INAUGURAL SYMPOSIUM

NOVEMBER 19, 2016 | 10:30 - 5:30 PM
COLUMBIA UNIVERSITY
MUDD 633 | SEELEY W. MUDD BUILDING

Institute
For
Visual
Intelligence

Institute for Visual Intelligence Presents Its

INAUGURAL SYMPOSIUM

SAT, NOV 19, 2016: 10:30 - 5:30 PM

PHILOSOPHICAL UNDERSTANDING OF VISUAL INTELLIGENCE

KEYNOTE SPEAKERS:

DR. AHMED ELGAMMAL
Director of the Art & Artificial Intelligence Lab at Rutgers University

DR. GARY HATFIELD
Director of the Visual Studies Program at the University of Pennsylvania

DR. SUN-JOO SHIN
Professor of Philosophy at Yale University

FACULTY HOST:

DR. ELLIOT PAUL
Assistant Professor of Philosophy at Barnard College, Columbia University

COLUMBIA UNIVERSITY
MUDD - Engineering Building
MUDD 633
500 W 120th St
New York, NY 10027

Institute For Visual Intelligence

WWW.INSVI.ORG
INSVI@INSVI.ORG
718-971-4808

SCHEDULE

10:30 - 45: Registration & Opening Remark

Philosophy of Psychology/Cognitive Science
10:45 - 12: Dr. Gary Hatfield, Louise Daoust
12 - 1: Lunch

Artificial Intelligence in the Domain of Art
1 - 3: Dr. Ahmed Elgammal, Diana Kim, Shahrzad Ziaee, Bingchen Liu
3 - 3:15: Break

Analytic Philosophy/Logic
3:15 - 4:30: Dr. Sun-Joo Shin, Jonathan Pang
4:30 - 5:30: Final Session with Keynote Speakers

Dr. Gary Hatfield

Director of the Visual Studies Program at the University of Pennsylvania

Adam Seybert Professor in Moral and Intellectual Philosophy and Director of the Visual Studies Program at the University of Pennsylvania. He works in the history of modern philosophy, the philosophy of psychology, theories of vision, and the philosophy of science. In 1990, he published The Natural and the Normative: Theories of Spatial Perception from Kant to Helmholtz; at HOPOS 2016, the 25th anniversary of the book was celebrated. In 2009, Perception and Cognition: Essays in the Philosophy of Psychology appeared from the Clarendon Press; a revised version of his book on Descartes' Meditations appeared in 2014. He is a member of the Center for Cognitive Neuroscience, the Penn Perception group, and the History and Sociology of Science Graduate Group. He has directed dissertations in history of philosophy, philosophy of psychology, and philosophy and history of science. He has long been fascinated by visual perception and the mind–body problem.

Dr. Ahmed Elgammal

Director of the Art & AI Lab at Rutgers University

Professor at the Department of Computer Science, Rutgers University. He is the founder and director of the Art and Artificial Intelligence at Rutgers, which focuses on data science in the domain of digital humanities. He is also an Executive Council Faculty at Rutgers University Center for Cognitive Science. Prof Elgammal has published over 140 peer-reviewed papers, book chapters, and books in the fields of computer vision, machine learning, and digital humanities. He is a senior member of the Institute of Electrical and Electronics Engineers (IEEE). He received the National Science Foundation CAREER Award in 2006. Dr Elgammal's recent research on knowledge discovery in digital humanities received wide international media attention, including reports on the Washington Post, New York Times, NBC News, the Daily Telegraph, Science News, and many others

FACULTY HOST

Dr. Elliot Samuel Paul

Assistant Professor of Philosophy at Barnard College, Columbia University.

He works mainly in early modern philosophy and epistemology. He also has interests in philosophy of mind and cognitive science, with a particular focus on philosophical issues surrounding creativity. He holds a BA from the University of Toronto and a PhD from Yale University, both in philosophy. From 2009-2011 Dr. Paul was Assistant Professor/Bersoff Faculty Fellow at NYU. He is co-editor of The Philosophy of Creativity: New Essays, published by Oxford University Press in 2014, and is currently completing a book on Descartes's epistemology.

Dr. Sun-Joo Shin

Professor of Philosophy at Yale University

At Yale Sun-Joo Shin teaches logic, philosophy of logic, history of logic, philosophy of linguistics and, philosophy of language.

In her book "The Iconic Logic of Peirce's Graphs" Shin explores the philosophical roots of the birth of Peirce's Existential Graphs in his theory of representation and logical notation. She demonstrates that Peirce is the first philosopher to lay a solid philosophical foundation for multimodal representation systems.

