The Challenges of Learning Machine Learning for Non-Technical Users or Novice Programmers

Robert Cinca
UCL Interaction Centre
University College London
London, UK
robert.cinca.14@ucl.ac.uk

Mirco Musolesi
University College London
The Alan Turing Institute
London, UK
m.musolesi@ucl.ac.uk

Stefan Kreitmayer
UCL Interaction Centre
University College London
London, UK
s.kreitmayer@ucl.ac.uk

Yvonne Rogers
UCL Interaction Centre
University College London
London, UK
y.rogers@ucl.ac.uk

Enrico Costanza
UCL Interaction Centre
University College London
London, UK
e.costanza@ucl.ac.uk

CHI, '19 Glasgow, UK
ACM 0-12345-67-8/90/01.
http://dx.doi.org/10.1145/2858036.2858119

Abstract
Recent breakthroughs in machine learning and AI have led to increasing interest in data science among the general public. This trend is evident in the abundance of new online courses, interactive websites like Kaggle and face-to-face code camps marketed to people with little or no programming experience. However, there appears to be a mismatch between how machine learning is commonly taught and the specific needs of this emerging user group. Our paper aims to explore this mismatch and highlight some of the challenges for explainable AI and interaction design. Our analysis draws on evidence from a study with 144 beginners at machine learning. Relevant bottlenecks were identified in four main areas, namely: machine learning, maths, programming and operating systems. Our preliminary findings suggest that educational systems should be designed to selectively "blackbox" and "whitebox" these areas in ways that are customisable and transparent to the learner.

Author Keywords
Learning machine learning; explainable AI; data science; blackbox; whitebox

ACM Classification Keywords
H.5.m [Information interfaces and presentation (e.g., HCI)]: Miscellaneous

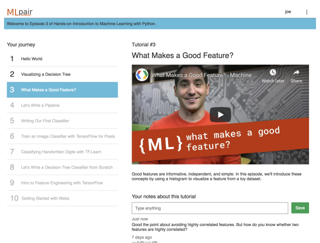

Figure 1: A simple user interface allowed participants to comment on each machine learning tutorial.

Introduction

Large-scale data analysis promises to change many areas, ranging from driverless cars in the automobile industry to smart thermostats within the home. Much of this data processing involves training machine learning models of ever-increasing complexity. Through those breakthroughs, we have observed an increase in the number of novice programmers that are taking up machine learning (ML), with the goal of applying it to their own fields. This emerging group often uses online tutorials and practical examples (among other resources) to teach themselves machine learning, with mixed levels of success.

In response, graphical user interfaces (GUIs) have emerged to help these non-technical audiences with their ML analysis. Weka, one of the most popular GUI-based ML tools, provides a platform to quickly train and visualise models without writing code [8]. Similarly, a wide range of libraries and APIs exist to help programmers use ML while hiding the underlying maths. Conversely, systems like Matlab and Octave can support a mathematical approach to learning ML, while shielding the novice from the details of computational implementation. Each of these approaches promotes clarity in some aspects while making others opaque for the sake of users making quick progress.

Black and White Boxes

The distinction between blackbox and whitebox systems is common in computing as well as education. In the context of human learning with computer models and simulations, Jonassen and Strobel [4] add a third category which they call the glass-box. Table 1 provides an overview of these different types of system.

Table 1: Types of systems according to Jonassen and Strobel [4]

System Type	Implication
Blackbox	The learner **cannot see** the model
Whitebox	The learner **can see** the model
Glassbox	The learner **can change** the model

Aims

This paper aims to contribute to the workshop in the following ways:

1. to describe the range of pedagogical needs and expectations among ML beginners, using empirical evidence from a recent user study
2. to contrast these needs with the current tool landscape and educational offerings for ML
3. to inform the design of explainable-ML systems for beginners

Method

This research draws evidence from a study where pairs of self-motivated adult learners followed a 10-part video tutorial series [3] on ML. Seventy-two pairs were recruited through meetup.com and word of mouth. Each pair was given access to the material through the custom web interface shown in Figure 1, allowing them to leave comments during and after each tutorial. Our qualitative analysis focuses on these comments as well as preliminary and final interviews with individual participants. We first performed open-ended coding on the data. Themes were then identified regarding individuals' motivations to study ML, their expectations of the learning experience and the challenges they encountered on their journey.

Findings

We identified several themes from the study. These are as follows: Backgrounds and motivations vary; Beginners enjoy real world problems; Beginners want to understand how models behave; Visualisations can engage and support understanding; Interests vary regarding coding details; Operation system details can frustrate.

Backgrounds and motivations vary

People approach machine learning from a diverse range of backgrounds and with different motivations. For example, one participant said that she worked in the field of natural language processing and wanted to communicate more effectively with the engineering team. Another, a physicist, said her aim was *"to broaden [her] data analysis skills"*.

Beginners enjoy real-world problems

Beginners expressed positive reactions when the tutorial involved analysing real data sets: *"We found this one to be probably the most useful because [the instructor] talks about real data sets"*.

