
Supporting User Interaction with Machine Learning through Interactive Visualizations

Jules François

SIAT - Simon Fraser University
250 -13450 102 Avenue, Surrey,
BC, V3T 0A3 CANADA
jfrancoi@sfu.ca

Frédéric Bevilacqua

SMTS Lab, Ircam-CNRS-UPMC
1, Place Igor Stravinsky 75004
Paris, France.
bevilacqua@ircam.fr

Thecla Schiphorst

SIAT - Simon Fraser University
250 -13450 102 Avenue, Surrey,
BC, V3T 0A3 CANADA
thecla@sfu.ca

Abstract

This paper discusses novel visualizations that expose the behavior and internal values of machine learning models rather than their sole results. Interactive visualizations have the potential to shift the perception of machine learning models from black-box processes to transparent artifacts that can be experienced and crafted. We discuss how they can reveal the affordances of different techniques, and how they could lead to a deeper understanding of the underlying algorithms. We describe a proof-of-concept application to visualize and manipulate Hidden Markov Models, that provides a ground for a broader discussion on the potentials and challenges of interactive visualizations in human-centered machine learning.

Author Keywords

User-centered Machine Learning; Visualization; Hidden Markov Models (HMM); Direct Manipulation.

ACM Classification Keywords

H.5.m [Information interfaces and presentation (e.g., HCI)]: Miscellaneous; I.2.m [Artificial intelligence]: Miscellaneous

Introduction

While consumer technology for capturing human movement are rapidly developing, designing interaction for movement remains challenging, because of the gap between sensor

data and meaning-making. This project belongs to a broader research agenda focusing on the use of machine learning for creative applications involving movement-based interaction, such as performing arts, new interfaces for musical expression [4, 7], movement sonification [6] and visualization, or gaming. In particular, designing for expression and creative applications requires experience and exploration of movement to bridge between the raw data and the potential for meaningful gestural languages (as a mode of 'experience design'). Machine learning provides a set of powerful tools for user-centered design, where the interaction process can be crafted by demonstration of the target behaviors rather than by programming a predefined set of rules. In these contexts, users might be designers, artists, instruments builders or hackers, who might not necessarily have background in machine learning theory.

The framework of Interactive Machine Learning, proposed by Fails and Olsen [3], introduced a fluid interaction workflow that integrates users at all steps of the process, from providing training data to training and evaluating machine learning models. Since then, the line of research of Human-center Machine Learning (HCML) have focused on improving user interaction with machine learning systems, using for example user feedback on target concepts in recommender systems [5], better programming environments [9, 4], visualization [10, 1], or pedagogical approaches that explain the system's decisions to users [8].

While this user-centered perspective on machine learning improved the usability of machine learning for end users, the underlying techniques and models often remain used as black-box systems where users get feedback only of the result of the inference process. This lack of transparency can be critically limiting for users without background in machine learning theory.

This paper discusses the potential and challenges of interactive visualization in machine learning for movement interaction design. We propose to extend the current feedback strategies to integrate visualizations of the internal values and behavior of machine learning models. Furthermore, we advocate *interactive* visualizations as manipulable artifacts that allow for both perceiving and crafting machine learning models. The power of user-centered machine learning resides in a bi-directional relationship between users and machines, a shared 'intelligence' that users construct by *learning to use models that learn*. We speculate that interactive visualization exposing the affordances of different models can help users build a deeper and more embodied understanding of machine learning's underlying mechanisms, eventually increasing their expertise, efficiency and expressivity.

User Interaction with Machine Learning

User-centered approaches to machine learning that emphasize a fluid workflow can allow novice users to design classifiers or regression models efficiently. For example, the Wekinator [4] encourages iterative design and multiple alternatives through an interaction loop articulating configuration of the learning problem (selection of features and algorithm), creation/editing of the training examples, training, and evaluation. Through several studies, Fiebrink et al. showed that users consistently iterate over designs, analyzing errors and refining the training data and algorithms at each step [4].

