

Computational Image Marking on Metals via Laser Induced Heating

SEBASTIAN CUCERCA, Max Planck Institute for Informatics, Germany

PIOTR DIDYK, Università della Svizzera italiana, Switzerland

HANS-PETER SEIDEL, Max Planck Institute for Informatics, Germany

VAHID BABAEI, Max Planck Institute for Informatics, Germany



Fig. 1. Laser marked images on stainless steel using our proposed method (the plates are 13×13 cm). Images © Piotr Didyk, Alexandr Ivanov, Sebastian Cucerca.

Laser irradiation induces colors on some industrially important materials, such as stainless steel and titanium. It is however challenging to find marking configurations that create colorful, high-resolution images. The brute-force solution to the *gamut exploration* problem does not scale with the high-dimensional design space of laser marking. Moreover, there exists no color reproduction workflow capable of reproducing color images with laser marking. Here, we propose a measurement-based, data-driven performance space exploration of the color laser marking process. We formulate this exploration as a search for the Pareto optimal solutions to a multi-objective optimization and solve it using an evolutionary algorithm. The explored set of diverse colors is then utilized to mark high-quality, full-color images.

CCS Concepts: • **Applied computing** → **Computer-aided manufacturing**; • **Computing methodologies** → **Computer graphics**;

Additional Key Words and Phrases: laser marking, color reproduction, multi-objective optimization, genetic algorithm, computational fabrication

©Publication rights licensed to Association for Computing Machinery.
This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in ACM Transactions on Graphics.

ACM Reference format:

Sebastian Cucerca, Piotr Didyk, Hans-Peter Seidel, and Vahid Babaei. 2020. Computational Image Marking on Metals via Laser Induced Heating. *ACM Trans. Graph.* 39, 4, Article 1 (July 2020), 12 pages.
<https://doi.org/10.1145/3386569.3392423>

1 INTRODUCTION

Creating visible patterns on surfaces using laser irradiation is a rapidly growing technology with many applications in object identification, customization, and authentication [Liu et al. 2019]. *Laser marking* is an environmentally friendly, low maintenance process with no consumables, dyes, or pigments. While mostly a monochromatic method, some materials exhibit a range of colors when treated with laser, as a result of complex physicochemical phenomena. Among such materials are stainless steel and titanium, some of the most industrially important metals.

Despite the great potential, the industrial adoption of color laser marking is almost non-existent due to its challenging *characterization*. In the absence of such a characterization, the relationship between the device's *design space* (laser parameters) and *performance space* (e.g. marked colors) is unknown. This relationship is too complex to capture with physics-based methods [Nánai et al. 1997]. Instead, the current practice finds design parameters that lead to “interesting colors” by trial and error measurements. These

primary colors are then used to mark simple motives and logos. This brute-force color *gamut exploration* scales poorly with the laser marking high-dimensional design space, resulting in neglecting some design parameters.

The goal of this work is to equip color laser marking with the same level of versatility found in color printers. Assuming a black-box model of the difficult device characterization, we design a measurement-based, data-driven performance space exploration. We explore different performance criteria including the color gamut and marking resolution by consecutive marking and measuring. For this, we uncover the process's *Pareto front* by formulating a multi-objective optimization problem and solving it using an evolutionary algorithm augmented by a Monte-Carlo approach. The optimization explores the hidden corners of our 7 dimensional design space in search of useful parameters that lead to a dense set of diverse, high-resolution colors. We also go far beyond the state of the art color image marking by introducing a complete color management workflow that takes an input image and laser-marks the closest approximation on metal surfaces. Our proposed color reproduction workflow adopts the principles of halftone-based color printing. It extracts a number of primary colors (akin to printer's inks) from the explored gamut and reproduces input colors by juxtaposing the extracted primaries next to each other in a controlled manner. Our fabricated color images enjoy high resolution, introduce no significant artifact, and demonstrate accurate color reproduction. The main contributions of this paper are:

- A discovery algorithm that automatically finds the desired design parameters of a black-box fabrication system.
- The first color-image reproduction workflow for laser marking on metals.