Varieties of Visual Intelligence

Gary Hatfield
University of Pennsylvania

What is visual intelligence? Is it a single thing? Presumably not. Just as the terms “intelligence” or “human intelligence” denote no one thing but are used to describe a variety of abilities, what might reasonably be meant by the term “visual intelligence” is multi-faceted. The term can be used to describe various skills in visual observation, from noting many details at once, to seeing things that many others do not see, to making clever inferences, Sherlock Holms style, from what is seen. In looking at an artwork, some may notice forms and features that others miss. “Visual intelligence” might also refer to the creativity and special skills that go into making visual artifacts, whether artworks, graphical design, engineering design, or scientific images.

After noting some instances out of this great variety, I will focus on the notion of visual intelligence as posited by many psychologists: the ideas that vision itself, the bare act of seeing, is intelligent or the product of intelligent psychological operations. As manifested among recent theorists in the psychology of perception, including Irvin Rock and Don Hoffman, the idea is that the visual system intelligently infers the visual world from relatively impoverished data supplied at the retina and by the ocular muscles and other associated systems. A central notion is that vision is a kind of problem solving, carried out unconsciously, and involving the application and logical interaction of rules of perception to “construct” the visual world. Focusing on the key notion of construction, I offer an alternative conception of visual intelligence drawn from the theoretical tradition associated with James Gibson: the notion of “smart mechanisms” that respond to visual stimulation in a “smart” way without engaging in inferences or problem solving. I then apply the theoretical conception to re-interpret one of the most celebrated attributions of intelligent operations to the visual system: Descartes’ natural geometry. According to the usual interpretation, stemming from George Berkeley in eighteenth century, Descartes held that the mind calculates the distance of things from the convergence of the optical axes, using the trigonometric relation angle-side-angle. I suggest, by contrast, that Descartes left all the “calculation” in this instance to smart mechanisms in the brain.

Deduction and Visual Representation

Sun-Joo Shin
Yale University

Deduction is a process of extracting certain information from given information. When we correctly spell out an extracting process in a convincing way, it is called a deductive proof. This simple definition of a deductive proof might sound trivial, but at the same time raises the following question: "Why is the process of extraction so difficult to carry out even when we know the conclusion follows from premises?" This mystery of deduction leads us in many directions: It could be a lifetime task for some epistemologists or could set up various debates in the philosophy of mind or cognitive science. Rather than focusing on issues involving the cognitive aspects of our mind, the present talk will relate the mystery of deduction to modes of representation. After locating a source of the mystery of deduction, I will show how visual modes of representation could relieve part of the mystery and will claim this advantage of the visual explains how diagrams aid us in the brainstorming stage of the problem-solving process.

Deduction contrasts with (so-called) ampliative reasoning, in which one gains new information not contained in the premises. Reasoning involved in a scientific theory is a prime example. The search for and justification of hypotheses have given rise to central controversies in the philosophy of science for decades. The reasoning involved in forming a hypothesis from given data is inductive or abductive, and is clearly not deductive, since in induction or abduction we need to go beyond information our data provides us with. The information stated in a hypothesis is more than the information of currently observable data, in order that a suggested hypothesis may predict future events. Thus, we easily see why carrying out an inductive argument is not an easy task, both in how to extend the given information and how to justify the conclusion.

By contrast, deductive reasoning makes us sure we do not conclude something more than the information provided by premises, and therefore, having a system with a finite number of valid rules is a way to assure each deductive step. Indeed, a system is sometimes simple enough to be checked by a machine, and even non-decidable systems do not pose any similar philosophical questions as we have in the case of ampliative inductive reasoning. To put it simply, the information conveyed in a deductive conclusion is "contained" in the premise-information. Nonetheless, how often have we been at a loss in making a correct next step in the middle of a non-trivial mathematical proof? This is the mystery of deduction.

Let me focus on three relevant aspects involving this difficulty: (i) At each stage of a proof, more than one valid step is available to us to get us from one stage to the next. (ii) We need to make a choice among multiple possible valid steps. (iii) There is no algorithm for the choice. That is, each valid step deductively follows from given information, but choosing a correct valid step itself is not a deductive process. How to make a correct choice for the next step is directly related to how our mind operates or how our mind is stimulated. I propose that the mode of representing given information helps stimulate our mind to find the right next step in our deductive reasoning. Citing several case studies (a Euclidean proof, a diagonalization proof, Turing machine representation), I show how visual representation guides us better in carrying out deductive reasoning than does linguistic representation.

