Information Theory and Art Abstraction Estimates

(Words and Numbers)

Gordon E. Lyon[1] and Merritt R. Lyon[2]

Keywords: art abstraction; information theory; measurement

A fine art discussion will occasionally involve estimating some work's degree of abstraction. Fortunately, a reliable "continuum of abstraction" convention—an ordered progression of ranks—has evolved to support art criticism, pedagogy and research. Realistic art objects sit at one end with highly abstract, non-representational works at the other. All works are scored by viewers, who choose from finite, ordered ranks such as '1—realistic,' '2—mixed,' and '3—abstract.' Subjective but extremely useful, the abstraction framework adapts well. For example, some investigators have used a hierarchical table whose five ranks range from I (most representational) to V (extremely abstract). Others have chosen nine- (1—9) or eleven-point (0-10) ranks. Still, limitations include the arbitrary nature of the format and ranks that are order data (i.e., do not support arithmetic). These points hinder inquiries such as, "Abstraction ranks signify what?" or "How do framework ranks relate?" We conjecture that capturing representational ambiguity helps address these questions. By representational ambiguity we mean collective, or viewership, uncertainty in describing a work via captions, Since highly abstract art is characterized as non-representational, might not independent viewers, having no realistic cues, assign such work the greatest assortment of captions? This caption variety might be a proxy for abstraction level.

Approach: Our experiment uses images of artwork, circa 1860 to the present, shown to viewers via the internet. For simplicity, the photographs are restricted to sculpture. Museums issue quality photographs of their sculptures as isolated objects, which assures that the internet viewers see and focus on the same thing. For each of 25 images, 60 remote viewers provide (i) a conventional, subjective abstraction score and (ii) a caption of 36 or fewer characters succinctly expressing whatever impression a piece evokes when the viewer first sees its image. Each image thus generates two sets—one of 60 scores and another of 60 brief descriptor captions. Captions for an image are reduced to their English simple subject—usually one word—and sorted into classification categories. Each category is then counted. The category counts constitute an image's partition, an ensemble of occurrence frequencies. Information-theoretic metrics use an image's partition to measure its representational uncertainties. The median of the 60 subjective scores establishes a sculpture's rank on the continuum of abstraction.

Results: Realistic sculptures do acquire many similar captions that get classified into a few categories. In contrast, the partition for a highly abstract work can have a majority of its classification categories holding but one or two captions. Three statistical metrics (including entropy, H) demonstrate that as sculpture abstraction rank increases, so too does representational uncertainty. And, as might be expected with human respondents, there is plenty of variation in subjective scores for an image. Nonetheless, viewer score/caption pairs add significant quantitative detail to the continuum of abstraction. They make possible statements such as, "With metric H, median uncertainty for rank 3 is 2.1 times that of rank 1" and "Sixty viewer captions generated 48 distinct categories for Moholy-Nagy's 1921 abstract work Nickel Construction."

[1] Montgomery College, Silver Spring, Maryland and independent artist, 390 Russell Ave., Gaithersburg, MD 20877 <gelyon@verizon.net>

[2] Department of Statistics, Columbian College of Arts and Sciences, George Washington University, Rome Hall 801 22nd St. NW, 7th Floor, Washington, DC 20052