Teaching the Foundations of Machine Learning with Candy

Anonymous Authors[1]

Abstract

Machine learning is ubiquitous in decision making processes across society. The prevalence of ML drives a need for improved education in key concepts at the secondary and tertiary levels that not only trains people to become informed citizens but also trains future researchers to be both principled and ethical practitioners. In this vein, we present a structured classroom activity that simultaneously teaches both supervised classification and critically thinking about ML applications and ethics. We use an active, object-based learning approach to teach supervised classification using a variety of candies, and a problem-based scenario to encourage critical questions about ethics in ML applications.

1. Motivation

Machine learning (ML) has become an integral part of decision-making across virtually all sectors and industries. ML has even been applied in areas where the ethics of its use are under scrutiny — for example, from public health (Ravì et al., 2017) and diagnostic medicine (Suzuki, 2017) to decisions about hiring and promotion (Raghavan et al., 2020). Moreover, the extraordinary pace of both development of new algorithms and their implementation in real-world decision making has left curricula lagging across the board.

At universities, ML is often taught only as part of dedicated computer science or data science degrees, while secondary school curricula are traditionally even slower to adopt new content (though an increasing number of proposals and programmes exist, e.g. Gavalda 2008; Ali et al. 2019). At the same time, an ecosystem of for-profit and non-profit "bootcamps", massive online open courses (MOOCs) and online tutorials have aimed to democratize education in ML in a more ad-hoc, haphazard fashion, leaving both learners and teachers with a wide range of materials and little guidance for its use. In addition, ML materials often stress one of two extremes: theoretical foundations or immediate application.

Teaching materials that focus on very theoretical foundations require a strong mathematical background and may not suit learners with less quantitative education experience. For example, the engineering and computer science curricula at the University of Florida contains two senior-level courses "Current Topics in Machine Learning" in their *Sustainable Model for Assimilating Research and Teaching* undergraduate program (Georgiopoulos et al., 2009). In these courses, students learn from and work with multiple faculty members whose research areas use machine-learning techniques. Prior to these courses, students are also provided with introductory modules on ML in their sophomore and junior years. Although the curricula provides a well-structured program for engineering and computer science students, the teaching materials and approach may be less accessible to people outside these fields. On the other extreme, there are ML materials that aim to quickly bring novice practitioners to a level where they can apply popular algorithms in a short time[1] . But this rushed application pays a price — the practitioners may use ML largely as a black box.

We believe there is a need for and interest in teaching machine learning with the twin objectives: (1) providing learners with the tools and understanding to become the next generation of researchers and practitioners, and (2) training learners to critically evaluate the impact that ML has on their and others' lives. The literature on best practices in teaching machine learning are still new, but as Sulmont et al. (2019) found through an exploratory study, "Student preconceptions include ideas that ML is important, but also not accessible. [...] Instructors reported students having difficulty appreciating the human decision-making aspects of ML, and overestimating the power of ML to solve real-world problems." Both of these findings speak to our twin objectives.

Creative and fun learning activities for teaching machine learning also exist. These approaches may gloss over the mathematical foundations in order to stress conceptual un-

[1]Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the European Conference on Machine Learning (ECML) and Principles and Practice of Knowledge Discovery in Databases (PPKD) Conference. Do not distribute.

[1]e.g. https://machinelearningmastery.com/machine-learning-in-python-step-by-step/

derstanding. For example, Rattadilok et al. (2018) present a gamified approach to teaching ML through use of a war-type game called "Clash of Clans". Learners are provided with different machine learning techniques that automate game-play, and are asked to test their ability to mimic the outcome of novice, intermediate, and expert human players. The authors argue that the real-time experience of using these techniques provides another avenue for learners to explore the concepts of ML without mathematical rigor.

Education research suggests that *fun* in the delivery of instruction is positively related to learner engagement (Tews et al., 2015; 2017), and is identified as a motivator to attend classes and seek out new skills (Lucardie, 2014). In addition, *object-based learning*, one of numerous *active learning* approaches, involves the active integration of objects into the learning environment in order to facilitate deep learning and stimulate the learner's imagination (Hannan et al.; Chatterjee & Hannan, 2016). In the field of statistics, there are many active-learning strategies that combine both fun and object-based learning by engaging learners with candy and chocolates (e.g., SherriJoyce & Alexander, 1994; Alexander & SherriJoyce, 1994a;b; Dyck & Gee, 1998; Richardson & Haller, 2002; Lin & Sanders, 2006; Downey, 2013; Schwartz, 2013; Froelich & Stephenson, 2013; Eadie et al., 2019).

