

Human-in-the-loop approach to digitisation of engineering drawings^{*}

Andrew M. Fagan, Graeme M. West, and Stephen D. J. McArthur

University of Strathclyde, Glasgow, United Kingdom
`andrew.fagan@strath.ac.uk`

Abstract. In the nuclear power industry, the high-performance black-box systems which are prevalent in modern AI research are difficult to match to applications which take advantage of their strengths. These systems generally require volumes of labelled or well formatted data, and will provide a high level of performance which cannot be easily understood, explained or audited. Indeed, most AI systems deployed in this industry are done so with constant oversight by a skilled human operator, negating many of their advantages in cost, speed and reliability. In many cases, the time taken to format data for these systems is prohibitive even when automation might be desirable. This paper presents a framework for deploying a variety of AI techniques to industries where human oversight is required. Instead of treating the user as an external element while automating the task, this framework incorporates them as an active participant in the process, augmenting their performance while leveraging their strengths to make the AI systems more reliable and user-friendly. The framework is concerned primarily with the problem of digitising Elementary Wiring Diagrams, an important class of engineering drawing. These are in regular use even as low-quality scans of paper documents, and while digitising them is desirable and has many potential benefits, the level of time investment required by skilled engineers is prohibitive. Mistakes in digitisation are also potentially costly, meaning that the skilled engineer must remain involved in the process.

Keywords: Human-in-the-loop · Digitisation · Engineering Drawings

1 Introduction

Engineering drawings serve several key purposes in design engineering, maintenance and asset management. The ability to quickly determine which components are connected to and affected by others can assist in tasks such as fault diagnosis and upgrading equipment, especially when designs are digitised with intelligent metadata and cross referencing between multiple drawings.

^{*} Copyright © 2021 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0). Supported by the Engineering and Physical Sciences Research Council (EPSRC), the Advanced Nuclear Research Centre (ANRC) and Bruce Power.

In the nuclear power industry, and in others such as the oil and gas industry, many assets have been in use for a significant period. With their associated drawings having been originally drafted on paper, they are now digitised in low-quality scanned formats. As it would take considerable effort by skilled engineers to manually redraft these as modern CAD drawings or to review them to find useful metadata, a system for intelligently parsing these drawings would be of significant value.

Unfortunately, while this problem has been approached for many classes of drawings, such as piping and instrumentation diagrams (P&ID), success has resulted from the existence of at least moderately sized labelled datasets, which are used as a starting point for symbol classification. In attempting to utilise this library of research on a class of drawing where no such dataset exists, such as elementary wiring diagrams (EWDs), this is the first and most manually intensive problem to overcome.

2 Background

2.1 Digitisation of Engineering Drawings

Digitisation of engineering drawings has been of interest to many since as early as the 1980s[1], and continues to be worked on. The most active research on the subject is being done on P&ID diagrams. These drawings are comparable in many ways to EWDs, in that they are composed of symbols, text and topological connections. They differ primarily in the family of symbols which they utilise, with P&ID diagrams featuring a more complicated variety of symbols, with embedded text and subtle lines which can change the meaning of a symbol.

In the domain of P&ID digitisation, Morena-Garcia identified three key challenges: quality, skewing and topology [2]. Briefly, the quality problem refers to the variance caused by hand drawn symbols and distortion caused by scanning paper documents. The skewing problem refers to class imbalance caused by some symbols being more common than others, while the topology problem is both the in-drawing topology of recognising lines and connecting symbols, as well as the meta-topology of drawings which connect to others. Again it can be seen intuitively that all of these problems transfer to the EWD domain, and so must be considered.

A wide variety of techniques have been applied to address these common problems, for example the utilisation of Generative Adversarial Networks [3] or class decomposition [4] to address the class imbalance problem. However, this research relies on a large dataset of labelled P&ID symbols[5], so applying this rich library of techniques to EWDs requires either a manually intensive amount of splitting and labelling drawings, or some interim techniques to allow the dataset to be expanded.

