**HCC: Intelligent and Explainable Interactions for Interactive Model Steering**

**Overview**

Machine learning (ML) augmented sensemaking enables analysts to organize and reason about data relative to the complex patterns and relationships identified by the ML model. This project will study how visualization paired with intelligent interactions can serve as a medium for domain experts who lack strong technical backgrounds (non-expert users) to understand model perceptions and express steering feedback based on prior knowledge. Intelligent interactions allow users to provide domain specific feedback to shape the model embedding space to their data and sensemaking task. However, most existing feedback methods take cluster/class-centric approaches, often at the data-point level, that require substantial knowledge of the dataset or labeled information, and may not capture the complexity of the user's prior knowledge. Complementary to interactions, intuitive explanations are critical for enabling users to confidently interpret results, including the effects of interaction. However, comparative explanation of high-dimensional changes remains an open challenge. This project aims to address these gaps, by 1) eliciting, categorizing, and validating a **taxonomy of feedback tasks** to lay a foundation for intelligent interaction design, 2) designing and studying an **intelligent interaction paradigm** that enables interactions on higher-level information spaces (rather than point-level interaction) to elicit learnable feedback of varying granularities and structure, and 3) designing and studying **explanations that deeply integrate with interactions** to explain the current model and the effects of feedback, drawing in context from the higher-level information space and connecting it back to the data space.

**Intellectual Merit**

The proposed work presents advances in both visualization and HCI research. First, it will provide the first formalization of the space of feedback tasks, distilled into a task taxonomy, which sets the stage for designing intelligent interactions that meet the needs of users. We will establish this taxonomy through a synthesis of prior literature into a baseline taxonomy, followed by empirical experiments to study how users structure information when given a feedback task. Second, drawing on this taxonomy, we will establish a paradigm for intelligent interactions that elicits feedback in a higher-level feature space, providing a more intuitive space for reasoning and expressing information, and translates information back to the data space for learning. This requires: 1) mapping visual interactions into rich, structured information, 2) automatically identifying supporting evidence in the dataset to tune on, and 3) automatically tuning models to the feedback, without overfitting or forcing relationships not-backed by the data. By addressing these requirements, we not only advance the state of knowledge on interaction design for eliciting information, but also provide guidance on balancing cognitive overhead and expressivity for non-expert user interactions as well as balancing user intent with data evidence. Finally, we will integrate explanations directly with the interaction paradigm, re-connecting the higher-level feature space back to the data embedding space, and introduce comparative explanations that illustrate the impacts of user feedback, a space previously unexplored. In sum, these components lay the framework for interaction-first designs that enable non-expert users to effectively employ and steer ML models for sensemaking tasks.

**Broader Impacts**

This project will support non-expert users in employing ML-augmented sensemaking pipelines, which applies to a range of domains where experts need to analyze domain data but may lack formal training. We will publish polished versions of research outputs as open-source tools for community use. Additionally, project outputs can be directly integrated into educational materials to teach non-expert communities how to engage with and employ ML. We will integrate these outputs into undergraduate Data Science courses and host workshops to introduce interactive ML to middle school aged students, conducted as part of Tulane's GIST and BATS programs.