Beginners want to understand how models behave

Several comments revolved around understanding how models behave. For example, one participant struggled with the random splitting of data: *"The two classifiers give different results when you run them every time, one outperforms the other at times"*. Another participant was unsure how to improve the model: *"testing the classifier gives about 90%; how do we improve that?"*. This is in line with previous findings that users are interested in model behaviour [6].

Visualisations can engage and support understanding

Many participants mentioned that they liked the visualisations and that the graphs helped them better understand the underlying model. These findings corroborate prior research in the field of Explainable AI [1, 5, 6, 7].

Interests vary regarding coding details

With regard to the coding implementation, levels of interest varied widely among the diverse, largely non-technical group of beginners. Some participants expressed a desire for full visibility: *"I'm a bit annoyed by the fact that they use Scikit Learn as a black box and never give definitions properly"*. Others questioned the need for coding details: *"not sure when we would practically write a classifier"*.

Operation system details can frustrate

Many participants expressed frustration about issues such as *"getting Anaconda installed on MacOS (had issues with it altering the PATH)"* or *"I wasn't able to install Docker in my machine"*. Such issues were generally perceived as distracting from learning about ML.

Discussion

While our findings are preliminary, they suggest a new perspective on the current landscape of tools for learning and using ML. Our discussion centres around: Current learning ML methods and whiteboxing attempts; Blackboxing the development environment; and How to choose the right tools and pathways when 're-painting the boxes'.

Several ways of teaching ML to non-technical experts have emerged. One such way, through the use of GUIs (such as Weka), actually blackboxes the ML by hiding away all the implementation details. Not all user want this. As one participant said: *"weka being GUI driven, would have been great to be able to do this all via CLI or via python programs"*. Although some users prefer that when coding for their own projects, at least for the initial learning process, the study showed a preference for a more detailed breakdown of the models.

Another way of teaching ML would be through an interactive visual of the model. TensorFlow's Playground [2] aims to do just that, whilst bringing a transparency to AI models. Although the program aims to visualise the insides of a neural network, it is just a demo, and therefore not a tool. In addition, it can still be considered a blackbox as it does not explain how the weights are decided. However, compared to Weka, it is a step in the right direction because it visualises connections between the weights.

Tools like Google Colab were appreciated by novices, as they prevented users from getting stuck, by blackboxing the operating system and development environment.

Conclusion

As machine learning techniques are rapidly adopted by a wide range of professional fields, new challenges emerge for educating a diverse audience of non-technical users and novice programmers. Our analysis of user comments and interviews revealed a better understanding of the needs and expectations among this important user group. While there are some apparent commonalities in users' needs (beginners want transparency in the ML concepts, whilst not having to worry about the development infrastructure), interest in computational and mathematical detail varied greatly across participants. We argue that by selectively black-boxing and white-boxing these aspects, educational systems can embrace the audience's diversity of backgrounds and goals when it comes to learning machine learning.

Acknowledgements

This work was supported by the UK Engineering and Physical Sciences Research Council (EPSRC) grant EP/R513143/1 for the University College London Interaction Centre (UCLIC). Our study was approved by the UCLIC Ethics Committee (UCLIC/1718/003).

REFERENCES

1. Ashraf Abdul, Jo Vermeulen, Danding Wang, Brian Y Lim, and Mohan Kankanhalli. 2018. Trends and trajectories for explainable, accountable and intelligible systems: An hci research agenda. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 582.

2. Google Developers. 2015. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. (2015). `http://tensorflow.org/` Software available from tensorflow.org.

3. Google Developers. 2017. Machine Learning Recipes. `https://www.youtube.com/playlist?list=PLOU2XLYxmsIIuiBfYad6rFYQU_jL2ryal`. (2017). [Online; accessed 19-July-2018].

4. D Jonassen and J Strobel. 2006. Modeling for meaningful learning. *Engaged learning with emerging technologies* (2006), 1–27.

5. Vassilis Kostakos and Mirco Musolesi. 2017. Avoiding pitfalls when using machine learning in HCI studies. *interactions* 24, 4 (2017), 34–37.

6. Brian Y Lim and Anind K Dey. 2009. Assessing demand for intelligibility in context-aware applications. In *Proceedings of the 11th international conference on Ubiquitous computing*. ACM, 195–204.

7. Brian Y Lim and Anind K Dey. 2010. Toolkit to support intelligibility in context-aware applications. In *Proceedings of the 12th ACM international conference on Ubiquitous computing*. ACM, 13–22.

8. Ian H Witten, Eibe Frank, Mark A Hall, and Christopher J Pal. 2016. *Data Mining: Practical machine learning tools and techniques*. Morgan Kaufmann.

Author Background

I am a 1st year PhD student at the UCL Interaction Centre (part of University College London), with a background in Computer Science. My research interest is in explainable AI. My motivation for participating in the Human-Centered Study of Data Science Work Practices Workshop stems from my interest in building tools or methods to support human activities in data science work.