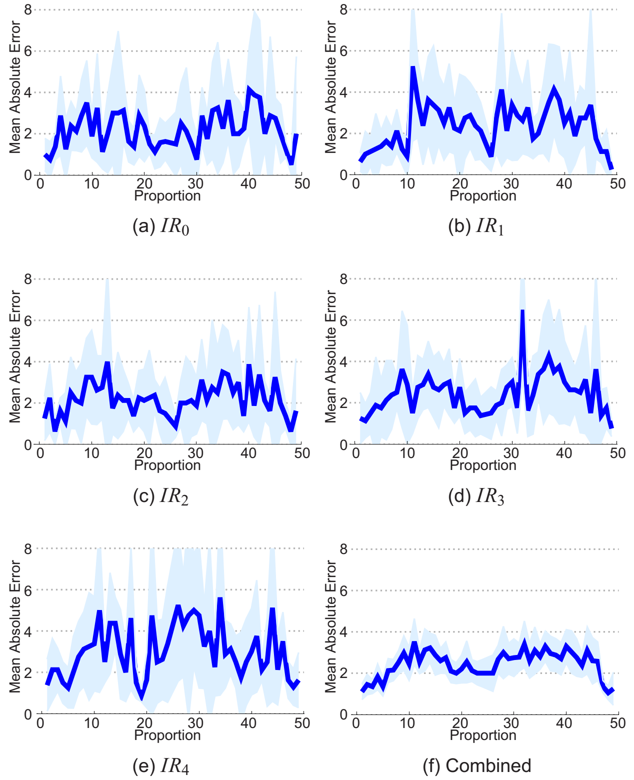

Figure 6: *Mean absolute error for each proportion and condition in Experiment 2. The light blue band represents the corresponding 95% confidence interval. Panel (f) presents the average value of the mean absolute error for all conditions.*

is their most preferred condition, and is also the one that they felt most accurate.

4.3. Discussion

The results of Experiment 2 show that the proportion accuracy of the doughnut chart is insensitive to the inner radius, except for the case of the thinnest doughnut charts, which gave the least accurate result. Removing IR_4 from the test shows that for the remaining conditions the radius does not have a significant effect anymore ($F(3,93) = 1.609$, $p = 0.193 > 0.05$). This indicates that the significant effect of the inner radius on accuracy is mainly the result of the thinnest doughnut chart (IR_4). These results align with similar results by Skau and Kosara [SK16] who investigate the effect of the inner radius and find that it makes no difference to the accuracy, except when the doughnut chart becomes extremely thin.

In our experiment, each proportion from 1% to 49% (excluding 25%) was tested the same number of times. This facilitates the comparison of the accuracy for different proportions. Figure 6 presents the mean absolute error for each proportion and for each inner radius condition. The results seem to support the idea that the accuracy improves for proportions near integer multiples of 25%. We see that for the conditions IR_1, IR_2 and IR_3 the mean absolute error becomes smaller near 0%, 25% and 50% giving a characteristic 'M' shape. This effect is less pronounced in condition IR_0 (pie) and does not appear in condition IR_4 (thin ring). We speculate this is a manifestation of 'anchoring' [SH87]. Moreover, the absence of the 25% 'anchoring' effect in condition IR_4 seems to be the main cause of the larger mean absolute error in this condition. Therefore, we conjecture that adding tick marks at the 25% anchor improves the reading accuracy. We study this conjecture in Experiment 3.

Table 5: *Mean response time (s) and 95% confidence intervals (CI) for conditions IR_0 to IR_4 in Experiment 2.*

Condition	Inner radius (%)	Mean	95% CI
IR_0 (pie)	0	8.718	[8.108, 9.329]
IR_1	24.5	8.596	[7.986, 9.206]
IR_2	49	8.719	[8.110, 9.329]
IR_3	73.5	8.767	[8.157, 8.996]
IR_4 (thin ring)	98	8.319	[7.708, 8.929]

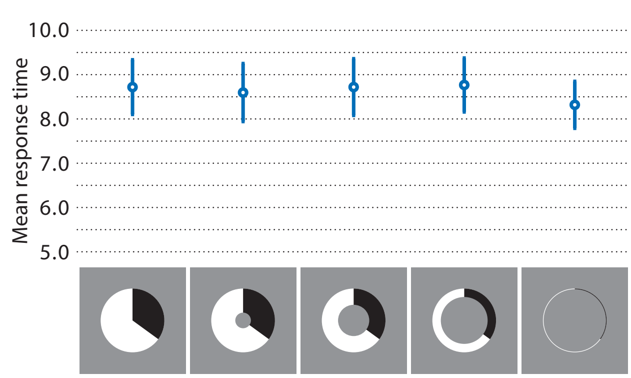

Figure 7: *Mean response time and 95% confidence intervals for conditions IR_0 to IR_4 in Experiment 2.*