The main question I would like to present to the audience is why our mind carries out (deductive) inference better through diagrammatic representation (at least in these cases) than through linguistic representation. Some fundamental differences, if any, among different forms of representation must be responsible for this discrepancy, and these differences can explain our on-going practice of brainstorming in stages, where diagrams have constantly been used. Moreover, I would like to conjecture these differences shed light on our understanding of visual inference, which is not limited to deductive reasoning, but extends to our general cognition involving images.

Dena Shottenkirk
Brooklyn College

To properly understand what it means to look at an artwork involves an examination of the connection between Perception, on the one hand, and Thought, on the other. This is often not believed to be the case, thus it is often under-estimated what it is to give an account of visual aesthetics. But the steps between the initial absorption of a visual entity and the final formulation in a thought – or a verbal account of that thought – are steps that, if correctly provided, explain much in many other fields of philosophy including both philosophy of mind and philosophy of language; it is at the heart of the very basic problems involved in providing a model for cognition.

But there are many ways to proceed in the analysis and the potential pitfalls come embedded in the formulation of the questions: What is the fundamental unit of the percept? Is it a physical object or is it sense data, or yet again, is it a phenomenal experience? What is the content of the perception? My proposal is to deal with the latter: The notion of content. The following is a basic statement of the problems.

Perceptions are thought to have content, but the word "content" is used frequently, in different contexts, and sometimes less than clearly. One most commonly thinks of representational content as found in perception, or propositional content as found in speech and semantic acts. So, before we look at content in perception it makes sense to step sidewise for a moment and look at content in semantic situations.

Speech acts have as their contents propositions. This is the view held by most including both those who advocate a structured propositions view and those who advocate a possible worlds analysis. To take the view from Frege, the sentence has its sense in the abstract but objective object that is the thought. This is the content. So, different sentences in different languages each express the same thought obtained now as an objective reality transmitted from person to person. In other words, the linguistic sentence mirrors the abstract object called a "proposition". The content e.g., the proposition, is timeless, eternal, unchanging and Platonic.

How does this relate to the idea of content in perception? As Frege and Russell thought grammatically structured sentences had their meanings in propositions, and were thus bearers of truth through their relationship to the similarly structured – though ontologically different – propositions, so many have argued that our perception has its contents in our grammatical representation of that experience, which in turn of course has its structure similarly concatenated in the abstract object propositions. And, once again, it in turn has its constitutive foundation in concepts. The tracing of this seems harmless. Beliefs have their foundation, largely, in perceptual experiences, and perceptual experiences are related in language, which has as it twin propositions. To say that perceptual experience is structured by concepts is to say that an understanding of perceptual experience is ultimately traceable back to concepts. Thus, 'conceptualism', as a term used within an analysis of perceptions, argues that for any perceptual experience, that experience has as its content a Fregean proposition.

There is a slightly different way of analyzing the contents of perceptual experience, and also giving us conceptualism. Instead of analyzing the process in terms of translations from one kind of experience (e.g. perceptual) to another yet equivocal kind of experience (e.g. linguistically and then ultimately propositionally) we instead start experientially from the basic, concrete moment of experience. In this very traditional way of approach the first question generally asked is, 'what are the immediate objects of awareness?'. That places us – the individual observer - in relation to an object. Now the question focuses on the nature of that object. Or to phrase it in the most-oft way: what is given in experience? Searle has of course ravaged the historically most frequent answer: sense-data – those private, non-physical entities that are experienced as sensory qualities have been thought by many, including Hume and Locke and Russell, to constitute our experience. Founded in our sensing, these entities though are ultimately founded in the empirical non-mental world – they are our bodies' translations of non-mental phenomena, e.g., the way we absorb or experience the physical world. And that physical world is antecedent to – and thus in that way a priori to – our experience of it.

This gets us to concepts is a slightly different way. Concepts become those things we construct out of the indirect data. They may, especially on something like a Kantian view, still be located in an a priori world, and thus very similarly to Fregean propositions, be content that is antecedent to the experience of the viewer. Regardless of that, sense-data theorists generally regard the contents of perception to be structured ultimately, through thought, by concepts.

