Human-Machine-Learner Interaction: The Best of Both Worlds

Eli T. Brown
DePaul University
Chicago, IL 60604
ebrown80@cdm.depaul.edu

Remco Chang
Tufts University
Medford, MA 02155
remco@cs.tufts.edu

Alex Endert
Georgia Institute of Technology
Atlanta, GA 30332
endert@gatech.edu

Abstract

This position statement presents a notional framework for more tightly integrating interactive visual systems with machine learning. We posit that increasingly, powerful systems will be built for data analysis and consumer use that leverage the best of both human insight and raw computing power by effectively integrating machine learning and human interaction. We note some existing contributions to this space and provide a framework that organizes existing efforts and illuminates future endeavors by suggesting the categories of machine learning algorithm and interaction type that are most germane to this integration.

Author Keywords

visual analytics; machine learning; interactive machine learning; sensemaking; data science;

ACM Classification Keywords

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

Introduction

The problem of the increasing volume and complexity of data is well known, and is an active research area across many disciplines, including machine learning and data visualization. In general, machine learning offers sophisticated algorithms for building mathematical models from

data. These techniques have proven effective, especially in the presence of ample training data, but may not leave the data stakeholder with any better data understanding. Alternatively, data visualization focuses on presenting interactive views of data to people in order to enable sensemaking and building of mental models for decision making and insight. The analytic process fostered by these tools is based on incrementally forming knowledge and understanding in the person exploring the data, primarily through interacting with the visualizations to answer questions, see different views, and gain insights [13]. Thus, building visual analytic tools that combine machine learning and interactive visualization techniques in a calculated way offers great potential to help people make sense of data. At the same time, consumer software and electronics increasingly advertise features that depend on machine learning, from self-parking cars to cell phones apps that alert users that when to leave for the airport to catch a flight. The fundamental understanding of how generally to create these integrated human and machine-learning systems remains an important, yet open challenge.

The machine learning community offers sophisticated algorithms for building models and making predictions from data. These algorithms may involve complex, opaque parameters, and present themselves to data-stakeholders as black-boxes. Expertise in these algorithms is often outside the realm of the analysts. They may be glad not to have to understand how the result is achieved, but they may also not inherently trust a result handed to them from an opaque process. Further, many open-ended problems are beyond the scope of what a computer can handle automatically, especially for cases in which ground-truth, expertly-labeled data examples are scarce.

Human reasoning is better suited to those open-ended problems, where the tasks and questions shift and the desired pattern, query, or question is unknown apriori. The visual analytics community enables data stakeholders to take control of their own analytics process by providing techniques and systems that allow them to leverage their expertise to explore the data and discover patterns on their own. However, visual analytics systems generally do not quantify these insights in the form of computational models for explaining patterns and re-using on other data (i.e., little, if any, training of machine learning models is generally supported).

In this position statement, we advocate for tightly integrating human-centered interfaces and visual data exploration tools with data-driven machine learning techniques. A spectrum of approaches exists for this combination, and so we provide a framework for how interaction methods and machine learning algorithms can be combined. Specifically, we present the framework of Figure 1 and explain the relationships between interactions and machine learning that it characterizes. We provide examples of existing work to illustrate these relationships in practice. We then provide a brief discussion of a series of examples of human-and-machine research and characterize their position in the framework.

The Beginnings of a Framework
In this section, we sketch a framework for systems that integrate humans and machine learning as groundwork for future progress in this growing field. The remainder of this section provides a discussion of our framework, diagrammed in Figure 1, and explains the types of interactions and machine learning techniques that can be used, with examples of how they have been integrated in the past.

Aside from the data and the potential output of the model
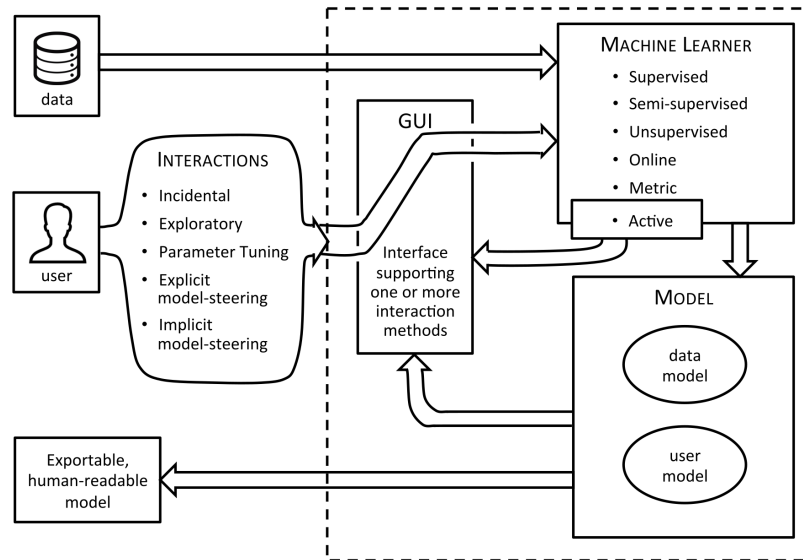
Figure 1: This diagram sketches a framework for interactive model-learning systems. It includes types of learning algorithm and types of interaction as described in in this paper.

