

CHARTER: heatmap-based multi-type chart data extraction

Joseph Shtok*, Sivan Harary*, Ophir Azulai, Adi Raz Goldfarb, Assaf Arbelle, Leonid Karlinsky**
IBM Research AI, Haifa Research Lab

ABSTRACT

The digital conversion of information stored in documents is a great source of knowledge. In contrast to the documents text, the conversion of the embedded documents graphics, such as charts and plots, has been much less explored. We present a method and a system for end-to-end conversion of document charts into machine readable tabular data format, which can be easily stored and analyzed in the digital domain. Our approach extracts and analyses charts along with their graphical elements and supporting structures such as legends, axes, titles, and captions. Our detection system is based on neural networks, trained solely on synthetic data, eliminating the limiting factor of data collection. As opposed to previous methods, which detect graphical elements using bounding-boxes, our networks feature auxiliary domain specific heatmaps prediction enabling the precise detection of pie charts, line and scatter plots which do not fit the rectangular bounding-box presumption. Qualitative and quantitative results show high robustness and precision, improving upon previous works on popular benchmarks.

KEYWORDS

charts detection, analysis and conversion

1 INTRODUCTION

Charts carry important part of the information content in many kinds of documents - financial reports, scholarly articles, presentations, to name a few. Recent interest in automatic document processing and conversion, such as in summarization and question answering tasks, has increased the importance of the extraction of underlying tabular data from chart images embedded in the converted documents. Chart analysis methods have evolved substantially in recent years from human-in-the-loop platforms relying on manual annotations [8, 15], through early data extraction algorithms [2], hybrid neural-algorithmic pipelines [7, 13], to end-to-end processing by a neural network [9, 12, 17].

Commonly a two-stage approach is used, first detecting the chart regions in the documents, and then applying some data extraction process to the detected charts. While the scope of the detection stage can be quite wide, including many types of charts [7], current tabular data extraction systems are mostly limited to the bar charts [2, 7, 9, 17], with few exceptions. One of the possible reasons is that standard object detectors, employed in recent works, better cope with (and enable easy inference from) objects like rectangular bars and text elements, less so with pie segments, while elements like line or area plots defy handling by box proposals. Additionally, matching of the chart legend to the corresponding chart elements, non-trivial due to small size of the color/texture chart component samples and the legend layout variability, remained out of scope for most of the recent works.

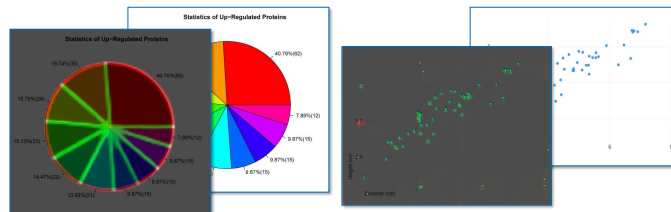


Figure 1: Heatmaps for non-rectangular graphical elements.

In this work we present a chart detection and analysis pipeline, addressing the mentioned limitations. We train a CNN object detector to produce predictions in the form of heatmaps, in addition to its standard bounding boxes proposals, thus enabling localization of graphical elements of arbitrary shape (as visualized in Figure 1). This allows us to better incorporate the geometrical domain knowledge of the chart structure into the detector during training, as well as extending the scope of chart conversion to bar, pie, line and scatter charts, all within the same (single) model. Beyond the improved detection of the graphical elements, we introduce algorithms for explicit label matching for bar and pie elements from colored legends, as well as from connector-based labels.

Our main contributions are: (1) introducing heatmap prediction for general graphical chart elements; (2) presenting novel algorithms for data extraction from pie and line charts; (3) presenting quantitative results and an ablation study of the proposed methods, showing improvement over present benchmarks.