Users generally get three types of feedback on their design: a feedback on the training data, the results of the training process (e.g. cross-validation accuracy), and the results of direct interaction with a trained model (direct evaluation). While evaluation methods such as cross-validation are widespread in the machine learning community, Fiebrink

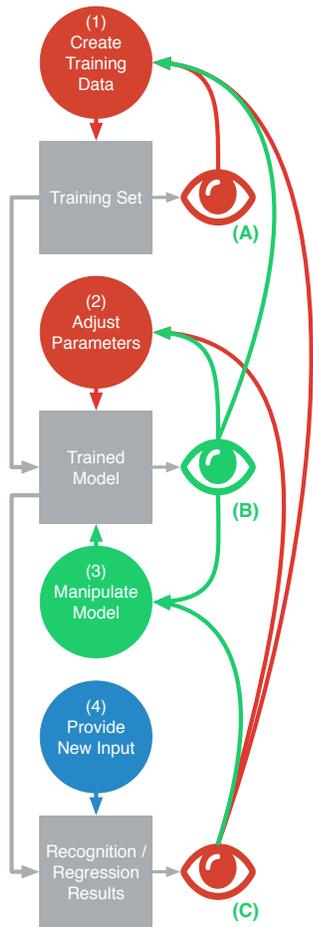


Figure 1: Proposed workflow of user interaction with machine learning. The colors blue, red and green represent the available user actions in conventional machine learning, interactive machine learning, and with the present proposal, respectively.

reported that users often prefer direct evaluation [4]. As a result, most of the feedback given to the user is result-oriented: the model remains a black-box and is understood through its output actions to new input.

Alternatively to result-oriented feedback, Kulesza et al. proposed “explanatory debugging” [8], an iterative approach where explanations of the system’s prediction are given to the user who, in return, can explain the system the necessary corrections. They showed that this approach resulted in both increased user understanding of the system and higher classification accuracy. This approach is powerful for discrete decision systems, such as ‘ElucidBug’, the text classification application evaluated in [8]. Nonetheless, providing explanation for more complex system for continuous interaction remains challenging, and visualization represents a promising solution to give insights on the internal values and behavior of the models. Furthermore, can we propose not only increased analytical understanding and higher classification accuracy but also creative and explorative solutions?

Interactive Visualizations

Users’ ability to improve a given design relies on their understanding of the behavior of the model. As a result, there is a critical need for improved feedback to the user’s design choices. We propose to use visualization at all steps of the interaction process: editing of the training data, training and adjustment of the model, and during interaction.

Visualizing Training Sets

Small training sets are common in machine learning for user-centered interaction design. Therefore, the quality of the examples significantly impacts the quality of the resulting model. Visualizing the training data can be very informative and support users in the design of the examples

(Figure 1, (A)). Many standard visualization techniques can be applied the visualization of the training set, that often consists of discrete data points or time series.

Visualizing Results during Interaction

Direct evaluation during interaction is often the preferred feedback modality: user can analyze a model in real-time by observing the system’s response to new input. In many cases, these predictions are mapped to the target interaction modality, such as audio or visual processing, and the model can be evaluated through the interaction itself.

Several authors proposed to visualize the results of the inference process (Figure 1, (C)). For example, *EnsembleMatrix* [10] displays the result of the classification process for different classifiers through confusion matrices. While such visualization allows users to compare the results obtained with different models, it might not inform them on *why* a given model works or not, and *how* to improve it. Model-Tracker [1] allows for monitoring performance in machine learning tasks during model building. Visualizations of classification errors over the training and test sets allow users to identify common errors related to flaws in the training set of data representation.

Because understanding the model’s behavior is essential to the gain of expertise of novice users with machine learning, we propose to use visualizations that follow an explanatory approach and expose *how* the model makes predictions. For example, Gesture Variation Follower (GVF) [2] allows for visualizing the particle distribution at each time step during recognition. Similarly, in the Max/MuBu implementation of XMM [7], as well as in the GaussBox application described thereafter, we display during performance the likelihood of each class, along with the state probabilities of Hidden Markov Models (HMMs) and the temporal progression in the gesture.