A world antecedent to the experiencing individual is a common theme, and generally forms the basis for truth-values. To reiterate, the Fregean proposition has a metaphysics based on eternal, platonic, structured and a priori entities – which is waiting for the experience of the individual; while in both the indirect and direct perceptual theories there is an already existing, structured, and inert external material world, which is also waiting for the experiencing individual. Both of these positions – the idealist one underwriting the propositional content of linguistic utterances and the materialist one underwriting the perceptual contents of experience – posit a static realm that is received by a passive observer.

Stalnaker in his "What Might Nonconceptual Content Be?" refers to Gareth Evans in the following, "Evans regards it as important to identify information-bearing states of perceptual systems with states of seeming since he is anxious to avoid the traditional epistemologist's picture according to which the subject receives, through the perceptual systems, sensory data that is 'intrinsically without objective content,' but which forms the basis for inferences about the world that causes them." (and here Stalnaker is quoting Evans) "'The only events that can conceivably be regarded as data for a conscious, reasoning subject are seeming's – events, that is, already imbued with (apparent) objective significant'".

I would argue that the contents of perceptual experience do not have to have the same kind of content or structure as do beliefs. And it is nonconceptual theories of perception that can get us there.

Philosophical Understanding of Visual Intelligence

Joanthan Pang
Institute for Visual Intelligence

Much of the time the objectivity of aesthetic judgment is in question. What causes one to say that a painting is "good," and what makes others dispute this judgment? Everyone has the capacity to make aesthetic judgments, yet some can claim to have better judgment than others, highlighting the putative exclusivity of good aesthetic judgment. Investigation into visual intelligence may provide clarity into our notions of objectivity as applied to aesthetic judgment. First, we must understand what intelligence is. "**Intelligence is the efficient use of cognitive, rational, mental resources, something which involves thinking, deliberation, reasoning, pondering, remembering, weighing alternative courses of actions and, therefore, the use of mental representations**" (Lanz, p.21-22). Visual intelligence is thus construed as an intelligence applied to aesthetic judgment of visual arts.

This kind of intelligence is more comprehensive than Gardner's definition of "spatial intelligence." (Spatial intelligence according to Gardner is the computational capacity to solve spatial problems, examples including navigation, visualization of objects, and notice of fine details.) For example, two people can have a similarly high level of spatial intelligence, but it is possible that the outcome of their aesthetic judgments on the same artwork could be very different. A high degree of spatial reasoning is fundamental to making good aesthetic judgments, but discrepancies between levels of visual intelligence will account for the difference in aesthetic judgment. This distinction tacitly reflects that aesthetic judgment is based upon spatial intelligence on a fundamental level, and that any objectivity found in the measurement of spatial intelligence can contribute to an argument for certain objective standards on what is called good aesthetic judgment.

Some key components of our aesthetic judgment, however, are not covered by spatial intelligence. There are emotion and feelings, and other senses that are not directly involved with spatial intelligence, yet important data that contributes to our aesthetic judgment. To reveal these components that visual intelligence encompasses and spatial intelligence lacks, we investigate specific cases.

Case 1: An artist is proficient in rendering perspective, and always draws the most minute details, thereby invoking an almost photo-realistic effect. However the artist has poor taste in color and the work is not visually pleasing. In this case the artist is spatially intelligent in the area of rendering, yet not spatially intelligent with regards to color scheme.

Case 2: Another artist is incapable of rendering things realistically, and has had no formal training to develop his spatial intelligence (or maybe just does not have a fully developed sense of space.) To us, the work from this artist is visually pleasing. What tells us this? It may seem his spatial intelligence is simply inadequate in certain areas such as rendering realistically, while he is spatially intelligent with regard to other areas, giving us enough incentive to enjoy his work. But can it be conceived that he is not spatially intelligent at all, while his work is still pleasing? The main cause of our attraction to the work must be attributed to something, which I claim is visual intelligence.

Case 3: Another artist is very talented at rendering photo-realistic work. The work is not visually appealing. One can say that the artist who can create photo-realistic work is spatially intelligent, yet something causes aversion to his artwork. We have two options- we can conclude that an aspect of the artwork reflects spatial intelligence in its well-rendered aspects, while the artwork does not appeal to us due to a flaw in its appearance, a flaw that we recognize through spatial intelligence. It can also be the case that it lacks in its appeal to us for some other reason, a reason which we would attribute to visual intelligence.

