Research Statement

Diana S. Kim

This statement is about my research journey during my doctoral studies within the last five years (2017–2022). As a computer vision researcher and an art enthusiast, I developed machine learning applications for art and engendered new insights and high-profile questions in both domains. Through my interdisciplinary research, I was able to clearly recognize the frontline of machine learning, computer vision, and artificial intelligence (AI) and confirm the potential of machine learning as a new analytic tool for art. Finally, I set up the next agenda to develop stronger AI applications for various domains beyond art.

I explored interesting art problems through machine learning methodologies; (1) demonstrating how art is changed visually over time from the machine's perspective and how that is related to art history methodologies [1], (2) finding the principal semantics for style recognition [2], (3) laying the groundwork for computational iconography, i.e. finding the co-occurrence and visual similarities among the content of fine art paintings [3], (4) quantifying fine art paintings with a finite visual semantics from style through language models without using direct annotations [4, 5], and (5) defining visual factors of our emotional reactions to abstract paintings (current study, a result presented in Figure 1).

I am truly fascinated by these problems because they seek connections between art and science. In general, these disciplines are separate in terms of their methodologies or goals; however, bridging them is necessary to advance human knowledge [6]. I focused on machine learning's ability to efficiently find the hidden structure of art data and some problems in art history looking for patterns in numerous paintings. Through cooperating with art historians, I laid groundworks (two representative outcomes: a ground-truth set for evaluation and a set of semantics necessary to describe art visually[1]) to bring art problems into a scientific and computational framework and proposed novel machine learning methodologies for them.

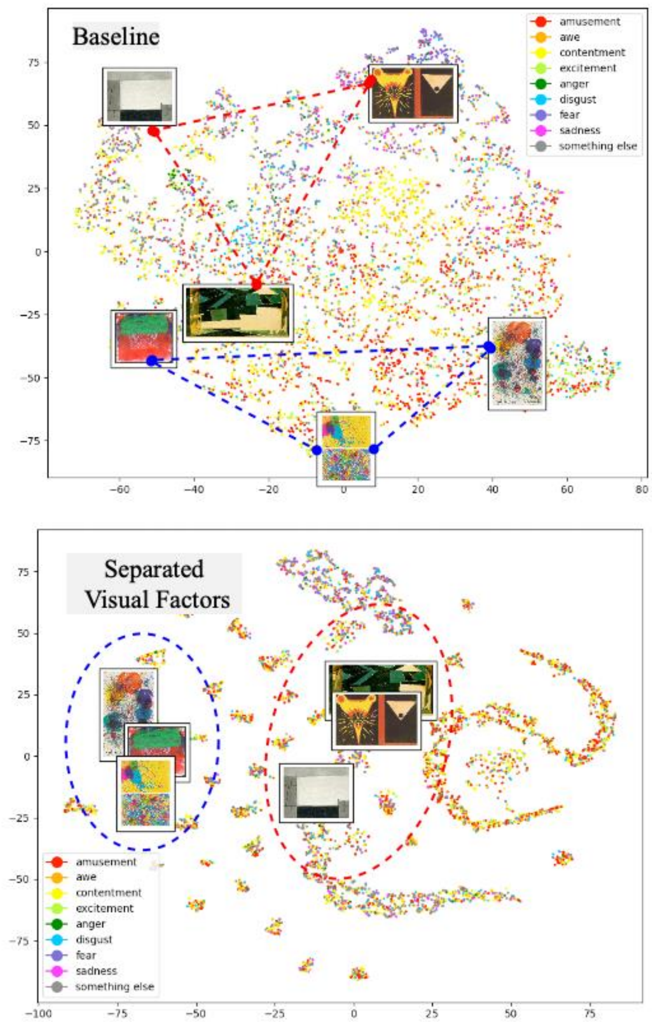

Figure 1 *"Defining Visual Factors to Emotional Reactions": Abstract artists focus on visual elements to create art with a strong emotional impact* [6]. *By using the external knowledge in brain science of bottom up and top-down processing of art appreciation, a new emotional classifier is developed (bottom figure), which generates discrete clusters according to a finite set of visual factors, while the baseline system was a continuous and entangled one to be interpreted.*