learning, the framework can be seen as two interacting components: the user and the system. The user interacts with the system through a set of curated interactions with a graphical user interface (GUI). Within the system, the flow is cyclical, as new information flows through the GUI to model-learning and is then reflected back in the GUI. The interactions of the user with the GUI are translated into information for the machine learning algorithm, which updates the internal models. The updated model is then reflected in the GUI, allowing further interaction from the user. This can happen in real time or with on-demand iteration. It can be based on explicit intent by the user to improve the model, or the user can reap the benefits while being completely unaware. For an explanation of the types of interactions included in the diagram, which are *incidental, exploratory, parameter tuning, explicit model steering and implicit model steering*, see the sidebar.

There are two types of models: (1) models of the data constitute the analytical results and affect how data is displayed as well as the final analytic product, and (2) models of users can be learned from their interactions and used to optimize the interface for that type of user (e.g. a future application of Brown et al. [5]). Once a set of interactions has been used to update one or both of the models, the GUI can be updated to reflect the new model information, presenting a new opportunity for feedback from the user.

In the machine learning component, we list the top-level

categories of machine learning algorithm (*supervised, semi-supervised* and *unsupervised*) as well as call attention to specialties that may be especially relevant (*online, metric* and *active learning*). Any of the different types of machine learning could potentially be used in an interactive context. Although unsupervised learning does not allow for any information about the data by definition, it can benefit from *parameter tuning*. Direct parameter tuning requires some expertise in the algorithm, as in iPCA [10], which allows a users to manipulate principle component analysis visually. Without user expertise, unsupervised algorithms can be used to cluster data either as part of its presentation, or in order to strategically display only subsets (e.g. WireVis [6]).

Supervised machine learning algorithms can be applied to a wide variety of systems, whether for directly working with a user's data, or working behind the scenes to improve performance. Because classifiers are already so broadly applied in data analysis, and produce important but sometimes complex results, one way that visualization cooperates with classifiers is to help visualize their output. Alsallakh et el. proposed a set of visualizations that help analyze the performance of classifiers that provide classification confidence scores for multiple classes of outcome [1]. *Parameter tuning* interactions can be useful in helping analysts get the most out of supervised algorithms. Gleicher [9] presented Explainers that allows a user to tune readability vs. accuracy tradeoffs in SVMs. The Baobab-View system by van den Elzen and van Wijk [14] allows the user to semi-automatically refine a decision tree. Finally, Muhlbacher et al. proposed a tool for generating regression models using a partition-based approach [12]. There are also numerous behind-the-scenes possibilities for these algorithms. It is important to note that for both *implicit* and *explicit model steering*, using a supervised algorithm re-

quires caution to ensure that the assumption of having all the needed labels is appropriate. One example of how to do this is ReGroup [2], in which a Naive Bayes classifier is continually used to help people build groups of their Facebook friends. In [5], the authors used supervised machine learning on *incidental* and *exploratory* user interaction data to build a classifier that can predict the performance of a user on a visual search task. That work demonstrates the possibility of detecting facets of a user automatically from interactions, but does not close the loop by following the arrow that updates the user interface based on the model.

Semi-supervised machine learning includes the ability to take into account partial labeling, or side-information like constraints about what data points are similar. This is a powerful concept for applications where the user may (1) not know the final labels of any of the points and (2) be providing information on data points incrementally, which are common problems with both *explicit model steering* and *implicit model steering*. Metric learning algorithms are especially germane and have seen wide use because they can be used in a semi-supervised fashion and result in human-readable models. For example, prior work has shown how user interactions with projected data points can be used to build a distance function that models how important different data dimensions are to a user's mental model [4, 8]. Similarly, user interactions on a spatial grouping layout can be used to learn distance functions to make suggestions [3] or provide computational support for the visual grouping [7].

There have been many successful efforts to use machine learning interactively with human users. However, further work is needed to flesh out the full extent to which these integrated systems are possible and guidelines for how to build them. Innovation will come from both the HCI and machine learning fields.

Challenges

Making integration of human and machine learning a ubiquitous feature of software and analytics will requires that both disciplines address certain challenges. In this section, we briefly outline some problems that remain unsolved.

Challenges for Human-Computer Interaction
The most important challenge for the interactive component of mixed-initiative systems is that the interface or visualization reflects the model. If updates to the model cannot be seen, the feedback loop of the user's iterative improvement will be broken. Second, though the learner is responsible for the update time of the model, there may be significant computation required to update the GUI accordingly. The interactive component must ensure that calculations for updates to the visualization are part of the limited time-frame available for showing the user the result of feedback at interactive speed. Finally, there must be some mechanism in the interface for collecting information that can be useful to the machine learning back-end. There are many forms from raw interaction logs to specific buttons for labeling data points. Perhaps among the most subtle challenges is that in cases where the user is not explicitly improving the model, i.e. the model is used to improve the user experience but the user is unaware of the model, the GUI must facilitate, even encourage, providing feedback without disrupting the main use of the software.

