Research Statement

Human-Centered AI through Scalable Visual Data Analytics

Minsuk Kahng
Georgia Institute of Technology

While artificial intelligence (AI) has led to major breakthroughs in many domains, understanding machine learning (ML) models remains a fundamental challenge. These models are often used as "black boxes," which could be detrimental. How can we help people understand complex ML models, so that they can learn their inner-workings more easily, build them more effectively, and use them more confidently?

My research addresses these fundamental and practical challenges in AI, through a **human-centered** approach, by creating novel **data visualization** tools that are *scalable, interactive, and easy to learn and to use*. With such tools, ML models that perform well will be trusted and understood, and those that do not can be interpreted and improved.

Through my research in *visualization* and *data analytics* over the past 9 years, I realized that the key to solving this problem of understanding ML models is bringing the *human* into the analytics process. Specifically, through designing and developing novel interactive visualization tools, I create new ways for AI to explain its reasoning, and for human to engage in the sensemaking process of ML models by amplifying their cognition and perceptual capabilities. My thesis research focuses on the following three interrelated research topics.

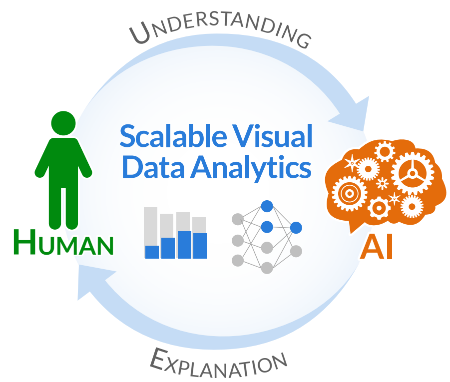

1. Visual Exploration of Complex, Industry-Scale ML Models

Despite the recent interest in visualization for deep learning models, the complexity of models and large-scale datasets used in industry pose unique design challenges. This motivated me to develop new scalable, interactive visualization tools for interpreting industry-scale deep learning models, ones that can work with a wide range of models and data, and can be readily deployed on machine learning platforms in industry.

ActiVis: Designing for Deep Learning Models at Facebook

I designed and developed a novel interactive visualization system, called **ActiVis**, for users to explore activation flows of complex deep neural network models (Figure 1). ActiVis, now **deployed** and **patent-pending**, was designed based on participatory design sessions with over 15 researchers and engineers across multiple teams at Facebook. I identified key design challenges in terms of model, data, and analytics patterns, and iteratively improved the ActiVis system based on their feedback. I presented ActiVis [1] at IEEE VIS'17, the top visualization conference; ActiVis was selected as one of the top four papers invited to present at SIGGRAPH'18.

ActiVis is the **first tool** that simultaneously supports instance- and subset-level exploration of deep neural network models. Both exploration strategies offer complementary analytical benefits: instance-level analysis shows how individual instances contribute to a model's accuracy, however it is tedious to inspect many instances one by one; subset-level analysis leverages input features to reveal relationships between input data attributes and ML algorithms' outputs. ActiVis provides a unified view for exploring data, both at instance- and subset-level — and is especially beneficial when dealing with large, industry-scale datasets,

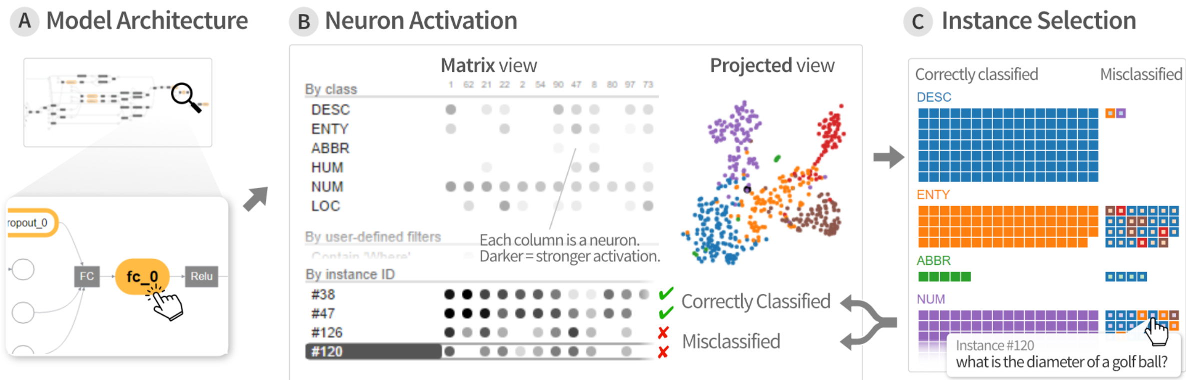

Figure 1: ActiVis's multiple coordinated views help engineers explore and analyze complex deep neural network models at both instance- and subset-level.

which may consist of millions or billions of data points. By exploring instance subsets and comparing them with individual instances, users can learn how models respond to many different slices within a dataset. Our case studies with Facebook engineers and data scientists indicate that ActiVis helps them for spot-checking models, exploring model architecture, and obtaining debugging hints.