2.2 Human-in-the-loop systems

In nuclear power applications, there is a consistent problem with utilising data intensive AI solutions. In many nuclear applications, including the problem of

digitising engineering drawings, there is a lack of quality data. Labelled data is usually minimal, hampering supervised techniques, and the data which is available is often low quality and poorly formatted, which creates difficulties applying unsupervised or semi-supervised methods. In addition to this, most modern AI solutions are “black boxes”, which may provide a high degree of performance but are not easily understood. This is also a poor fit for the nuclear domain, as most decisions must be justified and audited.

This leads to the main difficulty with nuclear applications of AI: a skilled human engineer must remain involved in any uncertain process, usually either checking the AI decision at the end or making their own decision based on AI suggestions. Therefore an AI must provide enough of a benefit to justify the time of a skilled engineer to format or label data, as well as the development and verification time required to design and implement such a system, and must perform better with supervision than an engineer working alone.

However, these restrictions also present an opportunity to explore ways in which the human expertise can be leveraged to improve performance, by taking advantage of the human in the loop rather than working around them.

Human-in-the-loop (HITL) systems refer generally to systems where AI and humans interact, though is more commonly used when referring to systems in which the human is not simply a supervisor of the AI, but when the human and AI cooperate to accomplish a task [6]. This human-AI cooperation is sometimes referred to as shared autonomy [7].

HITL has been applied to the process of designing AI systems in HELIX [8], a human centric tool for rapidly design and iteration on data science workflows. The design of HELIX involves providing the user with a high fidelity interface which suits their specialisation (in this case a programming language) and high quality visualisation tools to allow them to do the same work more effectively. The idea of the human user being a skilled expert in the target domain and the utilisation of their knowledge to improve the system transfers well to the nuclear domain, where the users will be experienced engineers, though the critical difference is that they will not necessarily have programming experience, but rather engineering and design experience.

3 Framework for human-in-the-loop digitisation

In order to address the problem of digitising EWDs using a human in the loop, there are several key considerations to address. To be of use to the end user, the system must be more efficient than an engineer simply digitising the drawing manually, even when the AI modules are at their weakest or have not yet been implemented. If the user is willing to spend a significant amount of time labelling data, they would prefer to redraw instead, in line with their expertise.

This also means that even once AI modules are implemented and are displaying a reasonable performance level, the user should not have to spend time tweaking until the AI digitises the drawing by itself. If the user makes corrections, those should be logged for future use, but the completely labelled and

connected drawing they have produced should not be discarded. It is the desired output of the system, and spending more time on a drawing once it has been digitised is wasteful.

The output of the system is the original drawing marked up with a label and bounding box on each symbol, blocks of text transcribed and topological connections represented by connecting lines. Additional information such as that originally contained in embedded tables would instead be attached to the relevant component in a text field. This intermediate state could in principle be used to automatically generate a new digital drawing, but could also be applied to other potential uses, such as a subsequent database application for spanning between components across an entire plant, allowing for many potential workflow improvements.