[1] Fifty-eight words were selected as a set of semantics to describe the visual form of paintings. One hundred and twenty paintings and their relevancies to the 58 visual words were annotated for quantitative evaluation. They are representative groundworks to implement machine learning frameworks for the art problems and quantitatively measure the systems' performance. (www.whitepage.space)

- **Review of My Research**

Two fundamental problems in art history–stylistic analysis (Figure 2 and Figure 4) and iconography (Figure 3)–were examined. My first research problem was to find a systematic relation between visual semantics and style. Using that relation I developed an automatic system to learn the visual information of the input paintings without using direct annotations in training [1] [2] [4, 5]. The secondary problem was to find the hidden connections or unexpected parallels between the themes, motifs, and visual styles of the works of art from different periods of time [3]. This iconography paper was outcome of a summer internship program with two undergraduate students I advised in 2018 as a research mentor. The paper defined the concept of computational iconography (CI) for the first time and explained why machine learning can be an efficient solution for CI. A system for CI can find co-occurrences and visual similarity relations among content in extensive corpus of paintings rapidly. For art historians, recognizing a repetitive motif in a number of paintings is an important task because it can provide reliable clues to the symbolic meanings of paintings' content. Improving the prototype CI system is one of my future research plans.

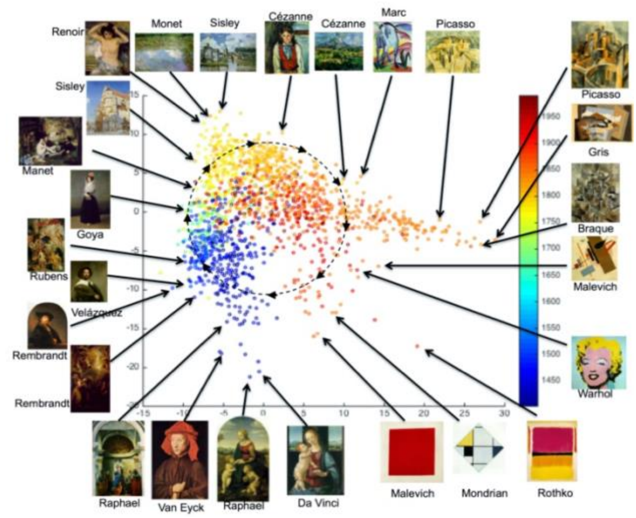

Figure 2 *"The shape of Art History in the Eyes of the Machine" [1]: A representation space that encodes painting chronical order is discovered and quantitively analyzed to see how the machine's representation is related to knowledge in art history.*

The problems are novel and have various intellectual merits. First, they address the fundamental problems of developing an AI system for art aligned with the first two principles of art, where the three primary aspects of understanding art are visual form (how it was made), content (the subject matter of art), and context (in what circumstances it was made) [7]. Second, the problems provide good research resources to study machine learning. The scarcity of annotations in art data and the complex depictions and compositions of paintings directly highlight the major problems in current machine learning study, such as the overreliance on data and the overfitting to the variations in style, texture, color, and the geometric deformations of shifting, transformation, and scaling. On the other hand, there are sufficient well-established theoretical knowledge about art and an extensive corpus of art-related text enable me to design a plausible architecture based on the prior knowledge and the well-defined theories of art.

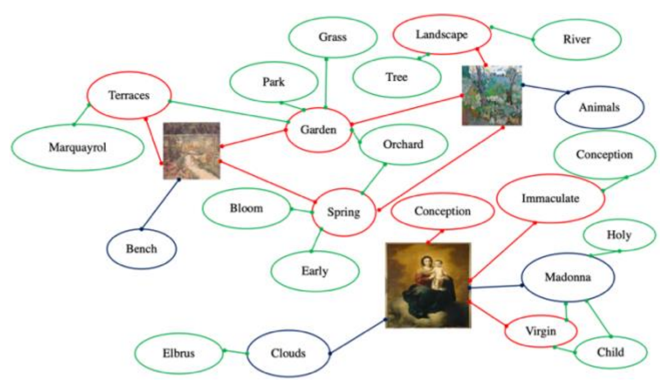

Figure 3 "*Computational Analysis of Content in Fine Art Paintings" [3]:co-occurrence and visual similarities among content of paintings are found. Based on the discovered connection, paintings are annotated by various words beyond their titles.*