Deployment. ActiVis has been **deployed on FBLearner**, Facebook's ML platform, used by more than 25 percent of Facebook's engineering team. Developers who want to use ActiVis for their models can easily do so by adding only a few lines of API calls, which instructs the models' training process to compute data needed for the visualization. ActiVis users at Facebook (e.g., data scientists) can then train models and use ActiVis via FBLearner's web interface, without writing any additional code. ActiVis is patent-pending.

Comprehensive Survey on Visualization for Deep Learning

Over the past few years, an increasing number of visualization tools for deep learning models, including ActiVis, have been developed. To provide a comprehensive overview of the literature to new researchers and practitioners in both visual analytics and deep learning, I co-authored the first survey article on visual analytics for deep learning [2], published in IEEE TVCG, the top visualization journal. We were invited to present the survey at IEEE VIS'18. Our survey describes the literature using a human-centered interrogative framework, focusing on the Five W's and How, and concludes by highlighting future research directions, helping people learn key aspects of this rapidly growing body of research.

2. Interactive Insight Discovery in ML Workflow

While visualization tools like ActiVis promote the understanding of a model itself, machine learning systems used in practice often consist of many different components (e.g., feature extraction, model selection). Thus, for researchers and practitioners who work in various stages of a ML pipeline, it is important to assist them in exploring how models process data in a ML workflow, and to help them discover insights.

ML Model Comparison via *MLCube*

One of the practical challenges in building ML models is to compare and select the best model among many candidates. While overall model accuracy can be used to select models, users often want to understand why and when a model would perform better than others, so that they can trust the model and know where it can confidently used. To help them analyze how models perform differently for certain groups of data instances, I developed **MLCube** [3], a novel visual exploration tool for comparing ML results by using *data cube* analysis. **Deployed by Facebook**, *MLCube* enables users to visually compare aggregate statistics over subsets of data instances and interactively drill down into models, helping them spot interesting patterns (Figure 2). For example, users may learn that *"this retrieval model performs better than the baseline model for teenage user groups, but poorly for recent newspaper articles,"* informing users' model selection process and leading to discovering insights for further improving the model. Our scalable system backed by Apache Spark has been deployed on Facebook's ML platform. In addition, *MLCube* has also influenced *TensorFlow Model Analysis*, an open-source library developed by Google and integrated into TensorFlow, the most popular deep learning library.

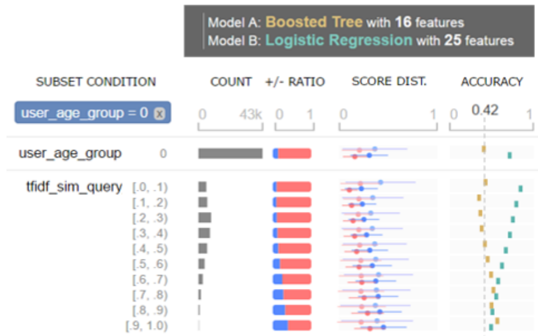

Figure 2: *MLCube* users can compare models by slicing and dicing datasets.

Making Sense of Raw Datasets in Databases

Data is a key ingredient that makes machine learning possible. Thus, before we extract features and prepare for training data, making sense of raw datasets is critical to building powerful ML models. To help data analysts rapidly develop deep understanding of unfamiliar datasets, I created the **ETable** system [4] for interactively exploring the complex relational structures of databases. *ETable*'s novel approach of constructing an extended table, based on *denormalization*, enables users to easily browse and navigate relational database tables. A controlled lab study with 12 participants demonstrated that users can construct complex queries much faster with *ETable* than a popular commercial database administration tool. I published this work at VLDB'16, the top conference in databases.

Besides relational databases, I have developed a number of interactive exploration tools for different types of data, including graph-structured datasets [5-6]. For instance, the *FACETS* system [5], which I developed with my colleagues, summarizes large graph data based on users' subjective interestingness for interactively exploring them, and was presented at SDM'16.