Thus, we were motivated to design a creative and fun activity that meets a balance between the two aforementioned extremes of theoretical foundations and immediate applicability of machine learning. Moreover, we sought to instill critical thinking skills about the ethics of using ML techniques in real-life situations. The development of our unit was motivated by three core questions we keep encountering in our teaching practice:

- ML is often framed in the popular discourse as the domain of mathematicians and computer scientists, which can make it intimidating to individuals without training in these fields. As ML practitioners and teachers, how can we demystify ML to a diverse audience of learners with varying levels of quantitative background?
- Given the prevalence of ML in the world, how can we encourage learners to engage with and understand how ML affects their own lives? Can we put their experiences with ML into the context of the larger, evolving societal discussion around the use of computation in human decision making?
- Machine learning is now easy to use and implement: commercial platforms have graphical user interfaces that abstract away much of the underpinning mathematical concepts. How can we encourage learners without a strong quantitative background to engage with the most important concepts and principles in ML? How can we enable them to make good (and ethical) decisions in the process, rather than use these technologies as a pure black box?

In this paper, we cover the approach, the learning outcomes, and also supply the materials to implement the activity in class.

2. Unit Data and Design: Candy, Concepts, and Critical Thinking

Our goal is to provide a hands-on, modular tutorial following The Carpentries' principles for lesson design[2]: working back from key concepts and learning outcomes, we designed a number of challenges that allow learners to engage with (and in some instances, self-discover) concepts important to designing and implementing ML models for real-world decision making.

The emphasis of our ML activity is on building intuition rather than mathematical rigor in a modular, hands-on design. At the same time, we include some mathematical concepts and advanced challenges appropriate for an audience with early college-level mathematics and programming knowledge. We envision this unit – enough content for about three to four hours, depending on depth and the audience's mathematical foundations – to be embedded in a larger formal (data) science curriculum and spread over multiple lessons, or used as a self-contained unit in a Carpentries-style workshop or summer school.

Many ML tutorials and materials are based around well-known data sets like the Iris data set (Anderson, 1936; Fisher, 1936). However, these data sets, while well-understood and curated, were often developed with algorithm bench-marking in mind, rather than teaching. The ability and speed with which learners pick up new concepts is correlated with the familiarity of the components with which these concepts are taught (e.g. Reder et al., 2016). Beginners may find it especially difficult to build their knowledge base about ML with these existing, standard data sets.

Instead of using a well-known training set, we built the unit around a problem wherein the learners must classify different types of candy (see also Figure 1). While we admit the problem is somewhat contrived, it is much closer to our learners' lived realities than many standard data sets. In addition, we believe the involvement of candies helps to establish a welcoming, non-threatening learning environment by introducing a playful character to the lesson design. It also allows learners to physically interact with (and, on occasion, eat) their training samples. While we were initially skeptical that this design would work well for learners

[2] https://carpentries.github.io/curriculum-development/

beyond the undergraduate level, we have in practice found that individuals of all ages seem to enjoy chocolate-based educational experiences.

Overall, the unit emphasizes the *process* of ML over individual algorithms, which are often the focus of traditional computer science-based curricula. We place particular importance on encouraging learners to ask critical questions about the real-life outcomes of computer-supported decision making early in the process of interacting with or designing a ML system, and keep asking these questions throughout the process. In particular, we place an episode on ethical considerations in building and producing models early in the unit, and refer back to it throughout the remainder of the unit. In this context, our design also dedicates significant discussion time to the process of constructing features for classification, and on evaluation metrics. Overall, we designed the unit to generate an engaging interplay between taught components and self-discovery of important concepts via challenges.

2.1. Overall Intended Learning Outcomes

By the end of the unit, learners should be able to ...