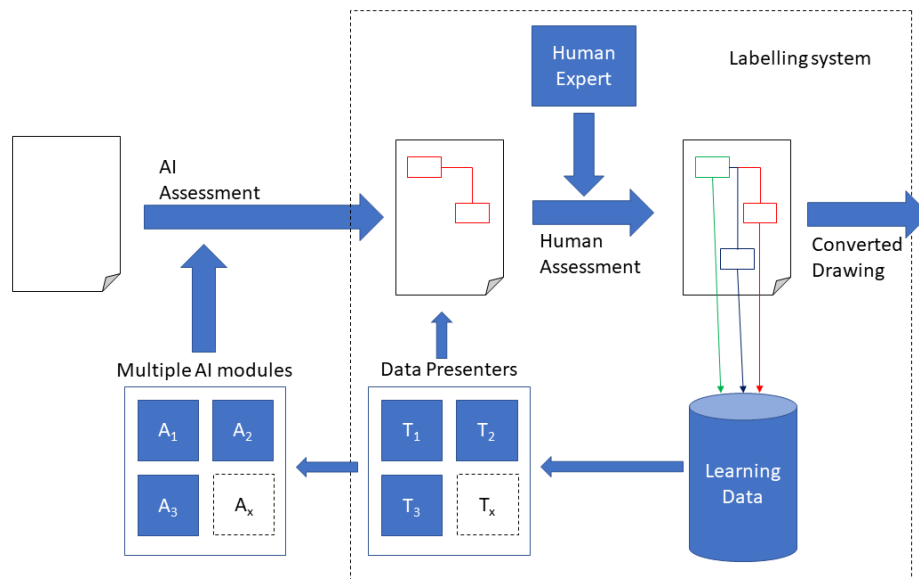


Fig. 1. HITL framework for EWD digitisation

The proposed framework is shown in figure 1. It consists firstly of an attempt by any number of AI modules to parse the drawing. These modules could be as simple as a classifier which identifies individual symbols or a rule based system governing how components can connect together, but could equally be complex holistic solutions of the type which exist for P&ID diagrams, as discussed in section 2.1. They return their findings using a common interface, represented on the drawing in the same format available to the user.

The modularity provided by a swappable bank of AIs moves the focus away from individual techniques and instead to the performance of the system as a whole and the interplay between human and AI. It also allows the addition

and removal of modules as appropriate. Early on, for example, a Convolutional Neural Network is likely to be of extremely limited value due to their reliance on volume of available data, but when a large dataset is available their performance is likely to be very high in comparison with other methods. We therefore might remove a CNN based method in the early stages, and only reintroduce it when it attains a reasonable level of accuracy.

In contrast, initial work utilising a Quality Assured Template Matching [9] based approach successfully identified 20% of symbols with no false positives, and required only one example per symbol class, but additional data and experimentation provided greatly diminished returns, suggesting it might be a useful module early on, but might fall out of favour as more data becomes available, or else might be relegated to a reliable check on the output of other modules. In principle, this selection and comparison of modules could be done automatically based on the accuracy metrics of modules, but in the simplest implementation of this framework it would instead be a choice on the part of the designer which they might revisit regularly.

The annotated drawing is then opened up to the human expert, who has the opportunity to modify the AI attempt, either by adding additional labels not flagged by an AI, or by modifying or deleting AI outputs they disagree with. The interface at this stage should be extremely simple and intuitive. The users actions at this stage, and therefore areas where AI modules were correct or incorrect, will be logged into the learning data repository, as well as areas where modules might have disagreed with one another.

The learning data will not be a single exhaustive dataset of the type that might be used to train a machine learning system, such as labelled pictures. It will instead be a log of user actions. If the user highlights a region and labels it as a relay, the data logged would be the name of the drawing in question, the coordinates highlighted, and the label “relay”. Likewise if the user deletes or modifies an area the AI flagged as a resistor, the coordinates would be logged either as a negative example or with the label flagged by the user. By saving this log instead of a single formatted dataset, we allow creativity when deciding what data would be valuable to a new AI module, by the use of the presenter modules.

The data presenters are another modular element to the system which allows the learning data to be parsed in many ways. A presenter in this case is a simple module which transforms a subset of the learning data into something usable by another part of the system. A simple example of this would be a presenter which collates all of the learning data relating to symbols, in the form of drawing name, coordinates and label, and extracts the pixels from the drawing to create a dataset for training a classifier. The presenters might be simple programs, or might themselves be more complex systems. An example of a different type of presenter is provided by the implemented subset of the system which, given a single coordinate point clicked by the user, performs simple image manipulation to find a potential area of interest which it flags on the drawing for the user’s further input. Another more complex presenter could take that same single co-

ordinate and run a classifier on a sliding window around it to identify the area of interest and its type more accurately.

The current implementation consists of the subset of the framework marked “Labelling System”. The lack of implemented AI modules in this version allows verification of the speed at which the operator will initially be able to digitise a drawing. By utilising the aforementioned presenter, the user can click a single point on a symbol, resulting in a best guess bounding box drawn around it based on the centre of mass within that region. The user can then pick from a list what the symbol is as well as adjust the bounding box, and their decisions are logged. It is then easy to add text to the symbols and draw links between them. This is a smoother process than manually segmenting each image, and results in more consistently sized images being identified. Compared to manually redesigning a drawing in a CAD format, which industrial partners suggest takes several hours, this process takes only minutes.

4 Conclusions and future work

The framework is still at an early stage of implementation, and requires further testing to verify the extent to which it improves the digitisation workflow. The framework does take steps towards all three of the challenges to digitisation of engineering drawings. The quality problem is addressed by minimising the impact of quality-based failures: the user has the ability to intercept mistakes made by the AI before they are output and allows the AI more leniency before enough data is available to potentially overcome the poor quality.