Case 4: We might see a photo-realistic painting and judge that it looks "cheap" although the painting clearly displays the artist's spatial intelligence. Again, we are averse to the work since we believe that it looks cheap- it is possible that we feel that the work shows lack of novelty. This feeling might also be attributed to visual intelligence.
[[Explication of visual intelligence provides clarity as to which components of our aesthetic judgment can be considered objective.]] **Philosophical investigation into the boundaries between visual and spatial intelligence <u>was conducive towards defining visual intelligence</u>.**

Multi-branch Convolutional Networks for Artistic Style Classification

Bingchen Liu, Mohamed Elhoseiny, Ahmed Elgammal
Art & Artificial Intelligence Laboratory
Dept of Computer Science
Rutgers University

In this work, we investigate the effect of multipath convolutional encoding of images for a challenging art-style classification problem. We comprehensively studied the state-of-the-art Convolutional Neural Networks models on this task for classification of paintings according to their artistic-style. Motivated by the nature of art-style classification, we introduce a convolutional-branched architecture that we can achieve a better performance compared to existing deep architectures like (AlexNet[1], VGGNet[2] and ResNet[3]), without increasing the capacity and sacrificing the training time.

Structure:

We integrating multiple filter sizes (e.g. 3, 7 and 16) on multiple branches separately in our model directly after the data input layer. The intuition behind using multiple branches is, different filter sizes allow our network to look at features from different scales, ranging from 3x3 windows (capturing finer details) to 16x16 windows (considering somewhat larger areas, while capturing composition). In the proposed architecture, the input image is processed by three convolutional layers simultaneously and in parallel. These 3 layers lead to their own convolutional branches with different depth. The output of these three branches is then concatenated and fed into 3 fully connected layers with the last FC layer is the output of class prediction.

Performance:

We evaluated our model on the publicly available art dataset from wikiart.org, and pick more than 80,000 images that are considered good representations of their artistic styles. Our current experiments are done with the consideration that the models will learn with no prior and external knowledge other than art, so they are not fine-tuned from any other dataset. Among all the architectures that we tested, our model achieved the best Top-1 Accuracy rate of 50.3% for classifying among 20 style classes, which is 4.2% better than AlexNet and 1.5% better than VGG-net 16 layers, since our dataset is relatively small compared to the ImageNet, so ResNet_50 only gets 43% of top-1 accuracy because of its large capacity that needs more data for a fully training. These results show that introducing convolutional branches regularized the learning process and enabled the CNNs to look at different receptive fields, which is important to improve style classification.

The model size is slightly smaller than AlexNet even though it has three branches and more convolutional layers. While the AlexNet model with kernel size 7 on first conv layer has size 282.2MB and a VGG-16 layers model has weights size 553.4 MB, our multi-branch-multi-filter model has 248.8 MB, which is 11.8% less than AlexNet. Such size difference also reflects on training time usage, where Multi-Branch (MB-net) model has the same training time as AlexNet and about 4 times faster than VGG-16 net.

Analysis and Discussion:

By visualizing the weights from all these models (AlexNet, VGG16, Multi-Branch), we can find some cues that relates to their performance. The proposed Multi-Branch (MB-net), the weights look more regular and shallower than VGGNet (figure 1), where VGG-net is still looks colorful and randomized (figure 2). Which shows that MB-Net with a smaller capacity can preserve the original image info better and find more useful features from the beginning of its processing procedure. Same observation can be found on the deeper layers, as we look at the last convolutional layers of these models, MB-net shows more neurons are activated (figure 3), which means it could provide more information to the decision layers (FC layer)

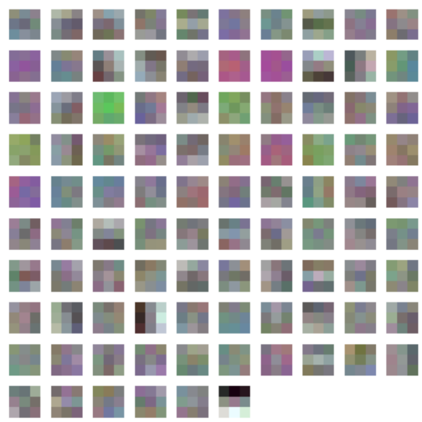

Figure 1 - Filters from the MB-net-branch 1-conv layer 1.
Each filter has formed relatively shallow and regular patterns

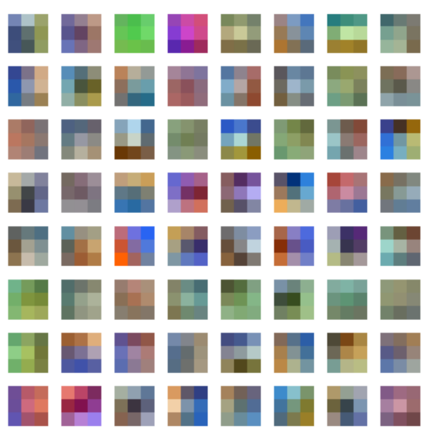

Figure 2 - Filters from VGG-net-conv 1.
Filters are still colorful and randomized, no obvious patterns can be find in these filters.