We did not find a significant effect of the inner radius on the response time. The mean response time for each of the conditions IR_0 to IR_4 is shown in Table 5 and Figure 7. We used repeated-measures ANOVA to test the mean response time of each person in each inner radius condition and found that the inner radius does not have a significant effect on the response time ($F(4,124) = 0.540$, $p = 0.707 > 0.05$). The completion times for each trial varied from 1.33 to 76.01 s, with the mean completion time at 7.71 s and 95% confidence interval [7.52 s, 7.90 s]. There does not seem to be a correlation between completion times and accuracy, that is, longer observations do not lead to more accurate estimations: Goodman and Kruskal's gamma test gives $p = 0.5629$ for the null hypothesis H_0 that there is no correlation.

As mentioned in Section 2, both Experiment 2 in this work and Study 2 in [SK16] study the effect of the inner radius to the proportion estimation accuracy. The results of both experiments agree. However, the context of the two experiments and the experimental methods are different. Skau and Kosara compare in [SK16] individual encodings (arc length, area or angle) in chart reading and their three studies aim to answer the question of which encoding is the most important one. In contrast, our work does not focus on individual encodings but aims to study the effect of doughnut chart design parameters (inner radius, outer radius and additional visual cues) on proportion estimation accuracy and what can be improved in the chart design (e.g. using additional visual cues). Moreover, our experiment was done in a laboratory setting instead of using crowdsourcing as in [SK16]. Compared to crowdsourcing studies, the laboratory setting allowed us to more easily control the experimental environment and closely observe the entire process at the cost of having fewer participants.

5. Experiment 3: The Effect of Visual Cues on Accuracy

In the first two experiments, we studied the effect on proportion estimation accuracy of the two most fundamental visual parameters of doughnut chart: size and inner radius. In Experiment 3, we explore whether we can improve the accuracy of doughnut charts for proportion estimation by using two additional visual cues, that is, by marking the centre or by adding tick marks at fixed positions.

The reason for adding tick marks is that during the previous experiments we observed some participants trying to make a gesture to help them visualize a proportion of 25%. This strategy was confirmed in the post-survey questionnaires: 54 of 60 participants reported they made their judgements comparing to virtual proportions of 25% or 50% in the doughnut chart. Moreover, recall that in Section 4.3 we discussed how the accuracy of the chart improves near 0%, 25% and 50%, manifesting as the 'M' shape in Figure 6(f). These observations indicate that proportions around 0%, 25% and 50% are easier to estimate. Thus, we hypothesized that the proportion estimate accuracy could be improved by providing a visual cue around such proportions.

Another observation during the previous experiments, was that some participants tried to figure out the position of the centre in doughnut charts and use it as a reference point. Moreover, Experiment 2 showed that the increase in inner radius has a negative effect to accuracy and we suspect that this occurs because, as the hole size increases, the centre position becomes more difficult to discern. Therefore, we test whether marking the centre improves the accuracy of proportion estimates.

Finally, pie and doughnut charts provide a natural visual cue—to estimate a proportion we compare it with the full chart. We were interested to explore the effect of removing this visual cue, that is, removing the part of the chart that is complementary to the estimated proportion. We thus introduce the 'anti-cue' of incomplete doughnut charts, where only the proportion to estimate is shown.

Summarizing, we considered the following two visual cues and a visual anti-cue (see also Figure 8).

Centre (C). The centre point is marked in doughnut charts.

Tick marks (T). A 'tick mark' (short black line) at radial direction is shown at the outer border of the pie or doughnut chart. We use four tick marks equally spaced along the circle at angles 0°, 90°, 180° and 270°, thus 25% intervals are visually implied.