Challenges for Machine Learning Techniques
This framework discusses several types of machine learning algorithms; in fact the machine learning community has produced myriad algorithms to choose from for creating integrated systems. However, in terms of supporting interactive learning systems, there are challenges yet to be fully met. The most important challenge is that interactive systems must support short round-trip times between user

feedback and newly learned models. Many machine learning methods are built to offer improvements in accuracy and may not run quickly enough for an interactive environment. The second challenge relates to the first – one way of creating algorithms that respond quickly is to formulate the algorithms to accept incremental information. A focus on *online* algorithms that get strong model-learning results one piece of feedback at a time fit this context perfectly. Finally, it is important in many use cases of integrated systems that humans be able not only to build a model interactively, but to understand the result (the topic of research in *comprehensible* machine learning [11]). Ideally, when the user is finished, the model can be exported as a product of the analysis and can itself be instructive. This is not a requirement for the learning process, but to get the best analytical results out of a human-computer partnership.

Working Together
Overall, it is important that the visualization and machine learning be chosen together so that they match up on data types and they can communicate. To ensure trust and interpretability, the models learned and the visualizations of data should be closely coupled. This means researchers across the fields must collaborate so the best of both disciplines can leverage the best of both worlds for the user.

Conclusion

As the size and complexity of data continue to increase, utilizing machine learning techniques in visual analytics systems will be ever more critical. Consumers will expect more and more automation in the software that controls their everyday devices. However, there has yet to be a systematic guideline for integrating user interfaces and interaction techniques with back-end machine learning techniques. This position statement presents a notional framework that categorizes interaction types and explains how they fit with a

variety of machine learning algorithms for this interactive context. Further, we present open challenges to both the machine learning and human-computer interaction communities to spur further research directions.

Acknowledgments

References

[1] B. Alsallakh, A. Hanbury, H. Hauser, S. Miksch, and A. Rauber. Visual methods for analyzing probabilistic classification data. *Visualization and Computer Graphics, IEEE Transactions on*, 20(12):1703–1712, 2014.

[2] S. Amershi, J. Fogarty, and D. Weld. Regroup: Interactive machine learning for on-demand group creation in social networks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 21–30. ACM, 2012.

[3] S. Basu, S. M. Drucker, and H. Lu. Assisting Users with Clustering Tasks by Combining Metric Learning and Classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 394–400, 2010.

[4] E. T. Brown, J. Liu, C. E. Brodley, and R. Chang. Disfunction: Learning distance functions interactively. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology (VAST)*, pages 83–92. IEEE, 2012.

[5] E. T. Brown, A. Ottley, H. Zhao, Q. Lin, R. Souvenir, A. Endert, and R. Chang. Finding waldo: Learning about users from their interactions. 2014.

[6] R. Chang, A. Lee, M. Ghoniem, R. Kosara, W. Ribarsky, J. Yang, E. Suma, C. Ziemkiewicz, D. Kern, and A. Sudjianto. Scalable and interactive visual analysis of financial wire transactions for fraud detection. *Information visualization*, 7(1):63–76, 2008.

[7] A. Endert, P. Fiaux, and C. North. Semantic interaction for visual text analytics. In *Proceedings of the 2012 ACM Annual Conference on Human Factors in Computing Systems (CHI)*, pages 473–482. ACM, 2012.

[8] A. Endert, C. Han, D. Maiti, L. House, S. Leman, and C. North. Observation-level interaction with statistical models for visual analytics. In *Visual Analytics Science and Technology (VAST), 2011 IEEE Conference on*, pages 121–130. IEEE, 2011.

[9] M. Gleicher. Explainers: Expert explorations with crafted projections. *Visualization and Computer Graphics, IEEE Transactions on*, 19(12):2042–2051, 2013.

[10] D. Jeong, C. Ziemkiewicz, B. Fisher, W. Ribarsky, and R. Chang. iPCA: An interactive system for PCA-based visual analytics. *Computer Graphics Forum*, pages 767–774, 2009.

[11] B. Letham, C. Rudin, T. H. McCormick, and D. Madigan. Building interpretable classifiers with rules using bayesian analysis. *Department of Statistics Technical Report tr609, University of Washington*, 2012.

[12] T. Muhlbacher and H. Piringer. A partition-based framework for building and validating regression models. *Visualization and Computer Graphics, IEEE Transactions on*, 19(12):1962–1971, 2013.

[13] P. Pirolli and S. Card. The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In *Proceedings of International Conference on Intelligence Analysis*, volume 5, pages 2–4, 2005.

[14] S. van den Elzen and J. J. van Wijk. Baobabview: Interactive construction and analysis of decision trees. In *Visual Analytics Science and Technology (VAST), 2011 IEEE Conference on*, pages 151–160. IEEE, 2011.