2 BACKGROUND AND RELATED WORK

Color Laser Marking. The coloration of different substrates using laser irradiation is an active field of research with a long history [Birnbaum 1965]. There are many color formation mechanisms employing different laser sources and different materials; see [Liu et al. 2019] for a recent review. Surface oxidation is one of these mechanisms where the heat (generated by a laser) facilitates the reaction of materials with oxygen¹. Oxidation-induced colors are believed to stem from multi-layer, heterogeneous mixture of structural colors (based on thin-film interference) [Del Pino et al. 2004] on one hand, and the traditional pigment-based color of oxides [Langlade et al. 1998] on the other hand. Despite a handful of initial efforts [Veiko et al. 2013], predicting the structure and composition of oxide layers is extremely difficult due to the complex thermodynamics of the laser marking process [Nánai et al. 1997]. Even with known material composition, predicting the surface color requires a challenging light-matter interaction model most likely based on an electromagnetic simulation [Auzinger et al. 2018]. Based on these observations, we adopt a black-box model of the process ruling out a physics-based prediction of the laser-induced composition of oxides.

For some popular metals, such as stainless steel and titanium, oxidation-based color laser marking has been extensively studied.

¹The same phenomenon turns motorcycles' exhaust pipes colorful.

This spans a range of laser-marked metals' behaviors, from electromechanical [Lawrence et al. 2013] to corrosion resistance properties [LeRcka et al. 2016]. More related to our work is a class of studies focused on the effect of various process parameters on the marked colors [Laakso et al. 2009]. Most of these works [Adams et al. 2014; Antończak et al. 2013, 2014] rely on sampling and marking the process parameter space uniformly. As laser marking is time and material consuming, these methods cannot cope with the dimensions of the design space and end up ignoring a large portion of process parameters. Moreover, we are unaware of any work on exploring the marking resolution, separate or jointly by color, a critical performance criterion for high-quality image marking.

It is worth noting that some empirical color discovery methods [Veiko et al. 2016] try to find different laser parameters that lead to the same color. But that color needs to be known *beforehand*. Moreover, these methods are restricted to interference-based colors within a limited range of laser energy and a limited number of parameters. This work introduces the first systematic color discovery algorithm for laser marking systems.

Performance Space Exploration. In computer graphics, formulating design problems based on multi-objective optimization and solving them by computing the Pareto front is known. Notable examples are simplifying procedural shaders [Sitthi-Amorn et al. 2011] or minimizing power consumption in real-time rendering [Wang et al. 2016]. In computational fabrication, exploring the performance space (or simply the *gamut*) of a process has attracted recent attention. With the advent of 3D additive technologies, these efforts are mainly focused on exploring the mechanical properties of 3D printed microstructures. As an example, Schumacher et al. [2015] precompute microstructures' performance space defined by mechanical metamaterial families in order to accelerate their heterogeneous topology optimization. They first populate the performance space by perturbing the initial designs (in the design space) and then fill the unpopulated regions of the performance space through either interpolation or inverse optimization. For a similar purpose, Zhu et al. [2017] combine a discrete, random perturbation of designs near the gamut boundaries with a continuous optimization that further expands the gamut by refining existing designs.

In a more general-purpose method, Schulz et al. [2018] further emphasize the importance of exploring the performance gamut's hypersurface (or Pareto front) instead of its hypervolume. A Pareto front captures a set of solutions in the performance space that are compromising different, potentially conflicting objectives.

Although these methods serve as important sources of inspiration, we cannot rely on any of them as they depend on closed-form, smooth characterization functions that connect the design and performance spaces. For example, the method of Schulz et al. [2018] requires a smooth (twice differentiable) forward characterization of the process and works only with continuous design parameters. Here, similar to Schulz et al. [2018], we cast our performance space exploration problem as a multi-objective optimization. But our solution relies on *online* marking, measuring of design-space parameters and navigating it using a modified evolutionary algorithm [Deb et al. 2002].

