Dissecting Pie Charts

Harri Siirtola, Kari-Jouko Räihä, Howell Istance, Oleg Špakov

► To cite this version:

Harri Siirtola, Kari-Jouko Räihä, Howell Istance, Oleg Špakov. Dissecting Pie Charts. 17th IFIP Conference on Human-Computer Interaction (INTERACT), Sep 2019, Paphos, Cyprus. pp.688-698, 10.1007/978-3-030-29384-0_41 . hal-02544607

HAL Id: hal-02544607

<https://hal.inria.fr/hal-02544607>

Submitted on 16 Apr 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Distributed under a Creative Commons Attribution| 4.0 International License

Dissecting Pie Charts

Harri Siirtola^a, Kari-Jouko Räihä, Howell Istance, and Oleg Špakov

Tampere University, Faculty of Information Technology and Communication, Tampere, Finland

ABSTRACT

Pie charts can be regularly found both in the popular media and research publications. There is evidence that other forms of visualizations make it easier to evaluate the relative order of the data items. Doughnut charts have been suggested as a variation that has advantages over pie charts. We investigated how pie charts and doughnut charts are affected by the number of sectors in the chart, the difference in sector sizes, and the size of the hole. We carried out an eye tracking study to find out how the charts are read. Our study reveals the distribution of visual attention for each type of chart. The results indicate that doughnut charts with a medium size hole have a slight edge over the other chart types we studied. We also show that contrary to common claims, for information extraction also the area and length of sector arc are used in addition to the angles of the sectors.

KEYWORDS

pie charts; doughnut charts; readability; effectiveness

1. Introduction

Pie charts are omnipresent. Whenever there is an election to be reported, a budget to be explained, or a poll to be published we see them, and usually many of them. Pie charts are practically the *de-facto* standard to represent how some part relates to a whole, or to other parts of the (same) whole. Perhaps the greatest strength of pie charts is that they are self-explanatory – or at least supposed to be so.

The pie chart is also one of the most controversial data graphic representations ever. Many prominent experts advice to avoid them completely because human visual system is better in perceiving length than angle (*‘The only thing worse than a pie chart is several of them’* (Tufte 1983), *‘Save pies for the dessert’* (Few 2007), *‘Pie charts are bad’* (Fenton 2009), *‘Death to pie charts’* (Nussbaumer 2011)), but there are advocates as well (*‘Why Tufte is flat-out wrong about pie charts’* (Gabrielle 2013), *‘In defense of pie charts’* (Kosara 2011), and also Spence and Lewandowsky (1991); Peck et al. (2013)). We are not trying to solve this debate or take sides – our approach is pragmatic: since pie charts are used anyway, we need to understand them better.

One possible explanation for the tumult over pie charts is that they are often mis-used. The most important guidelines with pie charts are that there should be no more than about 7 slices, slices should not be exploded (taken out from the pie), there should be no separate legend (requiring movement between slices and labels), and there should be no 3D effects (a generally bad idea in data graphics). Making a Google image search for “pie chart” shows that these rules are often recklessly violated.

We study in this paper how pie charts are read by reporting an experiment where

participants were requested to list the sectors of pie charts in a decreasing order of size while their gaze was recorded with an eye tracker. We also include the most popular variation of pie chart, the doughnut chart, and investigate how the size of the hole in the middle of the doughnut affects the reading. The chosen experimental task is perhaps the most common operation with pie and doughnut charts as they are really geared towards relative comparisons. Sometimes the slices of these charts can be freely sorted by magnitude which would render the task trivial, but often the variables represented by slices have a natural order which precludes this (e.g. political parties, ratings about something).

2. Previous work

Lima (2018) presents two explanations for the popularity of pie charts: historical and evolutionary ones. His examples suggest that the preference for round pie charts might be too deep-rooted to be removed with any kind of reasoning.

Eells (1926) did the first comparisons between pie charts and stacked bar charts. This was a justified experimental setup – both have parts as well as a whole. His experimental task was to estimate the percentage of the whole represented by an element, as a number. He concluded that pie charts could be read as rapidly as stacked bar charts, and that the accuracy of pie charts was better. As a response, von Huhn (1927) criticized Eells’s experimental task, because neither of the visualizations is meant to be used for absolute judgements, but relative ones. In addition, von Huhn suspected that the missing scales and labels from the stimulus ruined the ecological validity and the results. Despite von Huhn’s criticism, the same experimental setup was later repeated in a number of studies. A recent comparison of pie charts and bar charts (with relative judgments) is presented by Siirtola (2019).