3. Interactive Understanding of Deep Learning Models

As deep learning continues to attract attention, many people, including practitioners without formal computer science backgrounds, want to learn this new technology. I have created interactive, visual educational tools to broaden these people's education access to AI, which is one of my long-term research goals.

GAN Lab: Interactive, Visual Experimentation for Learning GANs

The complexity of modern deep learning models makes it difficult for people to learn about them. In collaboration with Google Brain's People+AI Research (PAIR) group, I built and open-sourced **GAN Lab** [7], a visual, interactive tool for non-experts to actively learn *Generative Adversarial Networks (GANs)*, one of the most popular, but hard-to-understand deep learning models. GAN Lab enables users to interactively create GANs and try out different combinations of hyperparameters, to learn how the models learn to generate data distributions. GAN Lab's visualization techniques work in tandem to help crystalize complex concepts in GANs, such as through animations that explain the generation process (Figure 3). I presented this work at IEEE VIS'18, the top visualization venue.

ML Education using Web Browsers. A critical difficulty in deploying such educational tools is that training deep learning models conventionally requires significant computing resources. Deep learning libraries, like TensorFlow, typically run on dedicated servers, which are not designed to support low-latency computation needed for real-time interactive tools, or large number of concurrent user sessions through the web. I overcome these challenges by using *TensorFlow.js*, an in-browser GPU-accelerated deep learning library developed by my colleagues at Google. As GAN Lab runs locally on a user's web browser, anyone can access the tool using their browsers without the need for a specialized backend, significantly broadening people's education access to AI. I also developed an in-browser ML education module for K-12 students with Prof. Hal Abelson at MIT, which is being deployed as a part of Google's engEDU program.

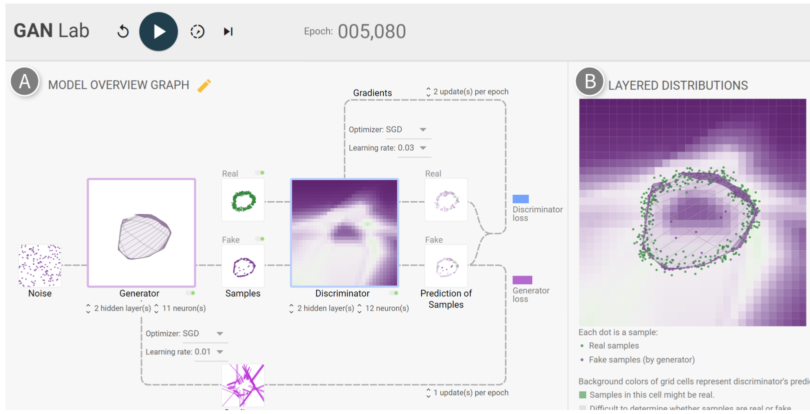

Figure 3: With GAN Lab, users can interactively train deep generative models and visualize their inner-workings.

Impact. I open-sourced and launched GAN Lab in September 2018. Within the first two months, more than 20,000 people from 120 countries tried it out. Its live system is available at https://poloclub.github.io/ganlab/.

Future Research Agenda

My long-standing research goal is to bring human-centered approaches to the field of AI and data science. To date, I have taken the first important steps toward this goal with my thesis research. For the road ahead, I hope to broaden and deepen this investigation, in close collaboration with my industry partners (Google, Facebook) and academic colleagues across disciplines. I will pursue this in the following four thrusts.

Interactive Model Building with Visual Guidance

While visualization tools promote people's understanding of ML models, they often do not directly support the important task of building well-performing models. This task is difficult even for experts, due to the countless combinations of parameters that need tuning. They often have to iteratively try different heuristics-based strategies to improve models, which is tedious and error-prone. To help them refine and debug their models, I plan to create new types of visualization tools that guide them to identify problems and discover actionable insights for debugging, by combining computational and interactive methods: scalable data mining techniques automate the discovery of problematic cases; interactive tools guide users to explore debugging strategies and make decisions. Through my Google PhD Fellowship, I have started discussing my ideas with Dr. Ed Chi, a principal scientist at Google AI.

