Subsurface Image Reconstruction by Near-Infrared Reflectance Analysis

ROBERT A. LODDER and GARY M. HIEFTJE*

Department of Chemistry, Indiana University, Bloomington, Indiana 47405-4001

A method is described for reconstructing decorative images concealed by layers of overcoatings. The method employs a nonparametric multivariate algorithm and a cellular-automaton construct to convert the near-infrared spectrum of a surface into digital images of the surface layers. The technique is potentially applicable to the restoration of historic buildings where aging, previous restoration attempts, and even political considerations have resulted in the alteration of the original historic surfaces.

Index headings: NIRA, image reconstruction; Automata; Paint.

INTRODUCTION

An important aspect of the restoration of historic buildings is the analysis of the paints and coatings used in the original construction. Unfortunately, it is unlikely that any painted surface will remain unchanged over time: smoke, dirt, aging of overcoats, and previous attempts to modify or restore the surface all complicate the problem of determining what was originally there. The investigation of historical surfaces must, therefore, attempt to determine the number of coating layers and the approximate date of each layer, the original colors (usually recorded according to the Munsell Color System, a standard system of describing and identifying colors), the type of coatings (oil- or water-base paints, stains, varnishes, wallpapers, etc.), and the nature of any decorative painting (such as stenciling or marbling).[1] Although the investigation of historical surfaces must always be considered in the context of the site as a whole (and generally only after the historical document research is finished), this report deals with only one aspect of surface-coating analysis: the detection and reconstruction of decorative images beneath one or more layers of obstructing paint.

The microscopic examination of paints and other coatings by restoration experts makes up the bulk of direct historical-surfaces analysis. Actual pigment identifications and chemical analyses to verify specific elements and compounds are seldom performed. Near-infrared reflectance analysis (NIRA), however, has the potential both to detect the presence of hidden decorative patterns and to provide chemical information about the surface layers.

The near-infrared scanning technique described here is important because it can make one aware of the need to restore a particular surface, and because the technique can perform the analysis relatively quickly, nondestructively, and inexpensively. The choice between restoring the original coatings or simply duplicating them on a fresh surface can then be intelligently made.

THEORY

NIRA makes use of pattern-recognition techniques (like multiple linear regression); consequently, the NIRA algorithm must be "taught" to interpret spectra using a "training set" of known samples that adequately represents the range of variability expected in the samples to be analyzed. This training-set requirement might seem difficult to fulfill when the samples are patches of paint on a wall: indeed, in many circumstances we may well be expecting, or at least hoping, that some unknown image will appear. How can this image (or perhaps more appropriately, its constituents) be detected and quantified?

Received 5 September 1987.

* Author to whom correspondence should be sent.

0003-7028/88/4202-0309$2.00/0

If the training set is taken to be a known, homogeneous section of paint (homogeneity can be verified simply by chipping—a process that would have been undertaken at some time anyway), the problem of detecting an unknown image beneath the paint becomes a form of the "false-sample problem." The false-sample problem arises whenever a pattern-recognition algorithm is presented with a sample unlike any the algorithm has ever seen before. A method of solving the false-sample problem, the BEAST (Bootstrap Error-Adjusted Single-sample Technique), has been described.[2] Modifying the BEAST procedure slightly to increase its stability with small training-set sizes (thereby minimizing the amount of chipping or other analysis that must be performed in creating a training set) and incorporating the BEAST into a simple two-dimensional cellular-automaton rule allow spectra stored in a 2-D cellular automaton to evolve into a meaningful image.

The Modified BEAST. The BEAST treats each wavelength in a spectrum as a dimension in hyperspace. A spectrum recorded with the use of *d* wavelengths appears as a point in a *d*-dimensional hyperspace, translated from the origin along each axis by an amount that corresponds to the signal observed at each wavelength. Samples with similar spectra produce clusters of points that project into similar regions of hyperspace.