According to a seminal paper by Cleveland and McGill (1984), the accuracy of judgements (from the most accurate to least accurate) in elementary perceptual tasks is length, angle, and area. However, this result is for extracting quantitative information from graphs (absolute judgement), and in our experiment only relative judgement is needed. In addition, the judgement based on the curved lengths in our experiment might be more difficult than on straight ones. The study was later replicated using Mechanical Turk by Heer and Bostock (2010) with similar results.

Spence and Lewandowsky (1991) give a comprehensive review of the earlier research on pie charts, and point out that the practitioners of display graphics keep using pie charts despite the harsh criticism from experts – and the practitioners probably have a good sense of what works and what doesn’t work.

Skau and Kosara (2016) did a comparison of pie charts, doughnut charts, and ‘angle only charts’ (two line segments depicting an angle), and argued that angle is not the primary or only factor when pie charts are read. They did not use eye tracking, but deconstructed pie and doughnut charts into their constituent parts. Their experimental task was absolute judgement (“*What percentage of the whole is indicated?*”) which we find unnatural in part-whole visualizations.

Although there thus is a wealth of research on pie charts and doughnut charts, it is targeted at their performance: how fast and how correctly can the information represented by the chart be extracted. There are no studies on how this extraction process takes place, i.e., where in the chart do viewers attend to. This can only be found out by tracking the gaze of the viewers. Ours is the first eye tracking study on this issue. Our goal is to shed light on *how* pie charts and doughnut charts are read.

Figure 1.: Variations of stimulus, from left to right: Pie Chart (no hole), Doughnut-25 (a doughnut chart with hole having radius of 25% of the pie radius), and correspondingly, Doughnut-50 and Doughnut-75 variations.

3. Method

3.1. Participants

In our experiment participants were recruited from an introductory course in human-computer interaction, where they received course credit for participating. 29 students volunteered to take part in the test. Reliable gaze data could not be collected for two of them: for one the tracker could not be calibrated, and for the other there were big gaps in the gaze point stream produced by the eye tracker. Thus 27 participants (17 male, 10 female) calibrated well and produced data that is reported in this paper.

The age of the participants ranged from 19 to 56 years, with median age of 23 years. All had normal vision or corrected to normal vision (7 wore eye glasses and 2 had contact lenses). Only one participant had previously used an eye tracker. Pie charts were previously familiar to all participants, but doughnut charts were equally familiar to only five participants, and somewhat familiar to another five participants.

3.2. Apparatus

A Tobii T60 eye tracker with a 17-inch TFT color monitor with 1280×1024 resolution was used to track the gaze. A PC running Windows 10 was used for the experiment. The stimuli were presented using the Tobii Pro Lab software.

3.3. Task

The participants were shown a sequence of pie charts in random order. The charts varied based on the number of segments (4, 5, 6 or 7). They also were of varying difficulty, with the difference between the value of the segments being depicted at least 6%, 10%, 14% or 18%. Finally, the radius of the hole was varied from 0% of the doughnut radius (corresponding to a full pie), to 25%, 50%, and finally 75%, corresponding to the slimmest doughnut (Figure 1). Altogether there were thus $4 \text{ (number of segments)} \times 4 \text{ (angle difference)} \times 4 \text{ (hole size)} = 64$ different charts.

The pies were centered on the screen and had a radius of 356 pixels. The sectors of the charts were labelled outside the perimeter of the pie with capital letters starting from A for the top right sector, and running clockwise from there on. The participants were asked to say aloud the order of the sectors from the biggest to the smallest by stating the labels of the sectors in that order.

3.4. Procedure

Upon entering the lab the participants first signed an informed consent form.

The experimenter then explained the task and showed on paper some sample images of pie charts and doughnut charts. The participant was told to speak clearly, and explained that they could revise their judgment of the order of the sectors during the presentation of a stimulus, as long as the order they eventually chose was clear.

They were asked to work quickly and accurately. As a motivation, the five best (using a combined measure of speed and correctness) were promised a monetary reward of 10 euros. The details of the metric used for ordering the performance were not revealed.