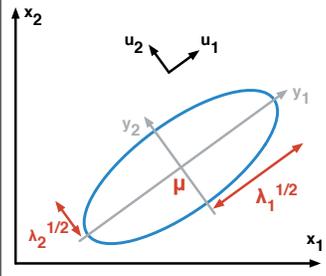
Visualizing Gaussians

Gaussian distributions can be visualized by plotting their contours of equal probabilities. For a multivariate normal distribution $\mathcal{N}(\mathbf{x}; \mu, \Sigma)$ of dimension n , these contours take the form of ellipsoids:

$$(\mathbf{x} - \mu) \Sigma^{-1} (\mathbf{x} - \mu)^T = \sum_{i=1}^n \frac{y_i^2}{\lambda_i^2}$$

where $y_i \triangleq \mathbf{u}_i^T (\mathbf{x} - \mu)$, and $(\lambda_i, \mathbf{u}_i)_{i=1}^n$ are the eigenvalues and eigenvectors of the covariance matrix Σ .

A multivariate Gaussian distribution can therefore be visualized by ellipses on a plane defined by two axes of the Gaussian. The ellipse is centered around the mean of the gaussian, and the angle and scale of the ellipse are related to the covariance.



Visualizing (Inside) Models

We described visualizations that emphasize the result of the interaction with the model. With such feedback, users build mental models of an algorithm's behavior as a black-box, from the input-output relationship inferred in direct evaluation. It is also essential to note that while in conventional use of machine learning, the classification accuracy might be the most common motivation, the ability to 'perform' the machine learning model might be more important from the user's perspective. Fiebrink noted that, in some cases, the shape and smoothness of the classification boundary defining the transition between several 'modes' was more important to the user than the actual recognition rate [4]. Therefore, visualization should inform not only on the results of the process but on the general attributes of the model, its boundaries, its potentials and limits.

We propose to improve the transparency of machine learning system through visualizations of the internal values of trained model (Figure 1, (B)). Such strategies have the potential to expose the behavior of particular models more clearly to users, bringing feedback on *how* the model is fitting data and making prediction. In the next section, we describe a prototype application for the visualization of Hidden Markov Models (HMMs) that exposes the values of each state in the model.

Interactive Visualizations and Model Affordances

Beyond the representation of machine learning results, visualizations have the potential to afford actions on the models themselves. We propose to define *interactive* visualizations as manipulable artifacts that allow for both perceiving and crafting machine learning models. They can facilitate a more fluid interaction between the user's actions, the outcomes of the model and the model itself. Ultimately, they reveal the *affordances* of a given model and enable users to

develop expertise and expressiveness with particular models that can be fine-tuned and even physically manipulated or sculpted.

GaussBox

GaussBox is a proof-of-concept application for interactive Hidden Markov Models (HMMs) visualization. It unveils the internal values of each states of a HMM through visualizations that respond to the user's actions: changes in training data, in model parameters, and direct manipulation.

Visualization and Direct Manipulation of HMMs

We propose to visualize HMMs with continuous observation distributions by projecting their states in two dimensions. The observation distributions at each state is modeled as a multivariate Gaussian distribution that can be displayed as a 95% confidence interval ellipse representing its mean and variance (see sidebar note). Each HMM chain is displayed as a set of ellipses on two given dimensions. Given the bijective relationship between the mean/covariance of a Gaussian distribution and its ellipse representation, the visualization straightforwardly allows for direct manipulation of the states' values. Changing the ellipse position modifies the mean of the Gaussian, and rotating or scaling the ellipse acts upon the covariance matrix. While understanding the meaning of the values of a covariance matrix is challenging, ellipses provide a meaningful visualization that allows for direct authoring.