- extract meaningful features from data sets, visualize these features, and draw decision boundaries between classes,
- describe how k-nearest neighbours and (optionally) logistic regression classify samples and draw decision boundaries, and apply accuracy as an evaluation metric while keeping its limitations in mind, use cross-validation to test algorithm performance,
- understand that ML is not a magical black box — it is a set of well-motivated, but limited mathematical principles for modelling data sets and predicting future instances using those models, and
- critically interrogate the ML algorithms with which they implement or interact.

3. Unit Outline

The unit starts with a problem statement: *You are preparing your house for a party and have a giant box of mixed candies for snacks, but you have just discovered that one of your guests has a peanut allergy. Therefore, you decide to separate out the candies that contain peanuts and put them in a separate bowl. However, the immensity of sorting through thousands of individual candies by hand is both daunting and impractical. Your friend suggests you use a ML algorithm to sort the candy.*

The above problem statement sets up the discussion around ML as a tool for classification, and leads into the larger

Figure 1. An example illustration from the unit depicting different types of candies, used to set up the problem and to engage learners in the process of feature engineering. In in-person contexts or virtual contexts where learners have candies available, they may interact with these objects directly. In other virtual contexts, the images and tables provided with the online materials may be used as stand-ins.

discussion around ethical considerations in ML. The next step in the unit guides learners through hands-on exercises with sets of candies, which can be carried out either with real, widely available types of candies of broadly similar shapes and colours (here, we use different types of M&Ms along with Skittles and Jellybeans), if classes are held in an in-person context where these are available, or with images and pre-recorded data sets available with our teaching materials (see e.g. Figure 1). In a context where real candies are available, learners can interact with the different types of candy to self-discover the process of feature-engineering and drawing decision boundaries.

We introduce K-nearest neighbours as one of the most accessible algorithms (see e.g. Figure 2 for an illustration), as well as logistic regression as an optional and somewhat more advanced method. In the context of these algorithms, learners explore accuracy and its limits as a tool for measuring performance, the effects of imbalanced training data on outcomes, and cross-validation. We follow The Carpentries design principles as much as possible by breaking the unit into modular episodes, each with at most two key concepts and key learning outcomes (though some include optional more advanced materials and challenges).

3.1. Episodes and Intended Learning Outcomes

Below we list the episodes for the unit and their intended learning outcomes.

1. **Introduction**: At the end of this episode, learners should have a working understanding of ML, and be

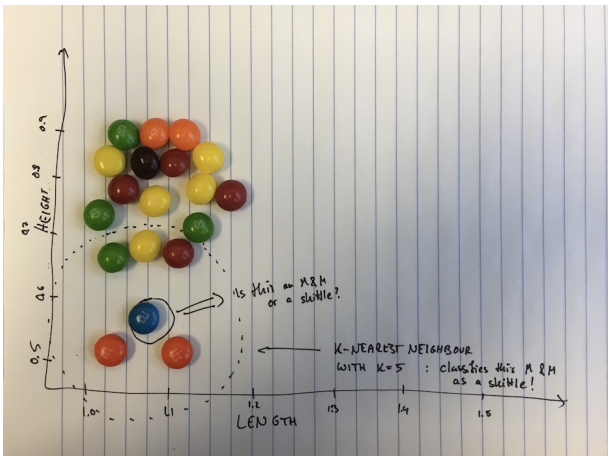

Figure 2. Example illustration from a hands-on exercise in the K-Nearest Neighbours episode: learners place candies on a graph of features, approximately according to their feature measurements, and subsequently perform the K-nearest neighbour algorithm by hand on a small set of new candies. This illustration also demonstrates the effect of imbalanced training data on classification results.

able to identify problems where ML may be applied.

2. **Problem Set-up: Classifying Candies**: Learners should understand the classification problem to be solved

3. **Incorporating Ethics into your ML Project**: At the end of this episode, learners understand that ethical considerations are integral to all ML applications, and are able to critically question their own motives and approach to solving their problem. Learners can also lay out worst-case scenarios in order to design effective counter-measures to misuse.

4. **Feature Engineering**: Learners discover that in order to classify objects, many approaches require well-structured information. They can identify features from a set of objects and quantify these features.