The BEAST takes a set of training spectra and develops a discrete estimate of the real-sample distribution (cluster) from the training set. A point estimate of the center of the real-sample distribution is also found. When a new sample is analyzed (in restoration, for instance, a new section of a wall), the new spectrum is projected as a point into the training-set hyperspace. A vector is then formed in hyperspace between the new spectral point and the computed center of the real-sample distribution. At this point a hypercylinder can be formed with this vector as the cylindrical axis. The population density of BEAST-estimated real-sample spectral points within this hypercylinder is used to construct a confidence interval along the axis of the hypercylinder. The use of an asymmetric nonparametric central 68% confidence interval produces BEAST distances that are analogous to standard deviations.

The radius of the hypercylinder can be varied to produce a balance between directional selectivity and distance stability.[3] When the training-set size becomes relatively small (e.g., 10 or 20 samples) the instability of the BEAST distance (the bias of the distance, in particular) can become significant. By increasing the radius of the hypercylinder to include all of the estimated real-sample spectral points within the cylinder, and by using the projected distance of the estimated real-sample spectral points on the hyperline to create the population-density distribution above, one can minimize the instability of the BEAST. The modified BEAST allows all of the spectra to contribute a certain amount of information to the distance estimate, not just the spectra in a certain direction. However, the contribution of a spectrum to the distance estimate is weighted by the cosine of the direction of the spectrum's displacement from the cylindrical axis vector and the center of the real-sample distribution. This weighting makes the effect of the loss of directional selectivity on the BEAST distance minimal for well-behaved training sets (training sets of relatively homogeneous composition).

Cellular Automata. A two-dimensional cellular automaton is simply a 2-D array of sites. These sites are often squares (each in contact with 8 adjoining sites), and each site can hold some sort of information, much like a mailbox. The information at the sites is updated simultaneously and at discrete time intervals according to a describable rule that applies to all sites. This rule generally specifies the new value of a site as a function of the site's previous value and of the previous values of some neighborhood of cells around the site. In addition, the rule requires that each site assume a finite set of values. Cellular automata can be thought of as idealizations of differential models in which time and space are quantized.[4]

Cellular automata have been applied to many problems for many reasons, but there is one major advantage to considering their use in image-reconstruction applications: cellular automata are analogous to parallel digital computers. The typical computer available today processes its data serially, organizing bits into bytes or words and processing these entities one at a time. In cellular automata, the initial array values can be thought of as data, the rule as a program, and the final configuration of cell values as the output result. Making each cell a simple processor turns the temporal evolution of the cellular data configuration into data processing, enabling calculations to be performed with previously unimagined speed.[5] The reconstruction of a subsurface image from pixels (picture elements) that are themselves overlapping spectra recorded at a large number of wavelengths is a computationally intensive task. The mapping of these spectra into a cellular automaton makes the reconstruction problem conceptually simple and opens up the possibility of employing high-performance parallel-processing units[6] to achieve real-time results.

The modified BEAST and a cellular automaton are melded together by incorporating the new BEAST into a cellular automaton rule. A coated surface is raster-scanned by a near-IR spectrophotometer, collecting pixel spectra and loading each of these spectra into a corresponding cellular automaton. The cellular automaton is then initialized by a rule that is trained on a small section of a known coated surface. The initialization converts the *d*-dimensional spectra in the cells into SDs, scalar quantities that are readily enhanced by algorithms like FACADES (Filtering Algorithms using Cellular Automata for Digital Enhancement of Signals).[7] The temporal evolution of a cellular automaton under a FACADES rule eventually produces a steady-state configuration (a class 2 cellular automaton) that is unchanged by further iteration.[4] This final configuration is readily represented with the use of contour plots or color graphics.