The participants were then seated in front of the eye tracker at a distance of about 60 cm from the screen. The eye tracker was calibrated using a 5-point calibration. The quality of the calibration was measured after the calibration. Both the accuracy and precision were less than 0.5 degrees, on average, and always at most 1 degree.

After calibration the experimenter started an audio recording using a Samsung A3 mobile phone and moved to another computer for entering the orders of sectors that the participant uttered. After the experiment the audio recording was used to double check that the experimenter had transcribed the participant's answers correctly.

The data collection then began. The participant advanced to the next chart by pressing the space bar. A dot with a 10 pixel radius was first shown in the center of the screen for 2.5 seconds. The next chart then appeared automatically. After uttering the order of the sectors the participant pressed the space bar again to move to the next chart, and the process was repeated with the dot appearing in the center.

Since viewing 64 charts in a row is a monotonous task, the participant was given information on progress. After the first four charts, instead of the picture with a dot, a circle containing the number 60 was shown for five seconds in the center to indicate that 60 charts still remained. Similar information was then given after every 10 graphs.

After finishing the task the participant was interviewed. Finally they were shown live visualizations of their gaze path when viewing some of the charts.

4. Results

4.1. Time and correctness

Two factors affect the performance of pie charts and doughnut charts inversely: number of sectors and difference of values depicted. The fewer sectors there are, the easier the task, and the smaller the difference, the more difficult it becomes.

For visualizing the distribution of the data points it is useful to define a metric that combines the effect of the two factors. We define an index of difficulty as

$$IOD = \ln(5 \times (\text{Number of Sectors})/(\text{Difference as Percentage})).$$

This function distributes the stimuli to cover the whole stimulus space so that all cases have a positive *IOD* value. The easiest case is one with 4 sectors and value difference of 18%, producing an *IOD* of 0.105. The most difficult case has 7 segments with difference of 6%, producing an *IOD* of 1.764. In the first case the data values that are represented by the sectors in the chart are 70.9, 83.6, 96.4 and 109.1. In the latter case the data values range from 43.6 to 59.3 with a difference of 2.6 between consecutive values, which can be expected to be a very difficult ordering task.

Figure 2 (on the left) shows the *IOD* versus the mean of task time (from appearance

of stimulus to its disappearance when the participant pressed the space bar) for each visualization type, aggregated per the levels of IOD . The overlaid curves are smoothed with loess local polynomial regression. The figure suggests that Doughnut-50 might be the fastest visualization to interpret in medium-to-difficult cases, but this difference is not statistically significant according to our mixed-effects modeling.

Figure 2.: IOD (Index of Difficulty) vs. mean task time (on the left) and number of errors (on the right).

Figure 2 also shows the IOD versus the number of errors for each visualization type (on the right). There is practically no difference between curves: the number of errors increases with the same rate in each visualization as tasks become more difficult.

Table 1 shows the mean and standard deviation for task execution time, and the error count per visualization type. Again, the numbers suggest that Doughnut-50 has a slight advantage in terms of time and errors, but there is no statistical significance.

Table 1.: Summary of task times and standard deviations, in milliseconds, and error counts.

Type	Mean of task time	SD	Error count
Doughnut-75	10,309	5,982	137
Doughnut-50	9,791	5,169	130
Doughnut-25	10,450	6,380	137
Pie Chart	10,423	6,925	142

4.2. *Distribution of visual attention*

In comparing the relative sizes of sectors in a circular visualization there are four features which can be used for comparison, illustrated in Figure 3. We can use the angle of sectors, either explicitly shown or imagined. In addition, we can use the length of sector arcs, either the outer one in case of pie chart, or inner and outer one in case of doughnut chart. We can also base our judgement on the area of sectors. In reality a combination of the features is likely to be used.

Figure 3.: What to look at when deciding the relative order of sectors, from the left: the angle, the length of inner or outer arc, and the area of sector.

In this section we use the raw gaze points and their empirical distribution to estimate the focus of visual attention. We do not cluster the gaze data into saccades and fixations but use the lowest level raw data available.