Furthermore, I found another important merit of machine learning for the art problems. In the early stage of my research, I sought to implement efficient and intelligent art systems that can perform the same level of art as art historians. Recently, however, this goal proved to be somewhat limited in explaining the true benefits of machine learning applications for the art domain, although it helped me to develop a high-quality vision system. In my research, I witnessed how machine learning can reveal the hidden dimensions or underlying structures of art that have not been accessible to the naked eyes or other traditional methodologies. I discovered that machine learning can efficiently reopen the multiple layers embedded in a large corpus of art so enables us to better understand art in an unprecedented way.

In my work, I visually demonstrated how paintings change over time [1], computationally simulated the hierarchical relations between visual concepts and styles [4, 5], quantitatively measured which visual semantics are important for style recognition [2], and finally found two associative patterns among the content in paintings [3]. Now, I envision machine learning as a new mechanism to objectively analyze the artistic process in ways that are different from the conventional tools, while anticipating its broader advantages for new art problems in my future research.

- **My Aha Moment**

Art data are typically annotated with artist's information and style, but they do not come with visual semantics, content information, or bounding boxes indicating the objects' locations. Thus, conventional supervised methodologies were not applicable, so the problems were fundamentally challenging. However, as I observed a hidden space of a style classifier (deep-CNN), a thing I wondered was how smoothly the input paintings were visually changed in the space and there existed evident semantics driving the smooth variations, which could be named as gestural, planar, abstract, geometric, and so on. Without using direct visual annotations, the latent space learned to quantify the visual semantics, while the machine was trained for style classification only (Figure 4).

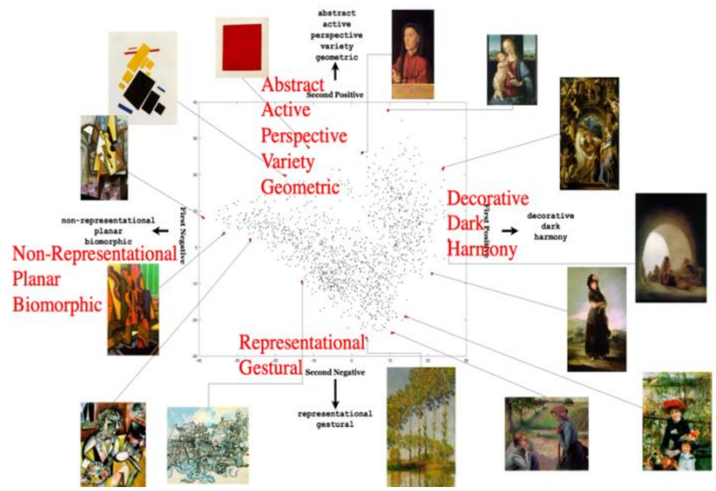

Figure 4 *"Finding Principal Semantics of Style in Art" [2]: The hidden space of a style classifier is interpreted by a set of visual concepts. They are proposed as the principal concepts (red-colored words) of style recognition.*

In the experiment, I found that an inherent structure of art that is embedded in art data, can be resonated, and revealed through a pre-defined architecture when its modeling reflects the original structure of the art well, even if it is only partial. Since the style classifier has an inherent hierarchical structure, it is synchronized with the hierarchical categorization process–synthesizing the form, elements, and qualities to determine the style later–so it can learn the visual forms and elements from its hidden layers without using direct annotations in training. That provided me with an important insight into how the architecture of machine learning can reduce the reliance on data through the integration with prior art knowledge. That idea became a cornerstone of my research.