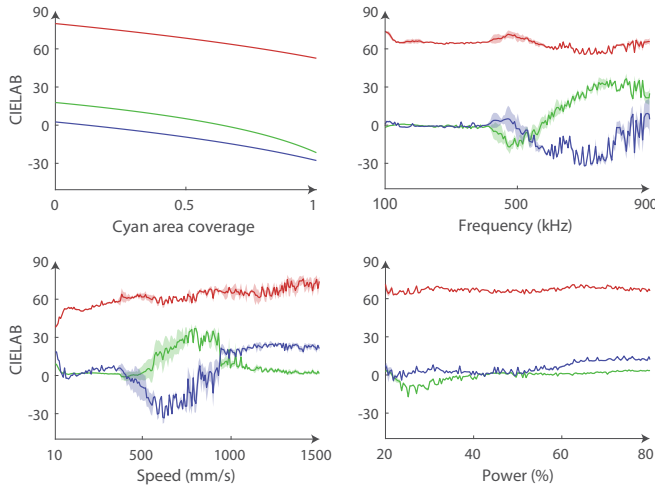


Fig. 2. The color response of our laser vs. that of a typical printer (top left). CIE L^* , a^* and b^* values are plotted in red, green and blue respectively. The laser-marked colors (on stainless steel AISI 304) are repeated three times. Their average (solid lines) and standard deviation (shaded region) are shown. Notice the non-smooth behavior of the laser-marked colors.

Custom Color Reproduction. Our work is closely related to color printing as our final objective is to reproduce color images through laser marking. Color reproduction, to the best of our knowledge, is absent in the laser marking literature. The closest practice is to mark simple logos and motives without systematic reproduction in mind. Our goal is to start from a source image (e.g., sRGB) and mark its most accurate approximation on metals using a custom color reproduction workflow. While the traditional color reproduction workflow is well validated for printing inks on paper, it falls short when dealing with unconventional circumstances, such as non-standard inks [Hersch et al. 2007] or substrates [Pjanic and Hersch 2013]. In such circumstances, custom color reproduction pipelines are required [Stollnitz et al. 1998]. We find these custom workflows to fit well to the problem at hand. In particular, as we will show in Section 4, the *juxtaposed halftone* color reproduction workflow can be adapted to our color reproduction setup. This workflow is designed to print with opaque inks and includes juxtaposed halftone synthesis [Babaei and Hersch 2012], color prediction [Babaei and Hersch 2016a] and color separation [Babaei and Hersch 2016b].

3 GAMUT EXPLORATION

Device characterization is the prerequisite for any color reproduction system including laser marking. In the absence of an analytical function that maps laser marking parameters to marked colors, we must rely on data-driven methods. In a first attempt, one can sample the design space, mark and measure the sampled design points, and construct a look-up table. This exhaustive strategy is subject to the curse of dimensionality given the relatively large number (7) of parameters involved in color laser marking. The fact that function evaluations require actual marking and measuring further slows down the process. Additionally, the non-smooth color response to laser parameters renders interpolation schemes ineffective. This

is shown in Figure 2 where color coordinates of marked patches may change abruptly in response to marking parameters. It is contrasted with the smooth response of a typical printer to its control parameters.

In this section, we propose a non-exhaustive performance space exploration of the laser marking system. Qualitatively speaking, our performance criteria favors *diverse*, *saturated* and *high-resolution* colors: the fundamental requirements for color images. Quantifying these objectives in Section 3.1, our main insight for solving this problem is to cast it as a multi-objective optimization (Section 3.2). Unlike a typical optimization, multi-objective optimization problems are evaluated based on multiple criteria. Very often, these criteria are in conflict. In our case, for example, some marked colors may be saturated but leave thick traces and lower the resolution. Hence, instead of a single optimal solution, there exists a set of optimal solutions, known as *Pareto optimal* solutions or *Pareto set*. The projected Pareto set into the performance space is called *Pareto front*. A member of the Pareto front is *not* dominated by any other point in the performance space in all criteria. In other words, it is more performant than all other points in at least one criterion.