The skewing problem is also addressed by this, but is further supported by the modularity of the framework. The class imbalance caused by the relative rarity of some symbols over others can eventually be addressed by the generation of additional artificial data, once enough labelled data exists, but can also be addressed in the shorter term by building smaller and more specific modules which identify only one symbol at a time.

The topology problem is in some ways the hardest problem to overcome, and features the least available research. It requires taking a more holistic view of the document rather than dividing into small units like identification of symbols or text, while also requiring a high level of fidelity in other systems before beginning. The framework addresses this problem both by allowing the user to make connections in the early stages before a module is written to overcome this problem, but also by the modularity of the system, by allowing identification to be easily abstracted away into another module.

In the future, the framework will be ported to other domains, initially to other similar drawings such as P&ID diagrams, but eventually to entirely different classes of problems. The domain of AI planning [10] in particular features scheduling problems which a skilled expert must undertake. Many techniques which produce partial solutions are available, but the human centred framework might provide an excellent way to turn these techniques into a more complete solution.

References

- [1] H. Bunke, “Automatic interpretation of lines and text in circuit diagrams.,” *Pattern recognition theory and applications. Proc. NATO ASI, Oxford, 1981*, pp. 297–310, 1982.
- [2] C. F. Moreno-Garcia and E. Elyan, “Digitisation of Assets from the Oil & Gas Industry: Challenges and Opportunities,” in *2019 International Conference on Document Analysis and Recognition Workshops (ICDARW)*, Sydney: Institute of Electrical and Electronics Engineers (IEEE), Nov. 2019, pp. 2–5.
- [3] E. Elyan, L. Jamieson, and A. Ali-Gombe, “Deep learning for symbols detection and classification in engineering drawings,” *Neural Networks*, 2020, ISSN: 18792782.
- [4] E. Elyan, C. M. Garcia, and C. Jayne, “Symbols Classification in Engineering Drawings,” in *Proceedings of the International Joint Conference on Neural Networks*, vol. 2018-July, Institute of Electrical and Electronics Engineers Inc., Oct. 2018, ISBN: 9781509060146.
- [5] E. Elyan, C. F. Moreno-García, and P. Johnston, “Symbols in Engineering Drawings (SiED): An Imbalanced Dataset Benchmarked by Convolutional Neural Networks,” in *Proceedings of the 21st EANN (Engineering Applications of Neural Networks) 2020 Conference*, Springer, Cham, Jun. 2020, pp. 215–224.
- [6] D. Sousa Nunes, P. Zhang, and J. Sa Silva, “A Survey on human-in-The-loop applications towards an internet of all,” *IEEE Communications Surveys and Tutorials*, vol. 17, no. 2, pp. 944–965, Apr. 2015, ISSN: 1553877X.
- [7] S. Javdani, H. Admoni, S. Pellegrinelli, S. S. Srinivasa, and J. A. Bagnell, “Shared autonomy via hindsight optimization for teleoperation and teaming,” *International Journal of Robotics Research*, vol. 37, no. 7, pp. 717–742, Jun. 2018, ISSN: 17413176.
- [8] D. Xin, L. Ma, J. Liu, S. Macke, S. Song, and A. Parameswaran, “Accelerating human-in-the-loop machine learning: Challenges and opportunities,” in *DEEM’18: Proceedings of the Second Workshop on Data Management for End-To-End Machine Learning*, New York, NY, USA: Association for Computing Machinery, Apr. 2018, pp. 1–4.
- [9] J. Cheng, Y. Wu, W. Abdalmageed, and P. Natarajan, “QATM: Quality-aware template matching for deep learning,” in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 2019-June, 2019, pp. 11 545–11 554, ISBN: 9781728132938.
- [10] T. Chakraborti, S. Sreedharan, and S. Kambhampati, “The emerging landscape of explainable automated planning & decision making,” in *IJCAI International Joint Conference on Artificial Intelligence*, vol. 2021-Janua, International Joint Conferences on Artificial Intelligence, Jul. 2020, pp. 4803–4811, ISBN: 9780999241165.