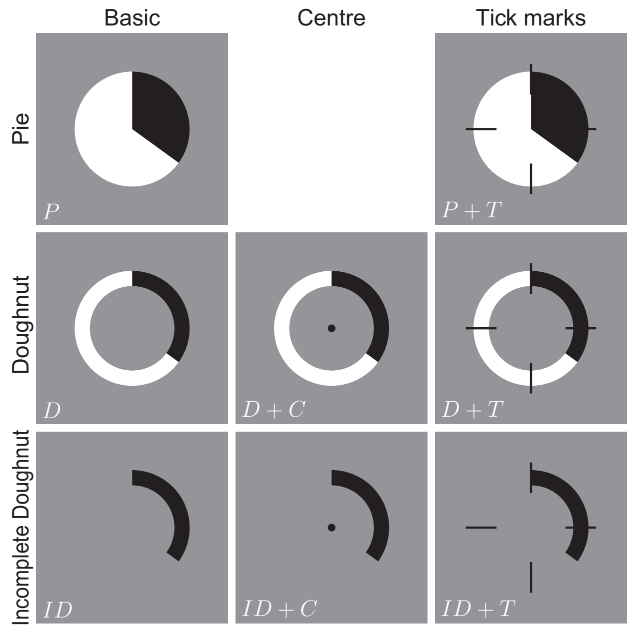

Figure 8: *Conditions in Experiment 3.*

Incomplete doughnut(ID). Show only the proportion to be estimated without showing the whole doughnut ring.

5.1. Design and procedure

Apparatus: The same setup was used as in Experiments 1 and 2.

Participants: A total of 32 participants (seven female) participated in Experiment 3. Their age ranged from 20 to 27 years (mean 23.7) and all reported normal or corrected-to-normal vision. Twelve of them were undergraduate students and the rest were postgraduate students; all participants studied Computer Science. Twenty participants reported to have no more than 1-year experience of using pie or doughnut charts and eight participants reported to have more than 3 years experience. Each of the participants received 10 Yuan (Chinese RMB) as reward.

Task: For baseline comparisons, we kept the traditional pie and doughnut chart (P, D) and added two additional visual cues (C, T). In addition, we considered incomplete doughnut charts (ID) and the effect of the additional visual cues on them. We chose the inner radius for the doughnut and incomplete doughnut charts as 73.5% of the outer radius. In total, we used eight different designs of charts, as seen in Figure 8:

P: pie chart;
$P+T$: pie chart with tick marks;
D: doughnut chart;
$D+C$: doughnut chart with centre;
$D+T$: doughnut chart with tick marks.
ID: incomplete doughnut chart;
$ID+C$: incomplete doughnut chart with centre;
$ID+T$: incomplete doughnut chart with tick marks .

Table 6: *Mean absolute error and 95% confidence interval (CI) for the eight conditions in Experiment 3.*

Condition	Mean	95% CI
P	1.8021	[1.5948, 2.0094]
$P + T$	1.4674	[1.3034, 1.6313]
D	2.0235	[1.8063, 2.2407]
$D + C$	1.7513	[1.5992, 1.9034]
$D + T$	1.4935	[1.3459, 1.6410]
ID	2.2924	[2.0361, 2.5488]
$ID + C$	2.0859	[1.8782, 2.2937]
$ID + T$	1.6509	[1.4599, 1.8419]

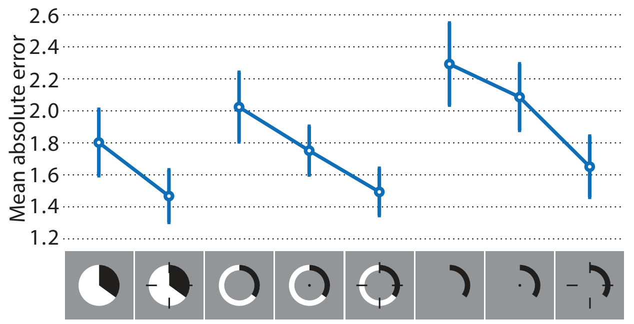

Figure 9: *Mean absolute error and 95% confidence interval (CI) for the eight conditions in Experiment 3. Note that even though the slice to estimate starts exactly at a tick mark in these pictures, both the starting point and the size of the slices were randomized in the experiment.*

We did not consider the combination $P + C$ since the centre can be accurately located in pie charts without any additional cues. For each of these eight conditions, we considered 12 proportion conditions, following the same setting as in Experiment 1, see Section 3.1. Overall, Experiment 3 consisted of

$$8 \text{ (conditions)} \times 12 \text{ (proportions)} \times 32 \text{ (participants)} = 3072 \text{ trials.}$$

Procedure: The procedure was the same as in Experiment 1 except that there were eight practice tests for training and the formal experiment for each participant was divided into four blocks each one containing 24 trials. For each test, we recorded the participant ID, the design condition, the proportion condition, the estimate given by the participant and the time in milliseconds from the moment the test appeared on screen to the moment the participant confirmed their answer. The typical time for the whole experiment, including the training session and the questionnaire, was about 30 min.