Our goal is to uncover a dense set of Pareto-optimal solutions to the color laser marking problem with the above objectives. To this end, we adopt a multi-objective evolutionary algorithm, a successful tool for finding Pareto optimal solutions [Fonseca et al. 1993]. The algorithm, called non-dominated sorting genetic algorithm (NSGA-II) [Deb et al. 2002] is well suited to our model-free characterization function, with both discrete and continuous parameters. At the heart of this method, is a sorting algorithm based on the members' presence in multi-level Pareto fronts. As we show, the NSGA-II *non-dominated* sorting is insufficient for our specific problem due to our hue diversity objective. Thus, we resort to a Monte-Carlo approach on top of the non-dominated sorting and introduce a new sorting method based on *front frequencies*. We call this algorithm Monte-Carlo, multi-objective, genetic algorithm or MCMOGA for short.

3.1 Color Laser Marking Performance Criteria

As we plan to adopt halftoning for color reproduction (Section 4.2), we need a set of primary colors that cover both achromatic and chromatic axes. In a divide and conquer strategy, we separate the chromatic and achromatic (black and white) explorations, starting here with explaining the former. High chromatic performance requires saturated colors corresponding to larger radii in the CIECH color space shown in Figure 4 [Schanda 2007]. Furthermore, it requires colors that span a range of different hues. Such colors *mixed* with black and white (through halftoning) generate a highly populated color gamut that can be utilized for color image marking.

The goal of our multi-objective optimization is to, eventually, *maximize* the following criteria:

- (1) *Chromaticity*, as marked colors with large chroma produce more saturated color images.

$$f_C(a^*, b^*) = \sqrt{a^{*2} + b^{*2}} \quad (1)$$

where a^* and b^* are the color coordinates of the CIELAB color space [Wyszecki and Stiles 1982].

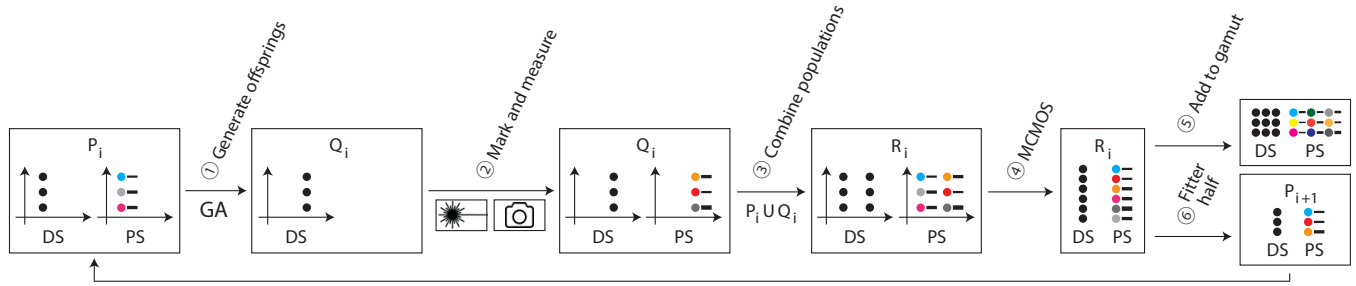


Fig. 3. One iteration of MCMOGA. ① The algorithm takes the starting population at iteration i (P_i) and uses the genetic algorithm to generate an offspring generation Q_i . ② Marking and measuring the design space (DS) yields the corresponding points in the performance space (PS). P_i and Q_i are combined into R_i ③ and sorted using the proposed MCMOS algorithm ④. R_i is added to the gamut ⑤ and its fitter half P_{i+1} is passed as the starting generation to the next iteration ⑥.

- (2) *Hue spread* (f_{HS}) ensures the presence of high-chromaticity colors at all hue angles.
- (3) *Resolution*; since we use a line-based halftoning, this criterion is evaluated by measuring the thickness of a line marked by a set of given laser parameters.

$$f_R(t) = \frac{1}{t} \quad (2)$$

where t is the line thickness.

- (4) *Performance space diversity* (PSD); it is measured for each point in the CIECH space as the reciprocal of the distance to its closest two neighbors.

$$f_{PSD}(p, P) = \frac{1}{\arg\min_{i,j} \|p - P_i\|^2 + \|p - P_j\|^2} \quad (3)$$

$i \neq j, 0 \leq i, j < |P|$

where p is the point to be scored and P its respective population in CIECH space.

- (5) *Design space diversity* (f_{DSD}); it is measured as the performance space diversity except in the design space.