2 RELATED WORK

Early works addressing automatic charts classification and data extraction [2, 14], used classical computer vision techniques, such as clustering image patches, Hough transform and OCR. More recently, [6, 7, 13] have presented hybrid neural-algorithmic pipelines, performing detection of the graphical objects with following extraction of numerical and textual information using OCR, Computer Vision techniques and rules; our approach belongs to this group of methods in terms of its general design. Other line of works [9, 12, 17] proposes an end-to-end analysis of the charts by a neural network. [17] develops an encoder-decoder architecture an attention mechanism for direct data extraction from bar charts by an RNN. [9] introduced a neural model for Question Answering about charts, and in particular they allow to extract tabular data by appropriate sets of questions. In [12] a standard object detector is equipped with a relation network to address the connections between the different chart elements, this model is able to produce bar heights and angles of pie segments (for single pie chart), and to match them against the legend entries. However, this approach requires an individual network design for each type of charts, and the inference process involves affine transformations of the input image for spatial alignment of features.

*authors contributed equally to this research.

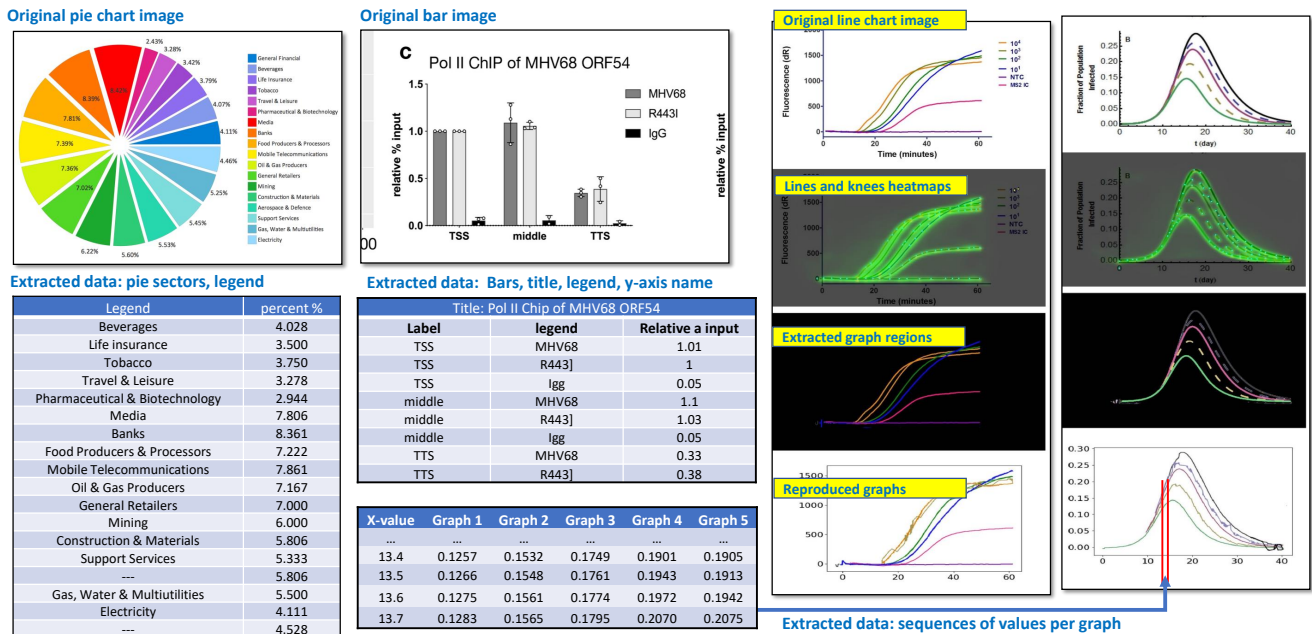


Figure 3: Examples of data extraction from different chart types.

3.4 Stage 4: Analysis module

The Analysis model builds upon recognized texts and graphical objects to discover the tabular data generating the chart. There is a dedicated logic for each type (bar, pie and line) of chart. Regrettably, we omit the details for line charts due to space limitation.

Bar charts We rely on Stage 2 detections of individual bars, while filtering them assuming equal bar width and a common y-level for one of the horizontal edges for all the bars. The bar heights are computed by first recovering the numerical y-axis (if available), via a Hough transform aligning the numerical OCR outputs on a vertical line, followed by interpolating the heights of the horizontal edges. Alternatively, we look for height values written inside or above the bars. Bar labels are read below the bars, or are retrieved from a legend by color matching; the treatment of textured legends is left for future work. The axis titles are detected by combining the detections of Stage-1 and Stage-2 models. Finally, the common chart elements, such as title and caption (found nearby by the Stage-1 detector) are added. The process was made robust by addressing the numerous special cases and issues encountered in the real world data, this robustness was then validated via results on the three datasets detailed in Section 4.