5.2. Results

Cleaning the data in the same way as in Experiments 1 and 2 resulted in removing nine trials (of 3072) from the results.

The mean absolute errors and 95% confidence intervals for each of the eight conditions are shown in Table 6 and Figure 9. The conclusion is that adding visual cues improves the accuracy of the corresponding 'base' charts (pie, incomplete doughnutand doughnut). Moreover, we observe that the effect of tick marks is more important than the effect of marking the centre for the doughnut (either complete or incomplete). In particular, note that for the pairs $(P, P + T)$, $(D, D + T)$ and $(ID, ID + T)$ the mean absolute error for each pair element lies outside the 95% confidence interval of the other pair element. In the case of the pie chart where the centre can be accurately located we observe that the tick marks also improve accuracy. Finally, note that making the doughnut incomplete is indeed an anti-cue: the incomplete doughnut has the lowest accuracy among all charts we compared. Nevertheless, and rather surprisingly, the incomplete doughnut chart with tick marks is the third most accurate chart with only the pie chart with tick marks and the doughnut chart with tick marks giving better results.

In post-survey questionnaires, 25 of 32 participants reported similar strategies as those in Experiments 1 and 2. The participants were asked to give their subjective opinion about which designs they thought were helpful in making accurate estimates; participants were not asked to rank the charts and they could identify multiple charts as being accurate. A total of 26 participants chose $D + T$, 21 chose $P + T$, 6 chose $ID + T$, 3 chose $D + C$ and 2 chose P. Note that more participants found $D + T$ more accurate than $P + T$ although the results tell a slightly different story. Besides, only two participants found that the plain pie chart P is accurate. This implies that, from the participants' point of view, adding tick marks is such a drastic improvement that charts without them cannot be considered accurate. In a similar question about which types of charts they preferred, 26 participants chose $D + T$, 22 chose $P + T$, 12 chose P, 10 chose $ID + T$, 4 chose $D + C$ and 3 chose D. Note that the plain pie chart is the third most preferred type of chart. We speculate that this is because of participants' more extensive familiarity with pie charts.

Of those participants who reported their estimation strategy, four reported that they would always try to estimate the corresponding angle of the ring sector, two reported that they would use the marked centre and four reported that they would leverage the tick marks.

5.3. Discussion

The main result of Experiment 3 is that tick marks consistently and significantly improve the accuracy of proportion estimates while marking the centre also improves accuracy, albeit not as drastically. Therefore, we suggest the use of tick marks for pie and doughnut charts when it is important to improve the accuracy of proportion estimates without using text annotations.

The finding that the incomplete doughnut chart gives the worst accuracy is not surprising and it signifies the importance of visual cues from a different point of view: taking away familiar visual cues results in a drop in accuracy. Nevertheless, what is surprising is that replacing such familiar visual cues with other ones (in this case, with the centre or the tick marks) can compensate to the extent that the incomplete doughnut with tick marks is at the same level of accuracy as the traditional pie chart.

A natural question is whether the tick marks improve the accuracy uniformly or only near the corresponding proportions 0%, 25% and

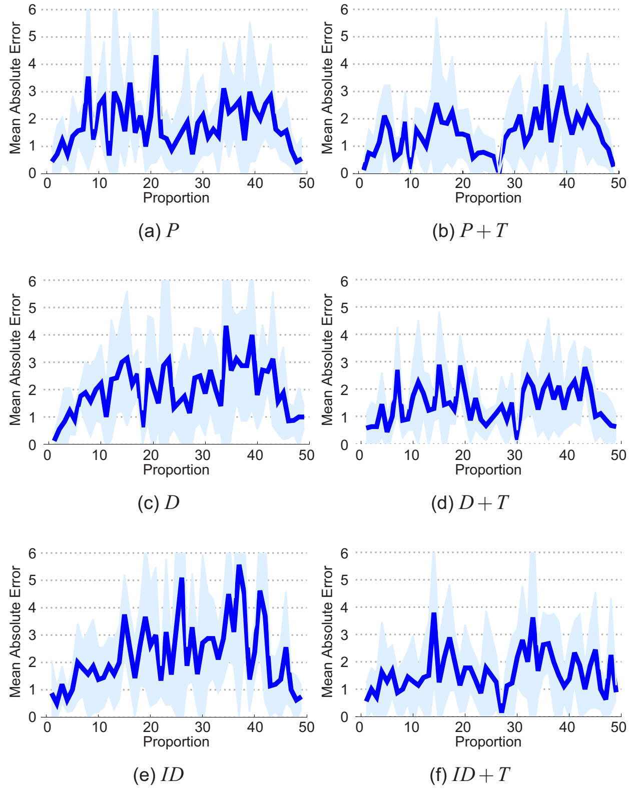

Figure 10: *Comparison of the mean absolute error for each proportion without (left) and with (right) tick marks for the three chart types in Experiment 3. The light blue band represents the corresponding 95% confidence interval.*