EXPERIMENTAL

The spectral data were collected at 18 wavelengths by a Technicon InfraAlyzer 400 filter spectrophotometer (Technicon Industrial Systems, Tarrytown, NY). The spectrophotometer was directly connected to a VAX 11/780 computer (Digital Equipment Corp., Maynard, MA). Analysis of the spectra and the data display were per-

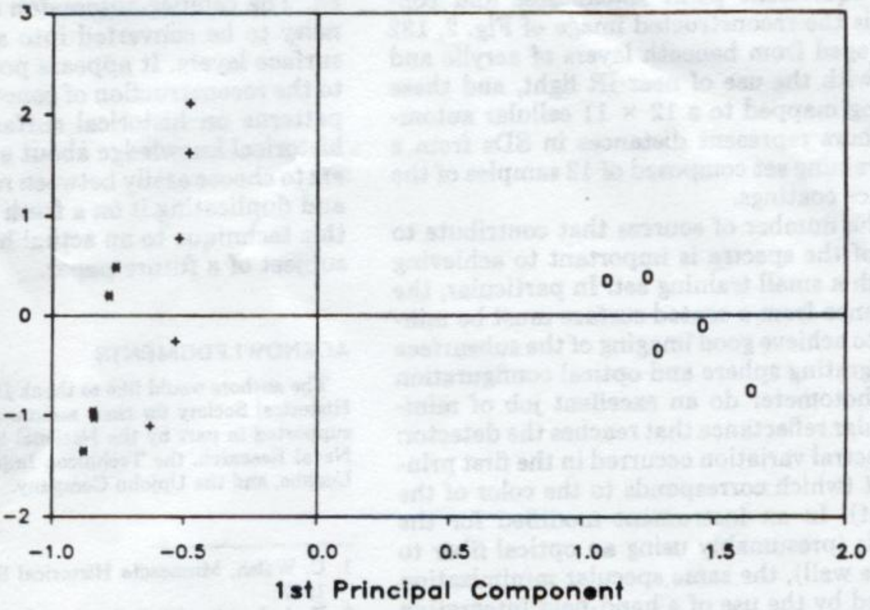

FIG. 1. Near-infrared spectra of red (+), green (O), and white (*) acrylic paint samples overcoated with white acrylic paint.

formed with the use of programs written in Speakeasy IV Delta Plus (Speakeasy Computing Corp., Chicago, IL).

We tested the feasibility of using near-infrared probe radiation for the reconstruction procedure by painting an image on illustration board. The image (one coat of the Indiana University logo in red, painted against a single-coat green background) was covered with two coats of white paint. The paints used were red (XF-7), green (XF-5), and white (XF-2) acrylics from the Tamiya Plastic Model Co., Shizuoka-City, Japan, and Krylon aerosol white enamel (#1502), Borden, Inc., Columbus, OH. The red and green layers were brushed onto the board, as was the first layer of white acrylic overcoat. The second layer of white enamel overcoat was sprayed on. The enamel layer smoothed the surface and ensured complete coverage of the underlying image. The illustration board was then numbered off into pixels on the back with a pencil and cut into pieces for analysis. Each pixel was supported on an aluminum block during scanning in the spectrophotometer.

RESULTS AND DISCUSSION

Test spectra of the three acrylic paints were collected before an image was prepared for analysis. The three paints were each separately brushed on five illustration-board squares. After this coat dried, all 15 squares were covered with two coats of the white acrylic paint. The 15 spectra of these squares were transformed to their principal components, and the results appear in Fig. 1. (The principal component axes are formed from linear combinations of wavelength axes and can be thought of as a sort of wavelength space, simplified by the deletion of wavelengths containing no useful spectral information.) The data in Fig. 1 suggest that image reconstruction using near-IR radiation is quite feasible.