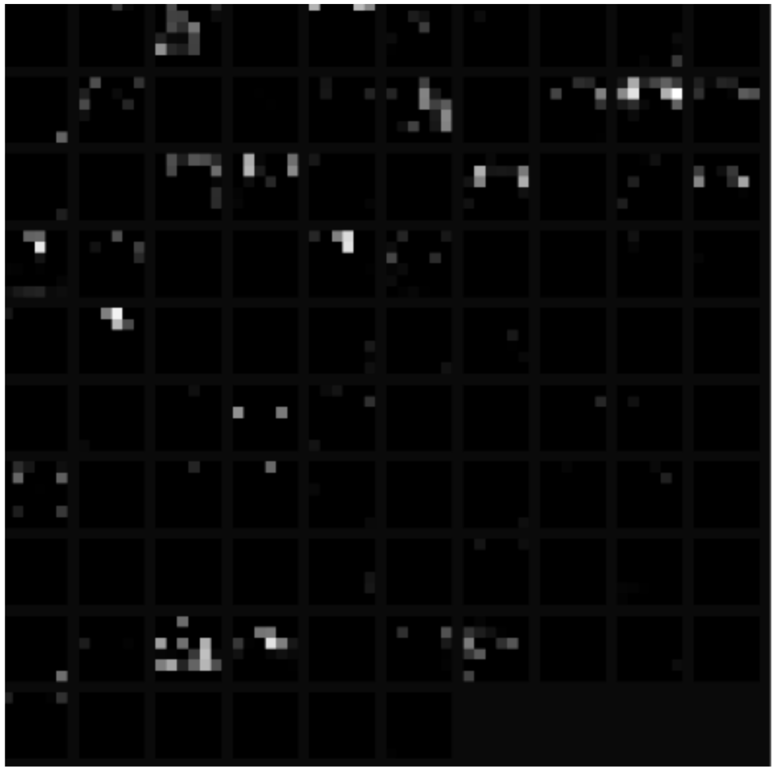

Figure 3 - Activations from the MB-net-branch-3-conv layer 3.

More neuron units are activated after passing an image into the network.

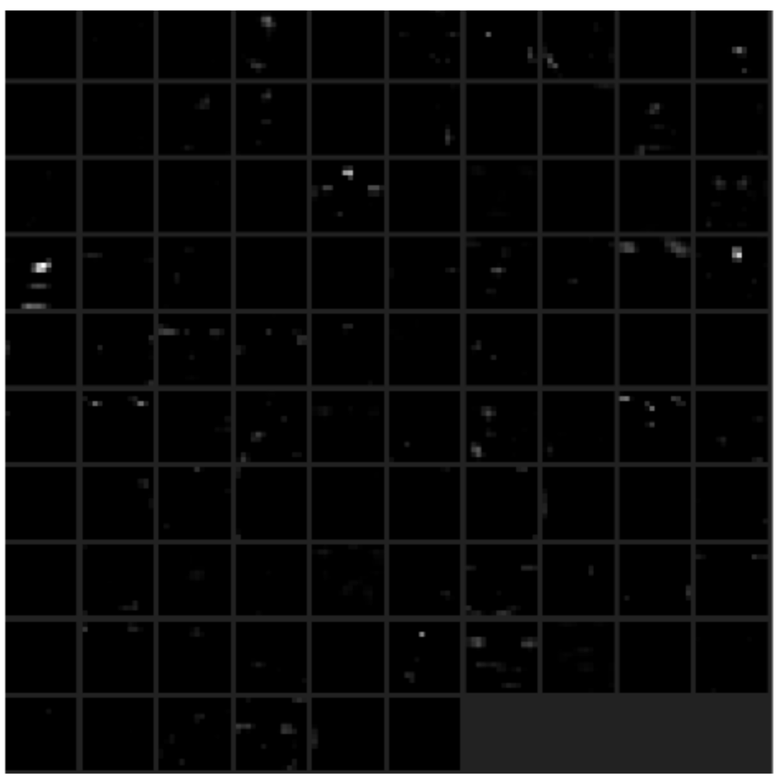

Figure 4 - Activations from VGG-net conv layer 5.