50%. Although we did not focus on this question in our experimental design one can check the mean absolute error for each proportion in conditions with and without tick marks. The results are presented in Figure 10 where we compare the accuracy of each of the three chart types without and with tick marks. For all chart types, we observe that the addition of the tick marks improves the accuracy around 25% while around 0% and 50% accuracy is already so good that tick marks have no clear effect. Moreover, tick marks tend to reduce errors in ranges where these are larger than the average.

We can reasonably assume that the accuracy of proportion estimates would further improve if we added more tick marks. Nevertheless, adding more tick marks (for example, every 10%) would either require that we explicitly provide information about the distance between tick marks or that readers count the number of tick marks and infer their distance. The choice of having only four tick marks at 25% distance has the benefit that the reader can easily infer this information without explicitly providing it. We would not be surprised if further studies show that this choice strikes a good balance between accuracy and ease of reading.

We also report how the additional visual cues affected the response time. Nevertheless, it should be taken into account that there was no time limit in our experimental design and we acknowledge that imposing a time limit, or even mentioning to participants that time is a significant metric, might have produced different results.

Table 7: *Mean response time and 95% confidence interval (CI) for the eight conditions in Experiment 3.*

Condition	Mean	95% CI
P	6.499	[5.938, 6.797]
$P+T$	7.880	[7.318, 7.927]
D	6.465	[5.903, 6.731]
$D+C$	6.106	[5.543, 6.524]
$D+T$	7.147	[6.585, 7.477]
ID	6.131	[5.569, 6.481]
$ID+C$	6.216	[5.655, 6.640]
$ID+T$	7.006	[6.442, 7.570]

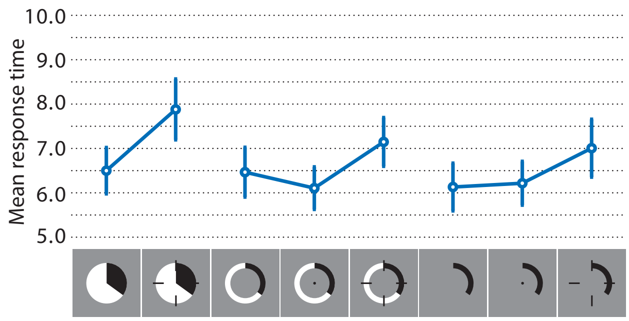

Figure 11: *Mean response time and 95% confidence interval (CI) for the eight conditions in Experiment 3.*

The mean response times and 95% confidence intervals for each of the eight conditions are shown in Table 7 and Figure 11. Surprisingly, adding tick marks increased the response times. Adding the centre did not appear to have any effect. Since no correlation (longer observation produced more accurate results) was found between accuracy and response time in the previous two experiments, we conjecture that the increase in response time is due to the additional visual cues. One possible explanation could be, that with the addition of the tick marks more participants have tried to estimate the given proportion through comparison with the tick marks and add/subtract operations thus using more time to come up with an estimate. In contrast, adding the centre in the charts did not increase the response time, since time is saved by not needing to estimate the centre position as a reference point anymore. Such explanations, naturally suggest that different strategies are in play during proportion estimations and more studies are necessary to clarify such aspects of our research.

6. Conclusions and Further Discussion

We conducted three experiments to test the effect of different design parameters to the accuracy of proportion estimates in doughnut charts. The first two experiments tested the effect of the overall size and the inner radius (the hole size). These experiments showed that the accuracy is insensitive to the overall size but is negatively affected when the inner radius increases. The third experiment studied the effectiveness of visual cues and showed that adding visual cues (such as adding tick marks) improves accuracy while removing visual cues (such as using an incomplete chart) reduces accuracy.