- **Proxy Learning**

The idea was realized specifically in my most recent work, "Proxy Learning of Visual Concepts of Fine Art Paintings Through Language Models'' [4, 5], which is a more advanced development of my first paper [2]**.** This paper proposed a novel machine learning framework that can quantify fine art paintings with a set of visual semantics but that does not require any visual annotation in training (Figure 5).

To achieve the holy grail, a neural net was modularized, interconnected, and constrained based on general art knowledge and well-founded theories in cognitive science. Inspired by prototype theory in cognitive science, a general and hierarchical relationship between style and visual semantics was mathematically formulated in an imaginary conceptual space, where each of the axes was aligned with their semantics. The general relationship was estimated from a large collection of art-related texts by developing new language models using natural language processing (NLP). Finally, the pre-learned relation was transplanted into the last part of a style classifier (deep-CNN) in order to make the neural net be synchronized with the hierarchical relation between visual semantics and style classification. The neural net was then guided to learn the visual semantics in its last hidden layer through the constraint, while the machine was trained only for style classification.

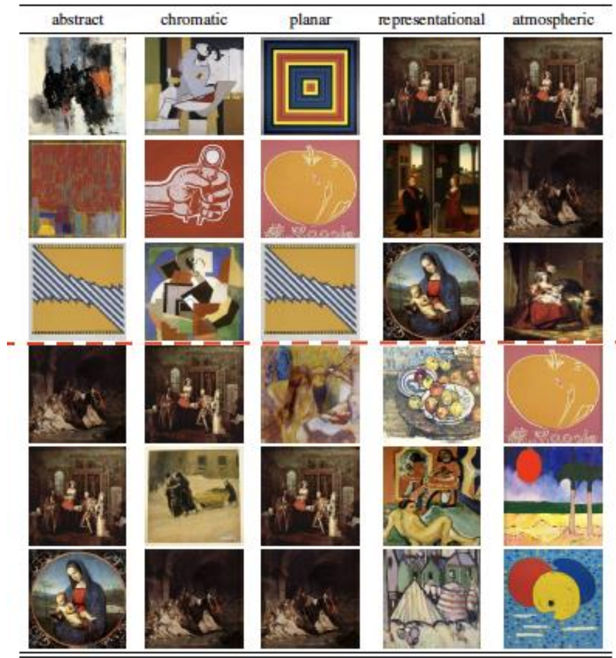

Figure 5 *"Proxy Learning of Visual Concepts of Fine Art Paintings from Styles Through Language Models [4, 5]: A neural net is trained to quantify the input paintings without using any direct annotation. The three most (1 - 3 rows) and three least (4 – 6 rows) relevant paintings are shown as the proposed system predicted.*

The main features of proxy learning are (1) functional modularity and (2) constrained and explainable architecture by prior knowledge. In proxy learning, a whole neural net is modularized, and each compartment has a certain functional meaning along the way toward the final style recognition layer. The pre-learned general relation between style and visual semantics is integrated with the system as a constraint and a bottleneck to force the system to learn visual information from its last hidden layer. The two features are important because they accomplished two important conditions for development of a stronger AI system with good generalization: (1) less reliance on data and (2) more transparent interpretability.