Pie charts: we use both box proposals and image-size heatmaps, produced by the CenterNet detector [18], for the detection of pie circle center, radius and the radial lines. Initial attempt to use only the sectors box proposals have proven insufficiently robust (it was difficult to filter proposals by the predicted score, especially for narrow sectors), although the circles were well fitted. Instead, we generate separate heatmaps for a circle circumference, the center point, the radial lines, and the sector corners; then the final center and pie circle are computed by initiating the centers local maxima of the center point heatmap (multiple pies are allowed), voting to a

range of radii with the circle points heatmap values (each pixel votes to its distance from the center), and update the center locations as centers-of-mass of pixels at predicted radius. Finally, we detect the radial lines (disc sectors) using the radial lines heatmap within the established circles. Using the recovered pie geometry, the sector labels are detected, if a legend is present, by matching the legend colors to those of sectors. Otherwise, we look for lines starting inside the pie disc and leading out (connected components crossing the pie edge), and extract the label using the text piece on the other end as a seed. Finally, if no connectors are present, we match texts located within/outside of the corresponding sectors.

Line charts: for line charts, we use heatmaps for the graph lines, and heatmaps for knee points in the piece-wise linear graphs the model was trained on. The individual lines are extracted by color-based clustering of the highlighted region ; non-continuous line types are supported by a stitching algorithm building the output line from available pieces. As color clustering is not always accurate, we use noise filtering and morphological operations to robustify the inference. We extract x- and y- numerical axes (in same way as for the bar charts) to reconstruct the numeric data (an example is presented in Figure 3) as (x,y) pairs.

4 RESULTS

In this section we provide the quantitative analysis of different modules and an ablation study for the heatmap-based data extraction.

4.1 Datasets for training and evaluation

For stage 1 detector, we have collected and manually annotated 743 document pages, of which 122 were crawled from the web, and 521 taken from PubMed COVID-19 related medical articles. They contain 1260 bar charts, 336 pie charts, 256 line charts, and 213

Table 1: Detection of charts and related elements in documents

Category	Bar	Pie	Line	Scatter
AP@0.5	98.0%	97.8%	84.4%	91.35%

Category	Legend	Caption	Title	X label	Y label
AP@0.5	77.8%	91.4%	71.8%	84.7%	94.9%

Stage 1 detector performance, measured as Average Precision (AP) with IoU= 0.5, averaged over 5 random train/test splits.

Table 2: Detection of graphical elements in chart image

Category	Horizontal bar	Vertical bar	Pie sector
[12]	–	80.2%	–
Ours	76.5%	90.5%	90.9%

Performance (in Average Precision (AP) with IoU= 0.5 of the graphical elements detector (box categories) on the real data.

scatter plots. Stage 2 detector was trained on a synthetic dataset, generated with FigureQA code, as detailed in Section 3.2. It consists of 30k chart images for each type (pie, vertical bar, horizontal bar, line, and scatter charts). Test data for this stage comprises of 169 real-world pie and bar charts, also partly crawled from the web, and partly taken from the COVID-19 documents. For the evaluation of the data extraction performance we have used 70 bar charts from ICPR2020 Competition data [1] and manually annotated 30 pie charts from this corpus.

4.2 Evaluation of Stages 1,2

We have trained the stage 1 detector in the k-fold fashion, using 5 random 90%/10% train/test splits of the 743 pages. The results, averaged over the 5 splits, are given in Table 1. As reflected by the AP values, the performance of the stage 1 detector is satisfactory for the rest of the system to rely on.

Stage 2 detector performance on the 169 real-world charts is given in the Table 2). It is compared to [12], where a similar evaluation was conducted (on their *Annotated* dataset).