5. **Decision Boundaries**: At the end of this episode, learners understand the relevance of different features for the success of classification. They also understand the concept of a decision boundary and how to draw one.

6. **K-Nearest Neighbours**: Learners should be able to describe in qualitative terms the K-Nearest Neighbours algorithm. More advanced learners should be able to write their own version of the algorithm in code, and use the data they generated to classify new data points.

7. **Model Evaluation**: At the end of this episode, learners know how to split training data into training and test sets. Learners can state the differences between true positives, true negatives, false positives, false negatives, can use these definitions to define precision, recall, specifity and F1 score, and understand when they are relevant.

8. **Logistic Regression**: Learners develop an intuition about logistic regression as a way to model two-class problems. More advanced learners can also write their own version of the algorithm in code, and use the data they generated to classify new data points.

9. **Cross-Validation**: Learners can apply cross-validation to their candy data set to test the validity of their algorithm. Learners understand the difference between overfitting and underfitting and the problems they cause.

10. **Final Thoughts**: Learners reflect on the unit, and identify ethical reasons why ML might not be an appropriate tool for this scenario.

3.2. Online Materials

Online materials and teaching notes are shared in a public (CC-BY 4.0 licensed) *git* repository, hosted on *GitHub*[3] and rendered via *GitHub Pages*[4], using the Carpentries Lesson Template[5]. The repository and the materials therein are open to contributions, sharing and adaptation for individual instructors' needs. Similarly to the Carpentries' style, the materials aim to both give guidance to the instructor and function as a set of lecture notes that learners can refer back to independently after the unit is completed. Pre-recorded data is shared as part of the repository to enable teaching the unit in contexts where candy are not available (e.g. in virtual settings).

4. Conclusion

Given ML's relevance to society, it is important to teach ML to a wide range of learners. Because many educational efforts focus on training *practitioners* and often require a strong quantitative background, we created an introductory unit aimed at teaching important concepts of ML to a variety of learners in a playful, intuitive manner. We envision this work as a resource in the niche between theory-heavy classes and very practical online tutorials.

[3] anonymized version: https://github.com/ecml2020-teaching-candy/machine-learning-with-candy

[4] anonymized version: https://ecml2020-teaching-candy.github.io/machine-learning-with-candy/

[5] https://carpentries.github.io/lesson-example/

References

Alexander, M. and SherriJoyce, K. Colored Candies and PROC MACONTROL: A Sweet Way to Produce Chi-Square Control Charts Part 2: Statistical Presentation. *SUGI 19: Proceedings of the Nineteenth Annual SAS Users Group International Conference*, pp. 1145–1150, 1994a.

Alexander, M. and SherriJoyce, K. Colored Candles and SAS/QC: A Sweet Way to Produce Chi-Square Control Charts. *SAS Conference Proceedings: SouthEast SAS Users Group 1994*, pp. 249–253, 1994b.

Ali, S., Payne, B. H., Williams, R., Park, H. W., and Breazeal, C. Constructionism, ethics, and creativity: Developing primary and middle school artificial intelligence education. In *International Workshop on Education in Artificial Intelligence K-12 (EDUAI'19)*, 2019.

Anderson, E. The species problem in iris. *Annals of the Missouri Botanical Garden*, 23(3):457–509, 1936. ISSN 00266493. URL http://www.jstor.org/stable/2394164.

Chatterjee, H. J. and Hannan, L. *Engaging the senses: object-based learning in higher education*. Routledge, 2016.

Downey, A. *Think Bayes*. O'Reilly, Beijing, 1 edition, 2013. ISBN 9781449370787.

Dyck, J. L. and Gee, N. R. A sweet way to teach students about the sampling distribution of the mean. *Teaching of Psychology*, 25(3):192–195, 1998. doi: 10.1207/s15328023top2503_6. URL https://doi.org/10.1207/s15328023top2503_6.

Eadie, G., Huppenkothen, D., Springford, A., and McCormick, T. Introducing bayesian analysis with m&m's®: An active-learning exercise for undergraduates. *Journal of Statistics Education*, 27(2):60–67, 2019.