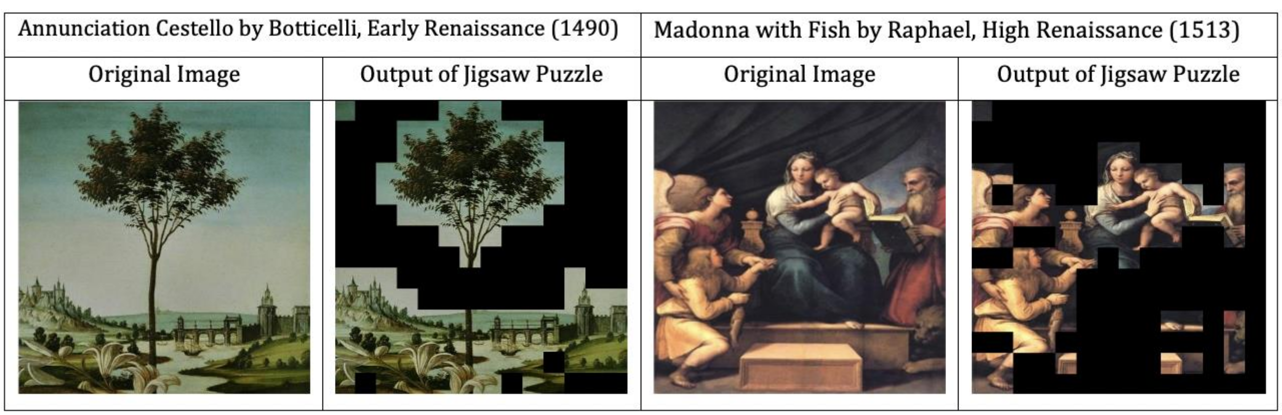

Figure 6 *A future Computational Iconography (CI) study named "Jigsaw Puzzle": Locating objects is fundamental to develop a CI system, but art data does not provide bounding box annotations for training. A prototype system is developed to capture the salient objects in paintings like capturing a moment of doing jigsaw puzzle.*

- **Future Research Direction**

In the near future (2022-2025), I will revisit the art problems that I have investigated. Some points can be improved and expanded, such as developing a new CI system (Figure 6), defining visual factors of our emotional reactions to abstract paintings, and improving the language models of proxy learning from linear to non-linear settings. As for the long-term direction of my research (2022-2030), I will investigate the ways of developing explainable machine learning frameworks that will be less reliant on data through highly modularized architectures based on the external knowledge of a targeted domain.

I believe that designing a well-defined structure or an architecture is as important as acquiring enough data to build a strong AI system. Due to raw data's multidimensionality, several functional regularities can potentially be embedded, but many of them could be irrelevant to the true target structure. In this sense, I am somewhat skeptical about machine learning trends that simply increase the complexity of systems to achieve better performance rates with predefined data sets on specific tasks, without sufficiently contemplating good architectures that reflect the true operation of a targeted task. That increased complexity is only valid when the data draws the real structure of an original problem in a very high dimensional space. However, a pre-defined data set does not have that fine granularity in general. Hence, I will pursue computer vision systems that have high modularity and interconnectivity aligned with well-founded common knowledge and achieve high generalization accuracy with less reliance on labeled data. Based on that direction, my research laboratory will focus on the following areas:

(1) Computer vision and machine learning applications will be the primary research fields. The research subjects will be broad to include pure 2D computer vision problems for static (image) and dynamic (video) images and to cover crossing boundaries problems in the humanities and the natural science of biology, physics, and chemistry, and so on. Through active collaborations with scholars in various fields, I will expand the research spectrum of my research, emphasize the practical merits of machine learning, and create various research opportunities for students. The relevant areas will be listed as digital signal and image processing (spectral theory), probabilistic modeling and learning, machine learning theory and applications, computer vision, AI, and NLP.

(2) The main direction of the research will be developing technical schemes to integrate well-founded theories or external knowledge into the designs of the machine learning framework's architectures or training procedures in order to make the system explainable, reduce its reliance on data, and finally improve its generalization. The machine learning system's architecture will be modularized and interconnected according to functionality and physically or statistically constrained based on referenced knowledge.

(3) NLP study will be one of the major topics. In particular, my laboratory will develop computational methodologies to extract general knowledge from the texts related to a targeted domain and explore the ways to efficiently represent the discovered knowledge. I see an extensive corpus of texts as a valuable, scalable, and objective resource to learn or extract knowledge independent from human annotations or theories from scholars with respect to contents and data formats. This NLP study aims to learn general knowledge that is directly applicable and compatible with existing systems or the frameworks under development.

(4) The spectral neural net will be explored to build computer vision that is robust to various geometric deformations of the targeted objects or words in 2D images or in 1D sentences. This plan is motivated by my previous research problem of computational iconography. Due to the high complexity of the content depictions and compositions in paintings, learning features invariant geometric deformation (scaling,

transform, and shifting) was very critical to achieve excellence in generalization. However, conventional vision systems based on the spatial domain were not efficient in achieving those goals.