4.3 Stage 4: extraction of source tabular data

We evaluate the quality of recovering the tabular data, following the experiment designs in [17] and [12]. The test is conducted on the real-world data of 70 bar charts and 30 pie charts from the ICPR2020 dataset, and is compared to similar Annotated datasets loc.cit. (30 charts in [17] and 10 bar charts + 10 pie charts in [12]). The accuracy of numeric information (bar heights and pie sectors angles) is measured as a portion of detected elements which numeric value is recovered within ϵ relative error: $|(val_{GT} - val_{pred})/val_{GT}| \leq \epsilon$, as in [7], for $\epsilon \in [0.01, \dots, 0.25]$. The accuracy of label reconstruction was not considered in [17], whereas in [12] TPs only elements with perfectly reproduced labels are counted; this is the reason part of the entries in Table 3 are empty. We measure accuracy of text label prediction by the ratio of Levenshtein distance L_{dist} [11], defined as $R_{Lev}(s, t) = (|s| + |t| - L_{dist}(s, t))/(|s| + |t|)$ for strings s, t . The exactly recovered labels correspond to $R_{Lev} = 1.0$, which we used as a condition for comparison to [12] results. We observe that our

Table 3: Accuracy of bar values in extraction of tabular data

Method	A.L	A.L	E.L	E.L	E.L	E.L
ϵ	0.02	0.05	0.01	0.05	0.01	0.025
[17]	67.0%	71.0%	–	–	–	–
[12]	–	–	28.4%	32.8 %	34.3 %	38.8 %
Ours	60.0%	74.2%	31.6%	55.8%	58.3%	60.3%

E.L - exact label prediction, A.L - any labels.

Table 4: Accuracy of pie segment angles in extraction of tabular data from real-world pie charts with exact labels

Method	$\epsilon = 0.01$	$\epsilon = 0.05$	$\epsilon = 0.1$	$\epsilon = 0.25$
[12]	28.9%	57.8%	60.0%	64.6%
Ours ($R_{Lev} = 1.0$)	44.9%	60.3%	61.0%	61.0%

Table 5: Ablation study of heatmaps impact for pie charts

Label accuracy	with sector boxes	with heatmaps
$R_{Lev} = 1.0$ (exact labels)	49.4%	59.7%
$R_{Lev} = 0.8$	58.2%	66.1%
$R_{Lev} = 0.4$	64.2%	71.0%
$R_{Lev} = 0.0$ (any labels)	93.3%	98.4%

numerical accuracy for bar charts is lower than that of [17] at the more strict 2% error restriction, but is better at 5% test; this may follow from the limited 512x512 resolution of the CenterNet input images. On pie charts (Table 4) our system outperforms [12] except at the low accuracy standard $\epsilon = 0.25$. Our accuracy saturates at $\epsilon = 0.1$, meaning that we don't have errors higher than 10%. This behaviour may be more suitable for real-world applications, where inaccurate numerical predictions can mislead the user systems.

4.4 Ablation study

We explore the impact made by introducing the heatmaps for detection of pie objects, which can be alternatively detected using sectors boxes. In Table 5 we present accuracy of predicted sector angles in pie charts, using only the sector box proposals or using the set of heatmaps detailed in Section 3.4. The accuracy is substantially improved (up to 10%) by using heatmaps-based computation of the graphical elements. This table also sheds some light on the distribution of label reconstruction accuracy among the reasonably accurate (5% rel. error) sectors.

5 SUMMARY AND CONCLUSION

We have presented the CHARTER - a practical system for document charts conversion to tabular data, which can extend the scope of real-world document conversion tools to include the content of the embedded document charts. The proposed use of heatmaps allows more accurate analysis of line, scatter, bar and pie charts within the same model. Moreover, our model can be easily extended to include additional chart types, which is in the scope of our future work.