4.2.1. Gaze distance from the origo

Figure 4 shows the overall distribution of the gaze point distance from the origo in our four conditions as percentage of visualization diameter. The density distributions are clearly different (Pie Chart: $Mean = 58.1\%$, $SD = 30.5\%$, Doughnut-25: $Mean = 60.6\%$, $SD = 28.1\%$, Doughnut-50: $Mean = 68.0\%$, $SD = 23.1\%$, and Doughnut-75: $Mean = 77.8\%$, $SD = 20.1\%$).

Figure 4.: Overall density of gaze distance from the origo per visualization type.

4.2.2. Proportion of time in areas of visualization

With the empirical distribution function we can also estimate what proportion of time the participants spent in each area of the charts. The areas of interest in our study are the surroundings of the origo, area around the lower ring, area within the band (between rings), and the upper ring (Figure 3). They correspond to estimating the

magnitudes of sectors by using angles, lengths, and areas.

Figure 5 shows the overall distance of gaze points from the origo as a binned density graph (with 6 bins). The black dot denotes the median value, and the dashed lines show where the ‘ring’ of the chart resides. With all charts types the participants had to read the labels outside the ring (appr. 100 – 130% of diameter) which shows as a similar bar far right. Personal variation was high: participant P20’s median attention is almost on the inner arc when participant P23’s attention is clearly on the outer arc.

Figure 5.: Overall density of gaze distance as percentage of diameter from the origo per visualization type. The black point indicates the median of distance from the origo, and the dashed lines show the ring of the corresponding visualization. P20 and P23 were the two extreme cases.

Finally, Figure 6 is a summary how the visual attention is allocated within the four chart types. This image summarises all gaze data from all participants in 10% steps.

Figure 6.: Overall density of gaze distance from the origo per visualization type. The gaze distance has been divided into 10% bands, and each quadrant shows one visualization. The color scale is from white to red – deeper red indicates higher amount of gaze hits in that band.

5. Discussion

Figure 4 shows that the information-extracting of Pie Chart is not concentrated on the origo, i.e. comparison of angles, which has been a common assumption (e.g. by Simkin and Hastie (1987)). This was observed by Skau and Kosara (2016) as well. The figure suggests that participants use angle, area, and arc length almost evenly. Another interesting point is the comparison of Pie Chart and Doughnut-25 – the densities are similar except in the vicinity of the origo. The use of angle for comparisons is reduced because of the hole, and even more so in case of Doughnut-50 and Doughnut-75.

Figure 5 shows the overall differences in the allocation of visual attention more vividly. There are several trends in visual attention as the hole in the visualization increases: the use of the angle decreases, the use of inner and outer arcs increases, and the median attention moves towards the inner edge of the visualization. It seems that participants prefer the inner arc over the outer arc for comparisons.

Figure 6 shows the distribution of visual attention as co-centric 10% circles for all participants and conditions. It is easy to see how the hole in the middle changes the attention allocation – it is easier to use the area of a sector or length of the arc to estimate the size. For Pie Chart the area of a sector appears to be the dominant method to compare size, not the angle. Overall, Pie Chart is the most evenly-allocated visualization type, and participants used all three methods for size comparisons. For doughnuts the hole decreases the use of angle for estimation.

The participants were also interviewed about their preferences and observations. They provided comments, e.g. P10: *“If there’s a hole, then the inner arcs are closer to each other than the outer arcs, so it is easier to compare them. . . . And the angle then, it had to be like divided in four sectors and then one could really use the angles.”*

Finally, it is important to be clear about the scope of this study. We have focused only on the graphical side of the pie and doughnut charts, and removed aspects that are essential parts of properly constructed charts. For the graphical side, we have followed the gold standard: no more than seven sectors and start the sectors from 12 o’clock, proceed clock-wise. We have named our sectors, but it is often useful to include the sector size in the label, especially if the values are close to each other. We did not use any color in the charts, as its benefits vary among participants. Thus this study is focused only on how the graphical aspects of pie and doughnut charts are read and perceived. Further research is needed for richer forms of pie charts.

6. Conclusions

Pie charts are one of the most common types of visualizations encountered in the media, so it is important to understand how readers extract information from them. We show that pie charts are not only used for comparing angles. Instead, the area and the outer arc are used as well in making judgments of the relative order of the sectors. This contradicts the claim made in the literature (e.g. Simkin and Hastie (1987)).