Machine Learning for Everyone

A growing number of non-experts who do not write code not only want to learn AI, but also want to build machine learning models for their products and data. Large technology companies have started providing these users with web-based services for building ML models without any coding. However, it is very challenging to develop such interfaces that are easy for non-experts to learn and to use. I envision the next generation of interactive systems for these users to easily interact with ML models with the help of visual explanations to avoid the users to use ML as black-boxes. I believe this new accessible way of building ML models (e.g., GAN Lab) will broaden people's access to AI technologies and ensure their appropriate use.

Making AI Systems Fair

Researchers have recognized that results from ML models can be unfair. For example, a model may produce much worse accuracy for certain groups of people (e.g., people of color) than for the overall population. This can result in many problems, as AI-powered systems continue to make important decisions across social domains. I plan to build visualization tools that reveal unfairness in models and help people explore solutions to ensure the models' decisions are not discriminatory. I have started designing new tools with Prof. Jamie Morgenstern and her student at Georgia Tech. Recently, my work directly contributed to an NSF grant proposal and a Google Faculty Award proposal. Besides fairness, I will work in collaboration with researchers from different backgrounds, to make AI accountable, transparent, interpretable, and safe.

New Interaction Paradigms between Human and AI Systems

The advent of deep learning models is rapidly reshaping many existing data-driven systems that have a long history of research and development, such as database systems and search engines. For example, traditionally, web search engines take only short text queries, however, they now recognize long natural language queries, images, voice, and more. This has led us to rethink their interaction methods that involve both computational and usability challenges. For instance, when a system fails to recognize a user's voice query, we may want to design interfaces that actively solicit user feedback to iteratively refine the query. I plan to study new interaction paradigms between these data-driven intelligent systems and human, drawn from my experience working in both research areas and their intersection.

In summary, my research pushes the frontier of AI through human-centered approaches, contributing novel paradigms, methods, and tools that advance people's understanding of AI, accelerate its development, and increase their access to new technologies. With my experience and knowledge across interactive visualization, data science, and machine learning [1-10], I hope to accelerate innovation across these disciplines, making positive impacts to people's everyday lives and our society.

References

[1] **M. Kahng**, P. Andrews, A. Kalro, and D. H. Chau. ActiVis: Visual Exploration of Industry-Scale Deep Neural Network Models. *IEEE Transactions on Visualization & Computer Graphics (VAST)*, 2017. [1 of 4 top papers invited to SIGGRAPH'18; Deployed by Facebook]

[2] F. Hohman, **M. Kahng**, R. Pienta, and D. H. Chau. Visual Analytics in Deep Learning: An Interrogative Survey for the Next Frontiers. *IEEE Transactions on Visualization & Computer Graphics* (Early Access), 2018.

[3] **M. Kahng**, D. Fang, and D. H. Chau. Visual Exploration of Machine Learning Results using Data Cube Analysis. *Workshop on Human-in-the-Loop Data Analytics at SIGMOD*, 2016. [Deployed by Facebook]

[4] **M. Kahng**, S. Navathe, J. Stasko, and D. H. Chau. Interactive Browsing and Navigation in Relational Databases. *Proceedings of the VLDB Endowment*, 2016.

[5] R. Pienta, **M. Kahng**, Z. Lin, J. Vreeken, P. Partha, J. Abello, G. Parameswaran, and D. H. Chau. FACETS: Adaptive Local Graph Exploration. *SIAM Data Mining Conference (SDM)*, 2017.

[6] C. Stolper, **M. Kahng**, Z. Lin, F. Foerster, A. Goel, J. Stasko, and D. H. Chau. GLO-STIX: Graph-Level Operations for Specifying Techniques and Interactive eXploration. *IEEE Transactions on Visualization & Computer Graphics (InfoVis)*, 2014.

[7] **M. Kahng**, N. Thorat, D. H. Chau, F. Viégas, and M. Wattenberg. GAN Lab: Understanding Complex Deep Generative Models using Interactive Visual Experimentation. *IEEE Transactions on Visualization & Computer Graphics (VAST)*, 2018. [Open-sourced with Google AI]

[8] **M. Kahng** and S.-g. Lee. Exploiting Paths for Entity Search in RDF Graphs. *SIGIR (poster)*, 2012.

[9] **M. Kahng**, S. Lee, and S.-g. Lee. Ranking Objects by Following Paths in Entity-Relationship Graphs. *Ph.D. Workshop at ACM Conference on Information and Knowledge Management (CIKM)*, 2011. [Best Paper Award]

[10] **M. Kahng**, S. Lee, and S.-g. Lee. Ranking in Context-Aware Recommender Systems. *WWW (poster)*, 2011.