Based on these findings, we can make the following suggestions for the use of pie and doughnut charts:

(1) It is not necessary to try to show such charts in very large size. Experiment 1 showed that in the tested range (diameter from 3.20 to 8.88 cm) the accuracy does not significantly change. Nevertheless, this does not mean that making the chart even smaller will not affect its accuracy for proportion estimation tasks. We can reasonably expect that there is a minimal size $D_{\min}$ such that for sizes smaller than $D_{\min}$ the accuracy starts decreasing. It would be interesting to study this question and determine a value for $D_{\min}$ if the latter indeed exists.
(2) Doughnut charts are as good as pie charts and, in general, the inner radius does not have a significant effect on the proportion estimation accuracy of doughnut charts, except when the charts become extremely thin. If there is a need to leverage the inner space of doughnut charts, it is not necessary to keep the doughnut charts very thick. However, we suggest avoiding the use of an extremely thin ring for representing proportions.
(3) If aesthetically acceptable, add tick marks. We speculate that adding more visual cues will further improve the accuracy for proportion estimation tasks although we do expect a situation of diminishing returns. In the future, we would like to study the combined effect of more than one visual cues and their effect when the number of proportions to be estimated increases.

The present work has focused on static charts. Nevertheless, mobile devices allow interaction and such possibility can significantly enrich the representation of proportions. As the most simple example, consider a pie or doughnut chart where touching a slice shows an accurate value for the corresponding proportion on screen. Therefore, in such context many new questions take precedence. What are the most effective ways to interact with such charts, especially taking into account the limited size of the interaction surface? Should we focus on proportion accuracy or on different measures of the effectiveness of such charts? What is the role of animation?

The paper focuses on studying the effect of specific parameters of doughnut charts. At the moment, a more general understanding of doughnut charts and the role they can play in data visualization tasks is missing. More work is necessary in that direction and we hope that the results we present in this paper will contribute in such studies by offering a basis for understanding how to effectively employ doughnut charts and for guiding the design of future experiments that answer more general questions.

Ultimately, our paper studies the effect of fundamental design parameters to the task of proportion estimates in doughnut charts and gives specific suggestions for the use of pie and doughnut chart. As such charts are very commonly used to represent proportions, we hope that our findings and suggestions will be applicable to real-world applications of doughnut charts.

Acknowledgments

The work was supported by National 973 Program of China (2015CB352503), NSFC (61761136020, 61502416, 61502132), NSFC-Zhejiang Joint Fund for the Integration of Industrialization and Informatization (U1609217), Zhejiang Provincial Natural Science Foundation (LR18F020001, LQ14E050010) and the 100 Talents Program of Zhejiang University. This project is also partially funded by Microsoft Research Asia. We thank Petra Isenberg from Inria and Steve Haroz from Northwestern University for helpful discussions.

References

[BSBE17] BLASCHECK T., SCHWEIZER M., BECK F., ERTL T.: Visual comparison of eye movement patterns. *Computer Graphics Forum 36*, 3 (2017), 87–97.

[CCM*14] CHEN H., CHEN W., MEI H., LIU Z., ZHOU K., CHEN W., GU W., MA K.-L.: Visual abstraction and exploration of multi-class scatterplots. *IEEE Transactions on Visualization and Computer Graphics 20*, 12 (2014), 1683–1692.

[CGS*11] CAO N., GOTZ D., SUN J., LIN Y.-R., QU H.: Solarmap: Multifaceted visual analytics for topic exploration. In *Proceedings of IEEE 11th International Conference on Data Mining* (Vancouver, BC, Canada, 2011), pp. 101–110.

[Cle94] CLEVELAND W. S.: *The Elements of Graphing Data*. Hobart Press, Summit, NJ, 1994.

[CLS*12] CAO N., LIN Y.-R., SUN X., LAZER D., LIU S., QU H.: Whisper: Tracing the spatiotemporal process of information diffusion in real time. *IEEE Transactions on Visualization and Computer Graphics 18*, 12 (2012), 2649–2658.

[CM84] CLEVELAND W. S., MCGILL R.: Graphical perception: Theory, experimentation, and application to the development of graphical methods. *Journal of the American Statistical Association 79*, 387 (1984), 531–554.

[CS27] CROXTON F. E., STRYKER R. E.: Bar charts versus circle diagrams. *Journal of the American Statistical Association 22*, 160 (1927), 473–482.

[CS32] CROXTON F. E., STEIN H.: Graphic comparisons by bars, squares, circles, and cubes. *Journal of the American Statistical Association 27*, 177 (1932), 54–60.

[CWS*17] CYPKO M. A., WOJDZIAK J., STOEHR M., KIRCHNER B., PREIM B., DIETZ A., LEMKE H. U., OELTZE-JAFRA S.: Visual verification of cancer staging for therapy decision support. *Computer Graphics Forum 36*, 3 (2017), 109–120.