(5) Using machine learning, my research laboratory will work on developing algorithms to reveal a hidden structure underlying a corpus of art. I will contiguously test what machine learning can discover about art and investigate how it is related to art history knowledge. On a related note, I discovered a platform called Generist Map, the result of a collaboration between the Metropolitan Museum of Art (the Met) and Microsoft, that shows that machine learning is able to reveal the process of artistic creation for the vases or ewers in the Met's collections.

(6) I believe a well-designed computational algorithm can be creative. Even though a machine follows a set of instructions, we cannot foresee how it will realize its outcome and the direction it will take. Through active collaborations with artists, my research laboratory will work on creating new forms of art, literature, music, and painting.

(7) Acquiring research funding from external sources will be helpful to support the students financially. the relevant funding resources based on my future research plans include but are not limited the ones listed below.

Funding Resources Relevant to Future Research Areas	
Faculty in Early Career	▪ Faculty Early Career Development Program (CAREER) ▪ Computer and Information Science and Engineering Research Initiation Initiative (CRII) ▪ Solan Research Fellowships
Knowledge Discovery	▪ IIS: Information Integration and Informatics (III)
Robust Intelligence	▪ IIS: Robust Intelligence
Explainable AI	▪ Explainable AI for decoding and modulating neural circuit activity Linked to Behavior

- **Closing Statement**

I believe that the best part of conducting research is that it raises new questions and problems to be explored next. Along the way to finding answers to the new questions, I can develop a better understanding of the subject matter, enjoy many opportunities to collaborate with prominent scholars, and encounter more ideas and problems beyond my imaginations, and use them for a new round of research. Understanding new scholarly works and proposing new schemes to fill the gaps between state-of-the-art systems and my goals have been the true joys of learning and contributing for me.

From my previous research, I gained a solid understanding of the frontline of machine learning as a foundational framework for building AI systems and new analytic tools for art. Furthermore, the challenges of the current state of the art systems due to the scarcity of art data led me to look for external knowledge of art problems from art history or cognitive science and apply them to designing the architectures of machine learning systems. Finally, the next problems that I look forward to working on with my future students and colleagues are developing (1) machine learning systems to reveal hidden structure for art, and (2) vision systems with high modularity and interconnectivity based on well-founded common knowledge and discovered information from NLP. However, I know that they might not be limited to the descriptions I presented here. My vision and ideas will evolve and expand along with future investigations.

Bibliography

[1] A. Elgammal, B. Liu, **D. Kim,** M. Elhoseiny and M. Mazzone, "The Shape of Art History in the Eyes of the Machine," in *In Thirty-Second AAAI Conference on Artificial Intelligence (AAAI)*, 2018.

[2] **D. Kim**, B. Liu, A. Elgammal and M. Mazzone, "Finding Principal Semantics of Style in Art," in *IEEE 12th International Conference on Semantic Computing (ICSC)*, 2018.

[3] **D. Kim**, J. Xu, A. Elgammal and M. Mazzone, "Computational Analysis of Content in Fine Art Paintings," in *In 10th International Conference on Computational Creativity (ICCC)*, 2019.

[4] **D. Kim**, A. Elgammal and M. Mazzone, "Proxy Learning of Visual Concepts of Fine Art Paintings from Styles through Language Models," in *In 36th AAAI Conference on Artificial Intelligence (AAAI)*, 2022.

[5] **D. Kim,** A. Elgammal and M. Mazzone, "Arxiv (2201.01819): Formal Analysis of Art: Proxy Learning of Visual Concepts from Style Through Language Models," 05 January 2022. [Online]. Available: https://arxiv.org/abs/2201.01819.

[6] E. R. Kandel, Reductionism In Art and Brain Science, New York: Columbia Univeristy Press, 2016.

[7] J. C. Van Dyke, Principles of Art, Fords, Howard, & Hulbert, 1887.

[8] W. Fenstermaker, "How Artificial Intelligence Sees Art History," 4 Feb 2019. [Online]. Available: https://www.metmuseum.org/perspectives/articles/2019/2/artificial-intelligence-machine-learning-art-authorship. [Accessed 10 11 2021].