REFERENCES

- [1] 2020. *No Title*. Technical Report. <https://chartinfo.github.io/index.html#task1>
- [2] Rabah A. Al-Zaidy and C. Lee Giles. 2015. Automatic extraction of data from bar charts. *Proceedings of the 8th International Conference on Knowledge Capture, K-CAP 2015* (2015). <https://doi.org/10.1145/2815833.2816956>
- [3] Jeonghun Baik, Geewook Kim, Junyeop Lee, Sungrae Park, Dongyoon Han, Sangdoon Yun, Seong Joon Oh, and Hwalsuk Lee. 2019. What is wrong with scene text recognition model comparisons? dataset and model analysis. *Proceedings of the IEEE International Conference on Computer Vision 2019-October* (2019), 4714–4722. <https://doi.org/10.1109/ICCV.2019.00481>
- [4] Youngmin Baik, Bado Lee, Dongyoon Han, Sangdoon Yun, and Hwalsuk Lee. 2019. Character region awareness for text detection. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition 2019-June* (2019), 9357–9366. <https://doi.org/10.1109/CVPR.2019.00959>
- [5] Mircea Cimpoi, Subhransu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. 2014. Describing textures in the wild. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (2014), 3606–3613. <https://doi.org/10.1109/CVPR.2014.461>
- [6] Mathieu Cliche, David Rosenberg, Dhruv Madeka, and Connie Yee. 2017. Scatter-act: Automated Extraction of Data from Scatter Plots. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 10534 LNAI, 1 (2017), 135–150. https://doi.org/10.1007/978-3-319-71249-9_19
- [7] Wenjing Dai, Meng Wang, Zhibin Niu, and Jiawan Zhang. 2018. Chart decoder: Generating textual and numeric information from chart images automatically. *Journal of Visual Languages and Computing* 48 (10 2018), 101–109. <https://doi.org/10.1016/j.jvlc.2018.08.005>
- [8] Daekyoung Jung, Wonjae Kim, Hyunjo Song, Jeong-in Hwang, Bongshin Lee, Bohyoung Kim, and Jinwook Seo. 2017. ChartSense. (2017), 6706–6717. <https://doi.org/10.1145/3025453.3025957>
- [9] Kushal Kafle, Robik Shrestha, Brian Price, Scott Cohen, and Christopher Kanan. 2020. Answering questions about data visualizations using efficient bimodal fusion. *Proceedings - 2020 IEEE Winter Conference on Applications of Computer Vision, WACV 2020* (2020), 1487–1496. <https://doi.org/10.1109/WACV45572.2020.9093494>
- [10] Samira Ebrahimi Kahou, Vincent Michalski, Adam Atkinson, Ákos Kádár, Adam Trischler, and Yoshua Bengio. 2017. Figureqa: An annotated figure dataset for visual reasoning. *arXiv* (2017), 1–20.
- [11] Vladimir Levenshtein. 1966. Binary codes capable of correcting deletions, insertions and reversals.. In *Soviet Physics Doklady*, Vol. 10. 707–710. [https://doi.org/10.1016/S0074-7742\(08\)60036-7](https://doi.org/10.1016/S0074-7742(08)60036-7)
- [12] Xiaoyi Liu, Diego Klabjan, and Patrick N. Bless. 2019. Data extraction from charts via single deep neural network. *arXiv* (2019).
- [13] Jorge Poco and Jeffrey Heer. 2017. Reverse-Engineering Visualizations: Recovering Visual Encodings from Chart Images. *Computer Graphics Forum* 36, 3 (2017), 353–363. <https://doi.org/10.1111/cgf.13193>
- [14] Manolis Savva, Nicholas Kong, Arti Chhajta, Fei Fei Li, Maneesh Agrawala, and Jeffrey Heer. 2011. ReVision: Automated classification, analysis and redesign of chart images. *UIST'11 - Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology* April (2011), 393–402. <https://doi.org/10.1145/2047196.2047247>
- [15] William R. Shadish, Isabel C.C. Brasil, David A. Illingworth, Kristen D. White, Rodolfo Galindo, Eden D. Nagler, and David M. Rindskopf. 2009. Using UnGraph to extract data from image files: Verification of reliability and validity. *Behavior Research Methods* 41, 1 (2009), 177–183. <https://doi.org/10.3758/BRM.41.1.177>
- [16] Francisco Massa Wan-Yen Lo Yuxin Wu, Alexander Kirillov and Ross Girshick. 2019. *Detectron2*. Technical Report.
- [17] Fangfang Zhou, Yong Zhao, Wenjiang Chen, Yijing Tan, Yaqi Xu, Yi Chen, Chao Liu, and Ying Zhao. 2021. Reverse-engineering bar charts using neural networks. *Journal of Visualization* 24, 2 (2021), 419–435. <https://doi.org/10.1007/s12650-020-00702-6>
- [18] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. 2019. Objects as points. *arXiv* (2019).