[CYW*16] CHEN S., YUAN X., WANG Z., GUO C., LIANG J., WANG Z., ZHANG X. L., ZHANG J.: Interactive visual discovering of movement patterns from sparsely sampled geo-tagged social media data. *IEEE Transactions on Visualization and Computer Graphics 22*, 1 (2016), 270–279.

[CZC*15] Chen H., Zhang S., Chen W., Mei H., Zhang J., Mercer A., Liang R., Qu H.: Uncertainty-aware multidimensional ensemble data visualization and exploration. *IEEE Transactions on Visualization and Computer Graphics 21*, 9 (2015), 1072–1086.

[Eel26] Eells W. C.: The relative merits of circles and bars for representing component parts. *Journal of the American Statistical Association 21*, 154 (1926), 119–132.

[FCB01] Friel S. N., Curcio F. R., Bright G. W.: Making sense of graphs: Critical factors influencing comprehension and instructional implications. *Journal for Research in Mathematics Education 32*, 2 (2001), 124–158.

[FIB*14] Fuchs J., Isenberg P., Bezerianos A., Fischer F., Bertini E.: The influence of contour on similarity perception of star glyphs. *IEEE Transactions on Visualization and Computer Graphics 20*, 12 (2014), 2251–2260.

[GND*18] Glueck M., Naeini M. P., Doshi-Velez F., Chevalier F., Khan A., Wigdor D., Brudno M.: PhenoLines: Phenotype comparison visualizations for disease subtyping via topic models. *IEEE Transactions on Visualization and Computer Graphics 24*, 1 (2018), 371–381.

[HB10] Heer J., Bostock M.: Crowdsourcing graphical perception: Using mechanical turk to assess visualization design. In *Proceedings of the ACM Conference on Human Factors in Computing Systems* (Atlanta, GA, USA, 2010), pp. 203–212.

[HPvU*18] Höllt T., Pezzotti N., van Unen V., Koning F., Lelieveldt B. P. F., Vilanova A.: CyteGuide: Visual guidance for hierarchical single-cell analysis. *IEEE Transactions on Visualization and Computer Graphics 24*, 1 (2018), 739–748.

[HS92] Hollands J. G., Spence I.: Judgments of change and proportion in graphical perception. *Human Factors 34*, 3 (1992), 313–334.

[HS98] Hollands J. G., Spence I.: Judging proportion with graphs: The summation model. *Applied Cognitive Psychology 12*, 2 (1998), 173–190.

[Kos06] Kosslyn S.: *Graph Design for the Eye and Mind*. Oxford University Press, New York, 2006. https://books.google.nl/books?id=dNe7GktaOF4C.

[KSB*09] Krzywinski M., Schein J., Birol I., Connors J., Gascoyne R., Horsman D., Jones S. J., Marra M. A.: Circos: An information aesthetic for comparative genomics. *Genome Research 19*, 9 (2009), 1639–1645.

[KTB*18] Kumpf A., Tost B., Baumgart M., Riemer M., Westermann R., Rautenhaus M.: Visualizing confidence in cluster-based ensemble weather forecast analyses. *IEEE Transactions on Visualization and Computer Graphics 24*, 1 (2018), 109–119.

[KZ10] Kosara R., Ziemkiewicz C.: Do Mechanical Turks dream of square pie charts? In *Proceedings of the BELIV Workshop: BEyond time and errors: novel evaLuation methods for Information Visualization* (Atlanta, GA, USA, 2010), pp. 63–70.

[LCW*15] Liu S., Chen Y., Wei H., Yang J., Zhou K., Drucker S. M.: Exploring topical lead-lag across corpora. *IEEE Transactions on Knowledge and Data Engineering 27*, 1 (2015), 115–129.

[LCWL14] Liu S., Cui W., Wu Y., Liu M.: A survey on information visualization: Recent advances and challenges. *Visual Computer 30*, 12 (2014), 1373–1393.

[LWL*17] Liu D., Weng D., Li Y., Bao J., Zheng Y., Qu H., Wu Y.: SmartAdP: Visual analytics of large-scale taxi trajectories for selecting billboard locations. *IEEE Transactions on Visualization and Computer Graphics 23*, 1 (2017), 1–10.

[MMWC17] Mei H., Ma Y., Wei Y., Chen W.: The design space of construction tools for information visualization: A survey. *Journal of Visual Languages & Computing* (2017).