Including a hole in the center of the diagram, i.e. using a doughnut instead of a pie, might seem like a step in the wrong direction, as it makes the angles of the sectors less prominent. However, with a suitable size of the hole the advantages may overcome this disadvantage. Concerning the time used for the judgments and the correctness of the judgments (Figure 2 and Table 1), there is a trend (but not statistically significant) that a hole that extends halfway through the radius of the diagram makes judgments slightly faster and less error prone than the other variations, including the standard pie chart.

Acknowledgements

This research was funded by the Academy of Finland, project *Private and Shared Gaze: Enablers, Applications, Experiences* (GaSP, grant number 2501287895). We acknowledge the use of the *Statistical System R* (R Core Team 2019) and *tidyverse* packages (Wickham 2017; Wickham and Grolemund 2016), especially *ggplot2* (Wickham 2010).

References

- Bar, M. and Neta, M. (2006). Humans prefer curved visual objects. *Psychological Science*, 17(8):645–648.
- Cleveland, W. S. and McGill, R. (1984). Graphical perception: Theory, experimentation, and application to the development of graphical methods. *Journal of the American Statistical Association*, 79(387):531–554.
- Eells, W. C. (1926). The relative merits of circles and bars for representing component parts. *Journal of the American Statistical Association*, 21(154):pp. 119–132.
- Fenton, S. (2009). Pie charts are bad. Personal journal on the web. <https://www.stevefenton.co.uk/2009/04/pie-charts-are-bad/>, visited 2019-01-10.
- Few, S. (2007). Save the pies for dessert. Visual Business Intelligence Newsletter. <http://www.perceptualedge.com/articles/08-21-07.pdf>, visited 2019-01-10.
- Gabrielle, B. (2013). Why Tufte is flat-out wrong about pie charts. Web log post. *Speaking PowerPoint – The new language of business*.
- Heer, J. and Bostock, M. (2010). Crowdsourcing graphical perception: using Mechanical Turk to assess visualization design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 203–212. ACM.
- Kosara, R. (2011). In defense of pie charts. Web log post. *eagereyes*. <http://eagereyes.org/criticism/in-defense-of-pie-charts>, visited 2019-01-10.
- Kosara, R. and Skau, D. (2016). Judgment error in pie chart variations. In Bertini, E., Elmqvist, N., and Wischgoll, T., editors, *Eurographics Conference on Visualization (EuroVis) 2016*.
- Lima, M. (2018). Why humans love pie charts – an historical and evolutionary perspective. Noteworthy – The Journal Blog. <https://blog.usejournal.com/why-humans-love-pie-charts-9cd346000bdc>.
- Nussbaumer, C. (2011). Death to pie charts. Web log post. *Storytelling with Data*. <http://www.storytellingwithdata.com/2011/07/death-to-pie-charts.html>, visited 2019-01-10.
- Peck, E. M. M., Yuksel, B. F., Ottley, A., Jacob, R. J., and Chang, R. (2013). Using fNIRS brain sensing to evaluate information visualization interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI ’13, pages 473–482. ACM.
- R Core Team (2019). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. <http://www.R-project.org>.
- Siirtola, H. (2019). The cost of pie charts. To appear in *23rd International Conference on Information Visualisation (IV19)*, July 2019.
- Simkin, D. and Hastie, R. (1987). An information-processing analysis of graph perception. *Journal of the American Statistical Association*, 82(398):pp. 454–465.
- Skau, D. and Kosara, R. (2016). Arcs, angles, or areas: individual data encodings in pie and donut charts. In *Computer Graphics Forum*, volume 35, pages 121–130. Wiley Online Library.
- Spence, I. and Lewandowsky, S. (1991). Displaying proportions and percentages. *Applied Cognitive Psychology*, 5:61–77.
- Tufte, E. R. (1983). *The Visual Display of Quantitative Information*. Graphics Press.
- von Huhn, R. (1927). Further studies in the graphic use of circles and bars: A discussion of the Eell’s experiment. *Journal of the American Statistical Association*, 22(157):pp. 31–39.
- Wickham, H. (2010). *ggplot2: Elegant Graphics for Data Analysis (Use R!)*. Springer.
- Wickham, H. (2017). *tidyverse: Easily Install and Load ‘Tidyverse’ Packages*. R package version 1.1.1.
- Wickham, H. and Grolemund, G. (2016). *R for Data Science: Import, Tidy, Transform, Visualize, and Model Data*. O’Reilly Media.