Workflow

We designed Gaussbox to allow fast iteration between data creation and editing, and model training and visualization. The interface allows for recording and labeling examples, which are visualized as two-dimensional trajectories (Figure 3, B). For clarity, the application presented in this paper is constrained to two-dimensional gestures recorded with

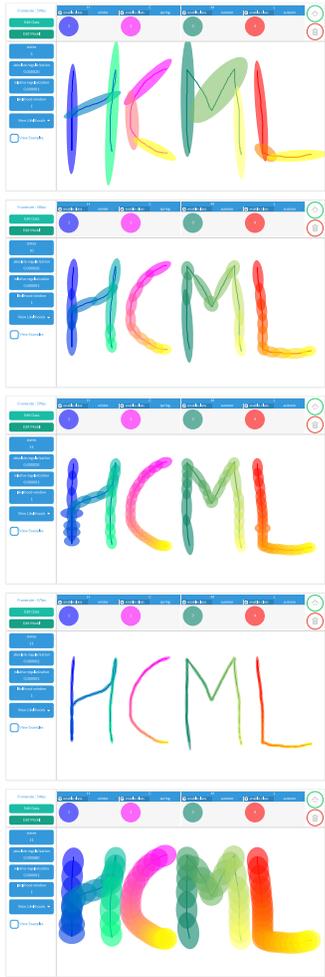


Figure 2: Screenshot of the Model Editor of GaussBox, for different values of the number of states of an HMM, and different level of regularization.

a mouse, where the user's actions can be co-located with visualizations of the training data and associated models in the graph area (Figure 3, A).

In editing mode, the models are visualized in overlay with the training examples. These visualizations are responsive to changes in the training set (e.g. if a training example is assigned to a new class) and to changes in the model parameters (Figure 3, C). The visualization also allows for direct manipulation: each ellipse is also editable in translation, rotation, and scaling, and each change is propagated to the underlying model. This gives users the possibility to visually craft, alter or fine-tune the trained models.

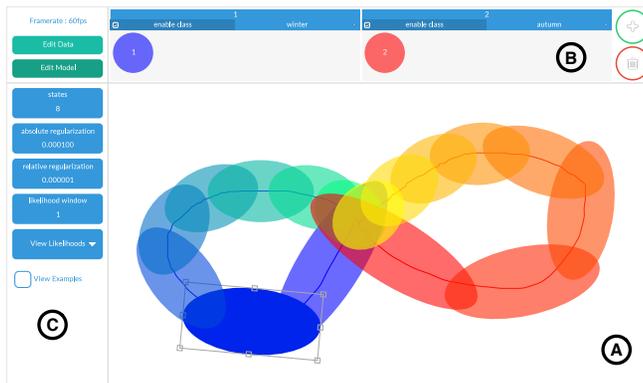


Figure 3: The GaussBox application. The top and left panels aggregate the different classes and display parameters, respectively. The main area is dedicated to the visualization and manipulation of both the trajectories of the training examples and the trained models

Finally, users can switch to 'performance' mode to evaluate the recognition of new input given the current models. In this case, we display not only the likelihood of each class, but also the activations of each state in the HMM (Figure 4).

This real-time visual feedback is very informative on how HMMs perform recognition: users can observe how probabilities propagate from one state to another, how these state probabilities give information on the temporal tracking of the gesture, or what is the influence of the temporal structure.

Discussion and Open Questions

This paper argues for the development of responsive and interactive visualizations to support user interaction with machine learning for interaction design. In particular, we propose to extend the current visualization strategies, that focus on the training data and the results of the recognition/regression, to integrate representations of the internal values and behaviors of models. Such visualizations have both a reflective and a highly pedagogical potential: they can expose the mechanisms underlying a particular algorithm in order to increase users' understanding, efficiency, and expressivity with machine learning. Moreover, we emphasize the creative potential of interactive visualization that unveil models' affordances, ultimately constructing a more embodied experience, exploration and understanding of machine learning's underlying mechanisms.