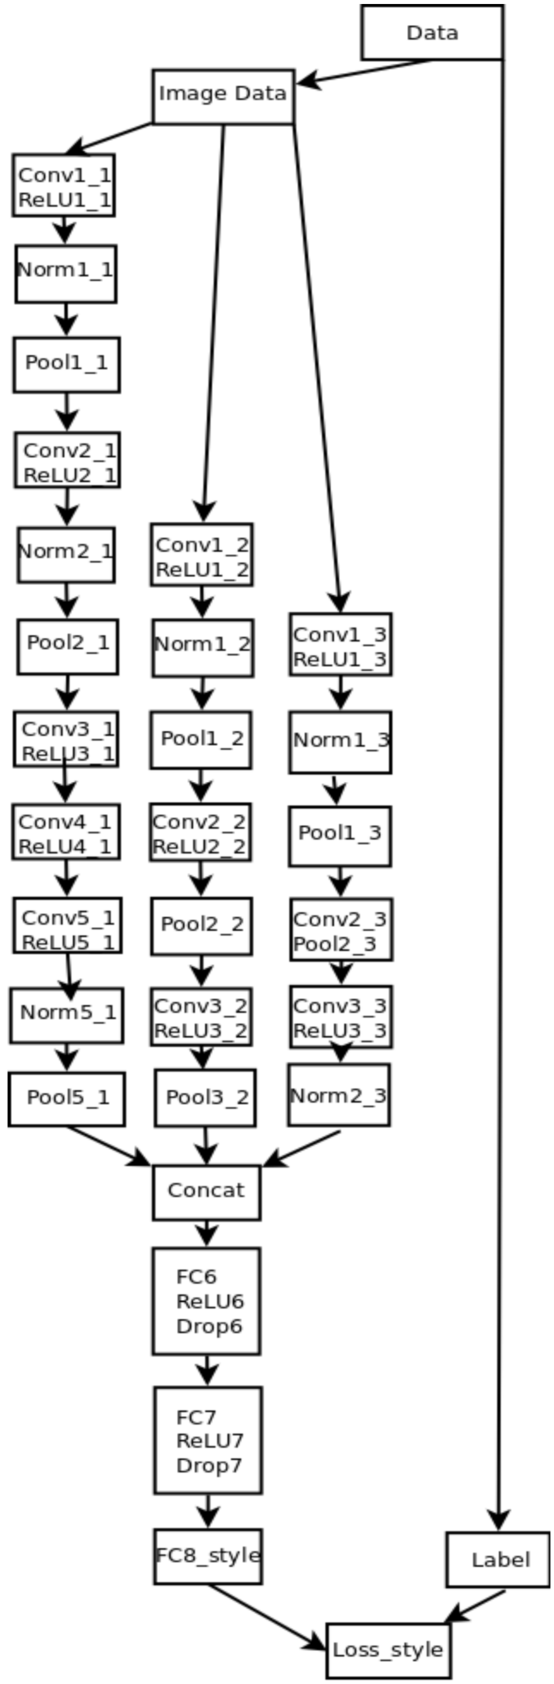

Figure 5 - The structure of Multi Branch Network.

The key-note is that each branch has a different initial filter size on its conv-1 layer (3x3, 7x7, 16x16), thus they can look at the image from different scales and process them differently simultaneously.

References:

1. Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. In Advances in neural information processing systems, 1097–1105.
2. Simonyan, K., and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. CoRR abs/1409.1556.
3. He, K.; Zhang, X.; Ren, S.; and Sun, J. 2015. Deep residual learning for image recognition. arXiv preprint arXiv:1512.03385.

Computational Semantic Feature Analysis for Art Styles

Ahmed Elgammal[1], Marian Mazzone[2], and Diana S. Kim[1]

[1]elgammal,dsk101@rutgers.edu
[2]mazzonem@cofc.edu
[1]The Art and Artificial Intelligence Laboratory
[1]Department of Computer Science, Rutgers University
[2]Department of Art History,College of Charleston

As defined by an American art historian, Meyer Schapiro(1904-1996), the style is a conceptual system synthesizing atomic forms, elements, qualities, and expressions which show consistency in the art of an individual or a groups. Finding the components which determine styles is an essential subject which historian has been invested because combination of the factors becomes fundamental grounds to measure uniqueness and creativity of the pieces as an art by finding consistent or distinctive features between comparison artworks.

