

An Investigation into the Scholarly Usage of Neural Network Systems Diagrams

Guy Clarke Marshall

Department of Computer Science, University of Manchester, UK
guy.marshall@postgrad.manchester.ac.uk

Abstract. Diagrams are often used to describe system architectures in Computer Science. Neural network systems are no exception. Without any framework or guidelines establishing good practices for their design, the current state is heterogeneous and evolving. Investigation of this field necessarily employs mixed methods. The research aims to improve diagrammatic representation of neural network systems by 1) deriving appropriate requirements and metrics 2) selecting appropriate representational guidelines and 3) create and implement an iterative framework for improving a set of diagram guidelines.

Keywords: Usability, Utility, Semiotics, Systems, Neural Networks

1 Background

In Artificial Intelligence, diagrams are often used to describe the arrangement of components and processes, known as system architectures. The majority of papers at conference proceedings describe a novel system architecture, and contain a diagram to describe this. At present, these diagrams are extremely heterogeneous, even within the same scholarly venue. This heterogeneity is both in terms of what is represented, and how it is represented. We hypothesise that this heterogeneity, given that the architectures are often quite similar (or architecturally correlated) is inefficient in terms of time taken and correctness, both for the creators and consumers of these diagrams.

2 Completed Work

An initial interview study has been conducted [5], to explore usage of neural network systems architecture diagrams. This showed that inconsistency in diagrams is causing problems for readers and authors. The main issues for readers are that they have to learn each other's way of representing things, whilst for authors they are designing their own visual encodings. The interview study confirmed our initial hypothesis that the current ways of using diagrams in AI systems have problems. Unexpectedly, we also found that when reading papers, three participants reported navigating to the diagram first, before reading the

text, suggesting diagrams have a unique and hitherto unexplored role in scholarly communication about novel systems.

A single interview study is insufficient to validate a hypothesis. Rather than validating our findings through a larger scale survey, it seems more useful to attempt to provide part of the solution and validate the utility of that work. The majority of interviewees asked for guidelines (and ideally a software tool) to support their diagram creation process. This request, combined with themes uncovered by the interviews, led to the creation of an initial set of guidelines for authors of these diagrams. We are not intending to create the diagramming software to support this as part of this thesis. Instead we will focus on an initial set of guidelines and a framework for engaging with the community and refining them collaboratively and scientifically [6], which could result in a more formal, but evolving, diagrammatic language [2].

Other completed work includes quantitative corpus analysis of diagrams found in conference proceedings [7,8], qualitative analysis using VisDNA [3], and a synthesis reasoning about the manifest heterogeneity [9]. Theoretically, there is the open question of what makes a good systems diagram, and to this end we explored quantification of diagram “quality” [4].

3 Expected Contribution

We are now assessing the usability and utility of a set of guidelines for natural language processing neural network diagrams. The complexity is both in terms of methodology and measures, neither of which are well established. The study and resulting measures span both creation and interpretation of the diagram, and describe pragmatic communicative properties based on the tasks the diagram is employed for. Tasks reported by interview participants include finding “how the system works” and “as an index to text”, and over 12 different tasks were reported from 12 participants. Adapting from specific User Interface Guidelines [1], our measures are based on two fundamental aspects:

1. Usability of the guidelines themselves: Do they result in the authors making changes to their diagrams?
2. Utility of the guidelines for readers: Does exposure of the author to the guidelines make their diagrams “better” for readers than without the guidelines?

In the domain of scholarly diagrams the accuracy of communication is important: Was the intended message of the author understood by the reader? This encompasses aesthetic preferences, vagueness, excessive variability, and any confusion or ambiguity for the reader, all of which could result in a negative sentiment and ultimately unsuccessful peer review. These reader-centric utility measures are inspired by the results of the interview study. In order to ensure the “correctness” of the reader response is measured as accurately as possible, we are asking the original author to assess the readers responses themselves. This also gives us an opportunity to get their end-to-end feedback on the entire process, and for the authors to identify any potential misunderstandings that have arisen as a

result of their diagram. The repeated engagement of the author in this method makes this a single study rather than two or three distinct studies. This choice is also pragmatic: The participants are required to have specialist technical skills that are in high demand and the potential participant pool is therefore relatively limited.

4 Aspects on which Advice is Desired

1. Increasing impact of diagram research, especially with relevant communities.
2. Approaches for quantitatively analysing diagrams through a corpus-based approach.
3. Technical system diagrams: Examination of, optimisation of, and developing tooling for.

Any discussion and potential future collaboration on these topics would be very welcome.

Acknowledgements

This work is being conducted under the supervision of Prof. Caroline Jay and Dr. André Freitas, who have my thanks. The PhD is funded by University of Manchester Department of Computer Science.

References

1. Al-Sa'di, A.: User interface guidelines for Tablet PC Arabic educational applications. Ph.D. thesis, Auckland University of Technology (2018)
2. Marshall, G., Freitas, A.: The Diagrammatic AI Language (DIAL): Version 0.1. arXiv preprint arXiv:1812.11142 (2018)
3. Marshall, G., Jay, C., Freitas, A.: Understanding scholarly neural network system diagrams through application of VisDNA. In: International Conference on Theory and Application of Diagrams. Springer, Cham (2021)
4. Marshall, G.C., Freitas, A.: Measuring diagram quality through semiotic morphisms. *Semiotica* **2021**(239), 125–145 (2021)
5. Marshall, G.C., Freitas, A., Jay, C.: How researchers use diagrams in communicating neural network systems. arXiv preprint arXiv:2008.12566 (2020)
6. Marshall, G.C., Jay, C., Freitas, A.: Diagrammatic summaries for neural architectures. In: Beyond static papers: Rethinking how we share scientific understanding in ML-ICLR 2021 workshop (2021)
7. Marshall, G.C., Jay, C., Freitas, A.: Number and quality of diagrams in scholarly publications is associated with number of citations. In: International Conference on Theory and Application of Diagrams. Springer, Cham (2021)
8. Marshall, G.C., Jay, C., Freitas, A.: Structuralist analysis for neural network system diagrams. In: International Conference on Theory and Application of Diagrams. Springer, Cham (2021)
9. Marshall, G.C., Jay, C., Freitas, A.: Why scholars are diagramming neural network models. In: International Conference on Theory and Application of Diagrams. Springer, Cham (2022)