Fisher, R. A. The use of multiple measurements in taxonomic problems. *Annals of Eugenics*, 7(2):179–188, 1936. doi: 10.1111/j.1469-1809.1936.tb02137.x. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1469-1809.1936.tb02137.x.

Froelich, A. G. and Stephenson, W. R. How much do m&m's weigh? *Teaching Statistics*, 35(1):14–20, 2013. doi: 10.1111/j.1467-9639.2012.00515.x. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-9639.2012.00515.x.

Gavalda, R. Machine learning in secondary education. *Proceedings TML*, 2008.

Georgiopoulos, M., DeMara, R. F., Gonzalez, A. J., Wu, A. S., Mollaghasemi, M., Gelenbe, E., Kysilka, M., Secretan, J., Sharma, C. A., and Alnsour, A. J. A sustainable model for integrating current topics in machine learning research into the undergraduate curriculum. *IEEE Transactions on Education*, 52(4):503–512, 2009.

Hannan, L., Duhs, R., and Chatterjee, H. J. Object based learning: a powerful pedagogy for higher education. *Museums and higher education working together: challenges and opportunities*, pp. 159–168.

Lin, T. and Sanders, M. S. A sweet way to learn doe. *Quality Progress*, 38(88), 2006.

Lucardie, D. The impact of fun and enjoyment on adult's learning. *Procedia-Social and Behavioral Sciences*, 142: 439–446, 2014.

Raghavan, M., Barocas, S., Kleinberg, J., and Levy, K. Mitigating bias in algorithmic hiring: Evaluating claims and practices. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, FAT* '20, pp. 469–481, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450369367. doi: 10.1145/3351095.3372828. URL https://doi.org/10.1145/3351095.3372828.

Rattadilok, P., Roadknight, C., and Li, L. Teaching students about machine learning through a gamified approach. In *2018 IEEE International Conference on Teaching, Assessment, and Learning for Engineering (TALE)*, pp. 1011–1015. IEEE, 2018.

Ravì, D., Wong, C., Deligianni, F., Berthelot, M., Andreu-Perez, J., Lo, B., and Yang, G. Deep learning for health informatics. *IEEE Journal of Biomedical and Health Informatics*, 21(1):4–21, 2017.

Reder, L. M., Liu, X. L., Keinath, A., and Popov, V. Building knowledge requires bricks, not sand: The critical role of familiar constituents in learning. *Psychonomic bulletin & review*, 23(1):271–277, 2016.

Richardson, M. and Haller, S. What is the probability of a kiss? (it's not what you think). *Journal of Statistics Education*, 10(3):null, 2002. doi: 10.1080/10691898.2002.11910683. URL https://doi.org/10.1080/10691898.2002.11910683.

Schwartz, T. A. Teaching principles of one-way analysis of variance using m&m's candy. *Journal of Statistics Education*, 21(1), 2013. doi: 10.1080/10691898.2013.11889662. URL https://doi.org/10.1080/10691898.2013.11889662.

SherriJoyce, K. and Alexander, M. Colored Candies and Base SAS: A Sweet Way to Produce Chi-Square Control Charts Part 1: Tutorial. *SUGI 19: Proceedings of the Nineteenth Annual SAS Users Group International Conference*, pp. 1470–1475, 1994.

Sulmont, E., Patitsas, E., and Cooperstock, J. R. Can you teach me to machine learn? In *Proceedings of the 50th ACM Technical Symposium on Computer Science Education*, SIGCSE '19, pp. 948–954, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450358903. doi: 10.1145/3287324.3287392. URL https://doi.org/10.1145/3287324.3287392.

Suzuki, K. Overview of deep learning in medical imaging. *Radiological physics and technology*, 10(3):257–273, 2017.

Tews, M. J., Jackson, K., Ramsay, C., and Michel, J. W. Fun in the college classroom: Examining its nature and relationship with student engagement. *College Teaching*, 63(1):16–26, 2015. doi: 10.1080/87567555.2014.972318. URL https://doi.org/10.1080/87567555.2014.972318.

Tews, M. J., Michel, J. W., and Noe, R. A. Does fun promote learning? the relationship between fun in the workplace and informal learning. *Journal of Vocational Behavior*, 98:46–55, 2017.