However, there is no exact formulation existed to represent the systematic relation between constant features and styles, and in practice, styles are not usually defined in a strictly logical way. Even though, there were some historical works to build objective elements determining styles, such as Swiss art historian Henrich Wolffin (1864-1945)'s 5 principals (linearly to painterly, plane to recession, closed form to open form, multiplicity to unity, and absolute clarity to relative clarity), the 5 elements are not fully universal, as they were initially devised to discriminate only for Baroque and Renaissance styles, also the degree of each element cannot be evaluated accurately by human perceptions. Thus, in constructive purpose, if we can find conceptual machinery features which could classify the styles with acceptable accuracy rate, then this research will provide reliable baseline to disclose universal forms consisting of styles.

Hence, we are proposing a framework to discover semantic features by using number of sophisticated techniques in machine learning and data science. We will show how our method can find conceptual features in style recognition of artificial intelligence by investigating visual machinery features from fully connected layer of CNN (Convolutional Neural Network), and by restricting a scope of artwork to fine art paintings. We will also verify how these findings are in boundary of human concepts, based on a current state of art mechanism and process of formal assessment with art historians.

Our approach is motivated by some correspondence between computer vision of deep convolutional neural network and biological human vision network, so we are expecting that our machinery semantic features are going to provide fundamentals to discover essential features for human perception in style recognition. Along with our approach, to extract conceptually interpretable features, we will perform ICA(Independent Component Analysis) on the intermediate results. Also, we will place first norm sparsity constraint on one of the layers of CNN during neural network training practice. In general,the sparsity constraint degrades a bit of classification performance but it is a necessary step to conduct more reliable ICA and to find concise semantic features.These semantic features in this research will be clarified by mapping them to human words through a componential analysis in a pre-trained word embedding space, called the global vectors for word representation (GloVe).

Patterns of Artistic Handedness Applied to Formal Attributes And Authorship

Shahrzad Ziaee[1], Ahmed Elgammal[1] and Marian Mazzone[2]
[1]Art & Artificial Intelligence lab, Computer Science Department, Rutgers University, NJ, USA
[2] Art & Architectural History Department, the College of Charleston, SC, USA

Motivation: It is a popular belief that proportionately more artists and musicians are left-handed than the general population. Left-handers are thought to be more creative, intelligent, and to have better access to the right hemisphere of the brain where artistic inclinations are housed. Scientifically, however, the proof for this is very slight. Despite hundreds of studies on left-handedness as it relates to intelligence, sports skills, creativity, health and more, empirical evidence is conflicted or lacking.

Goal: Instead of resolving this popular debate, our task is to characterize the pattern between handedness and art creation on the level of formal attributes when handedness is known. We propose to discover if there is any pattern inherent within the use of one hand or the other that we can detect, first by determining handedness characteristics in the formal elements via computational analysis. There is a precedent within connoisseurship for being able to ascertain the handedness of an artist based on hatching patterns and directionality of strokes.
Computational connoisseurship adds to this tradition by analyzing elements such as stroke directionality, compositional patterning, line weight, etc. across large numbers of works of art for statistical legitimacy. What one careful connoisseur could do can be multiplied a thousand-fold through computational speed and accuracy. This information can be part of a useful toolkit to determine the handedness of a work's creator, and by extension the legitimacy of claimed authorship of a work of art. The research questions we mostly intended to answer in the preliminary study were the followings:

- Is there a visual signal that can automatically detected that can correlated with left or right-handedness of artists?
- How reliable is the existing public information about handedness of an artist?

Methodology: Relying on anecdotal evidence of the handedness of artists and how they manipulated materials is not sound, because accurate information about historical figures is mostly absent. Although Michelangelo, for instance, has often been described as left-handed, the historical proof for that is non-existent. Aimed at performing accurate computational analysis, we evaluated three categories of visual encodings consisting of global features, local-to-global directional features and local texture information. *The preliminary results showed that the first and the second types of signals can act as predictors for handedness.* We also propose to survey the handedness and formal practices of living artists to create a database of accurate information about how handedness affects art creation. We then can match this data back against the results of computational formal analysis to make more accurate connections between characteristics such as handedness and any specific formal attributes.

Potential in Authentication: Linking formal attributes with artist handedness and the art making process promises a variety of applications, including the ability to affirm who painted or drew a work of art by comparing known examples to unknown ones, to detect one hand versus others in studio or workshop productions, or to discern whether authorship holds up across a number of works purported to be by a single artist. Our analysis may provide a key to unlocking authorship on the formal level that would be difficult to consistently evade across the number of features that could be used for detection.