[Rob05] Robbins N. B.: *Creating More Effective Graphs*. Wiley, Hoboken, NJ, 2005. https://books.google.nl/books?id=aKnuAAAAMAAJ.

[SH87] Simkin D., Hastie R.: An information-processing analysis of graph perception. *Journal of the American Statistical Association 82*, 398 (1987), 454–465.

[SHK15] Skau D., Harrison L., Kosara R.: An evaluation of the impact of visual embellishments in bar charts. *Computer Graphics Forum 34*, 3 (2015), 221–230.

[Sii14] Siirtola H.: Bars, pies, doughnuts & tables: Visualization of proportions. In *Proceedings of the 28th International BCS Human Computer Interaction Conference on HCI* (Southport, UK, 2014), pp. 240–245.

[SK16] Skau D., Kosara R.: Arcs, angles, or areas: Individual data encodings in pie and donut charts. *Computer Graphics Forum 35*, 3 (2016), 121–130.

[SL91] Spence I., Lewandowsky S.: Displaying proportions and percentages. *Applied Cognitive Psychology 5*, 1 (1991), 61–77.

[Spe90] Spence I.: Visual psychophysics of simple graphical elements. *Journal of Experimental Psychology: Human Perception and Performance 16*, 4 (1990), 683–692.

[Spe05] Spence I.: No humble pie: The origins and usage of a statistical chart. *Journal of Educational and Behavioral Statistics 30*, 4 (2005), 353–368.

[SWLL13] Sun G.-D., Wu Y.-C., Liang R.-H., Liu S.-X.: A survey of visual analytics techniques and applications: State-of-the-art research and future challenges. *Journal of Computer Science and Technology 28*, 5 (2013), 852–867.

[SZ00] Stasko J., Zhang E.: Focus+ context display and navigation techniques for enhancing radial, space-filling hierarchy visualizations. In *Proceedings of IEEE Symposium on Information Visualization* (Salt Lake City, UT, USA, 2000), pp. 57–65.

[TK75] Tversky A., Kahneman D.: Judgment under uncertainty: Heuristics and biases. In *Utility, Probability, and Human Decision Making*. Springer, Dordrecht (1975), pp. 141–162.

[Tuf01] TUFTE E. R.: *The Visual Display of Quantitative Information* (2nd edition). Graphics Press, Cheshire, CT, 2001.

[VH27] VON HUHN R.: Further studies in the graphic use of circles and bars: I: A discussion of the Eells' experiment. *Journal of the American Statistical Association 22*, 157 (1927), 31–36.

[WLL*16] WANG X., LIU S., LIU J., CHEN J., ZHU J., GUO B.: Topicpanorama: A full picture of relevant topics. *IEEE Transactions on Visualization and Computer Graphics 22*, 12 (2016), 2508–2521.

[WLS*18] WU Y., LAN J., SHU X., JI C., ZHAO K., WANG J., ZHANG H.: iTTVis: Interactive visualization of table tennis data. *IEEE Transactions on Visualization and Computer Graphics 24*, 1 (2018), 709–718.

[WWL*10] WU Y., WEI F., LIU S., AU N., CUI W., ZHOU H., QU H.: OpinionSeer: Interactive visualization of hotel customer feedback. *IEEE Transactions on Visualization and Computer Graphics 16*, 6 (2010), 1109–1118.

[WZM*16] WANG X.-M., ZHANG T.-Y., MA Y.-X., XIA J., CHEN W.: A survey of visual analytic pipelines. *Journal of Computer Science and Technology 31*, 4 (2016), 787–804.

[XCH*14] XIE C., CHEN W., HUANG X., HU Y., BARLOWE S., YANG J.: VAET: A visual analytics approach for e-transactions time-series. *IEEE Transactions on Visualization and Computer Graphics 20*, 12 (2014), 1743–1752.

[YYQ17] YANG Y., YAO Q., QU H.: Vistopic: A visual analytics system for making sense of large document collections using hierarchical topic modeling. *Visual Informatics 1*, 1 (2017), 40–47.

[ZK10a] ZIEMKIEWICZ C., KOSARA R.: Beyond Bertin: Seeing the forest despite the trees. *IEEE Computer Graphics and Applications 30*, 5 (2010), 7–11.

[ZK10b] ZIEMKIEWICZ C., KOSARA R.: Implied dynamics in information visualization. In *Proceedings of the ACM International Conference on Advanced Visual Interfaces* (Roma, Italy, 2010), pp. 215–222.