We illustrated these concepts with a proof-of-concept application for interactive visualization and direct manipulation of Hidden Markov Models. The pedagogical potential of this application relies on visualizations of the internal states of the HMMs, their relationship with the training data and parameters, and their activation to new input that explains *how* the recognition works. The choice of 2D gestures is obviously a good pedagogical example because of the congruence between users' actions and the visualization. However, how meaningful and useful are these visualization with other types of data, with larger feature spaces, or with a high number of classes? These limitations draw attention to more general challenges in interactive visualization for

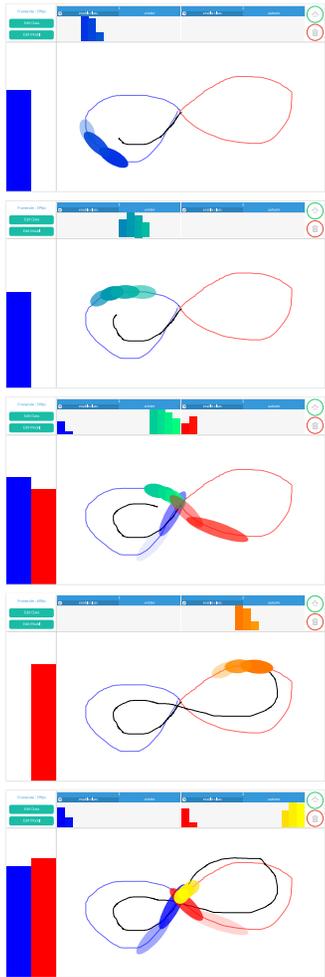


Figure 4: Real-time visualization of class likelihoods, state activations, and temporal progression during interaction.

human-centered machine learning:

- What should we visualize in various machine learning models?
- What are the affordances of different models and how to meaningfully manipulate them?
- Can such visualization build meaning of movement itself, beyond the system's design? What visualization can be standardized and which are specific?
- What is the creative potential of bi-directional visualizations?

This position paper will create a space to consider these questions and the multiple ways in which human-centered machine learning can provide new insights into embodied interaction.

References

- [1] Saleema Amershi, Max Chickering, Steven M Drucker, Bongshin Lee, Patrice Simard, and Jina Suh. 2015. ModelTracker: Redesigning Performance Analysis Tools for Machine Learning. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, 337–346.
- [2] Baptiste Caramiaux, Nicola Montecchio, Atau Tanaka, and Frédéric Bevilacqua. 2014. Adaptive Gesture Recognition with Variation Estimation for Interactive Systems. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 4, 4 (2014), 18:1—18:34.
- [3] Jerry Alan Fails and Dan R Olsen. 2003. Interactive machine learning. In *Proceedings of the 8th international conference on Intelligent user interfaces (IUI'03)*. 39–45.
- [4] Rebecca Fiebrink, Perry R. Cook, and Dan Trueman. 2011. Human model evaluation in interactive supervised learning. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'11)*. ACM, Vancouver, BC, Canada, 147.
- [5] James Fogarty, Desney Tan, Ashish Kapoor, and Simon Winder. 2008. CueFlick: Interactive Concept Learning in Image Search. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*. ACM, Florence, Italy, 29–38.
- [6] Jules François. 2015. *Motion-Sound Mapping by Demonstration*. PhD Dissertation. Université Pierre et Marie Curie.
- [7] Jules François, Norbert Schnell, Riccardo Borghesi, and Frédéric Bevilacqua. 2014. Probabilistic Models for Designing Motion and Sound Relationships. In *Proceedings of the 2014 International Conference on New Interfaces for Musical Expression (NIME'14)*. London, UK, 287–292.
- [8] Todd Kulesza, Margaret Burnett, Weng-keen Wong, and Simone Stumpf. 2015. Principles of Explanatory Debugging to Personalize Interactive Machine Learning. In *Proceedings of the 20th International Conference on Intelligent User Interfaces - IUI '15*. Atlanta, GA, USA., 126–137.
- [9] Kayur Patel, Naomi Bancroft, Steven M Drucker, James Fogarty, Andrew J Ko, and James Landay. 2010. Gestalt: integrated support for implementation and analysis in machine learning. In *Proceedings of the 23rd annual ACM symposium on User interface software and technology*. 37–46.
- [10] Justin Talbot, Bongshin Lee, Ashish Kapoor, and Desney S Tan. 2009. EnsembleMatrix: interactive visualization to support machine learning with multiple classifiers. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'09)*. Boston, USA, 1283–1292.