Figure 2 depicts the image to be reconstructed from near-infrared spectra. The Indiana University (IU) logo was painted in red against a green background to provide

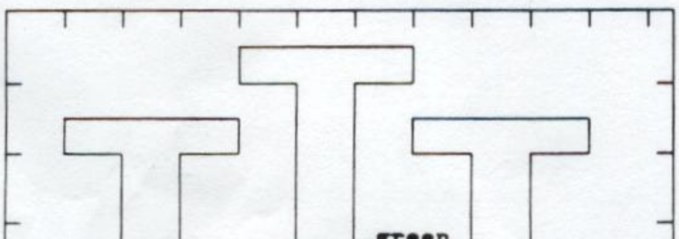

FIG. 2. The Indiana University (IU) logo was painted in red against a green background to equalize paint thicknesses and contrast. This image was then covered with a layer of white acrylic paint and a layer of white enamel paint.

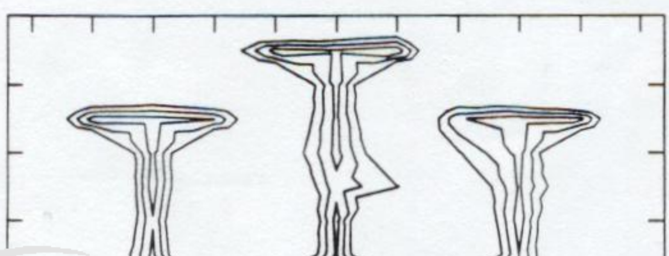

FIG. 3. The reconstructed image of the logo in Fig. 2. The contours (appearing at 5, 10, and 15 SDs) can be thought of as giving probabilities that the subsurface layers are different from the top layer.

approximately equivalent paint thicknesses and contrast. Figure 3 is the reconstructed image of Fig. 2, 132 pixels being imaged from beneath layers of acrylic and enamel paints with the use of near-IR light, and these pixels then being mapped to a 12 × 11 cellular automaton. The contours represent distances in SDs from a homogeneous training set composed of 12 samples of the two white surface coatings.

Controlling the number of sources that contribute to the variability of the spectra is important to achieving good results with a small training set. In particular, the specular reflectance from a coated surface must be minimized in order to achieve good imaging of the subsurface layers. The integrating sphere and optical configuration of our spectrophotometer do an excellent job of minimizing the specular reflectance that reaches the detector: 99.8% of the spectral variation occurred in the first principal component (which corresponds to the color of the subsurface paint). In an instrument modified for the scanning of walls (presumably using an optical fiber to carry light to the wall), the same specular minimization could be achieved by the use of a hand-held integrating sphere/detector module at the sample end of the optical fiber. Ordinary wire cable could easily be used to carry electrical information back to the instrument and a portable computer.

CONCLUSIONS

The use of near-infrared reflectance to locate hidden inhomogeneities in coated surfaces has been demonstrated. The cellular-automaton construct allows inhomogeneity to be converted into a meaningful image of subsurface layers. It appears possible to apply this method to the reconstruction of concealed decorative images and patterns on historical surfaces, thereby adding to the historical knowledge about a site and permitting restorers to choose easily between restoring the original coating and duplicating it on a fresh surface. The application of this technique to an actual historical surface will be the subject of a future paper.

ACKNOWLEDGMENTS

The authors would like to thank Dr. Cass Welsh and the Minnesota Historical Society for their assistance in this project. This work was supported in part by the National Science Foundation, the Office of Naval Research, the Technicon Industrial Systems Division of Bran Luebbe, and the Upjohn Company.

1. C. Welsh, Minnesota Historical Society, personal communication, 1987.
2. R. A. Lodder, M. Selby, and G. M. Hieftje, Anal. Chem. **59**, 1921 (1987).
3. R. A. Lodder and G. M. Hieftje, paper submitted for publication in Anal. Chem.
4. S. Wolfram, Sci. Amer. **251**, No. 3, 157 (1984).
5. W. D. Hillis, Physica **10D**, 213 (1984).
6. T. Toffoli, Physica **10D**, 195 (1984).
7. R. A. Lodder, M. Selby, and G. M. Hieftje, Presented at the Pittsburgh Conference, Atlantic City, New Jersey (1987), Paper No. 026.