Criteria (1) to (4) are measured in the performance space while criterion (5) is measured in the design space. Criteria (1) to (3) are the qualities we are directly seeking for laser marked images. Criterion (4) improves the convergence rate and criterion (5) helps avoiding local extrema by promoting *solo* points in the performance space.

3.2 Monte-Carlo Multi-objective Genetic Algorithm

Our algorithm, navigates the laser's design space in directions that lead to a dense Pareto set, i.e., the set of designs (laser parameters) that improves the above performance criteria. We start with a random population in the design space, mark it and measure its performance, then iteratively evolve it into a larger population with as many Pareto optimal solutions as possible. At each iteration, represented schematically in Figure 3, we promote the Pareto set of its population to be passed along to the next iterations using the genetic algorithm. We stop iterating when we do not observe a significant improvement in the Pareto front.

3.2.1 Monte-Carlo Multi-objective Sorting. As in almost any genetic algorithm (GA) method, a fitness measure should be assigned to each member of the population. Fitter solutions are selected by

GA and used to create the next generation. Given the difficulty of assigning a single fitness value to multi-criterion objectives, the non-dominated sorting algorithm [Deb et al. 2002] *sorts* the members of a population according to their presence in multi-level Pareto fronts (Appendix B). It starts with finding the *first non-dominated front*, i.e., all solutions in a population that belong to the Pareto front. This is done by comparing each solutions' performance objective by objective to every other solution in the population. If a solution is more performant than *all* other solutions in *at least* one criterion, it is labeled as a first-front solution. The second non-dominated front is computed by temporarily discarding the first front and repeating the above procedure. This procedure is continued until all members of the population are labeled with their respective fronts. This results in a number of disjoint subsets making up the whole population, each with its front label. Note that, in the spirit of the Pareto concept, members inside the same front are not sorted.

This type of sorting is not sufficient when considering the hue spread objective. This criterion helps the color gamut grow in all angular directions in a balanced manner. Without the hue spread criterion, the algorithm may explore some specific hue angles more than others resulting in a non-uniform growth of the chromaticity gamut. For achieving a gamut with balanced hue spread we cannot evaluate a single solution but rather in combination with other solutions. This can quickly lead to nontrivial computation: for 10 angular samples in a population of 200, we must perform $\binom{10}{200} \approx 10^{16}$ evaluations.

We avoid this combinatorial explosion by resorting to a Monte-Carlo method. A *hue wheel* splits the CIECH space into a random number of circular sectors (Figure 4). Within each sector, we sort the solutions using the described non-dominated sorting algorithm based on all criteria *except* the hue spread. We repeat this procedure each time with a randomly chosen number of sectors, and with a random angular offset. After each *turn* of the hue wheel, every individual solution is assigned to a potentially different front. At the end of this loop, every single solution is characterized by its *front frequency* vector that represents the frequency of its presence in the first front, second front and so on. We stop "turning" the Monte-Carlo hue wheel when the change in front frequencies is below a certain threshold. The population is sorted based on the frequency of their "top" fronts where a single first front is worth more than any

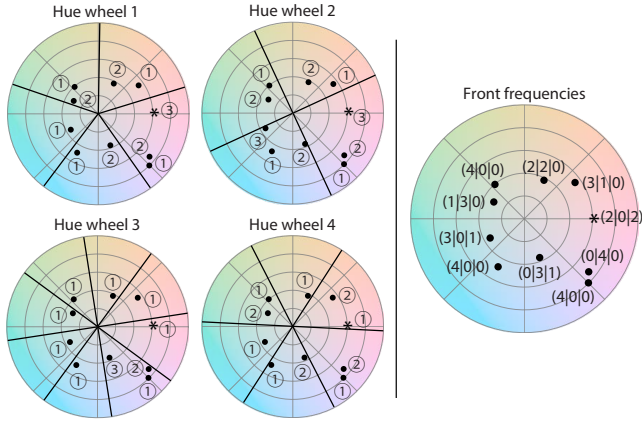


Fig. 4. Examples of different hue-wheel configurations (left) used in MCMOS. First, the points within each circular sector of each hue wheel are assigned to front labels (encircled numbers next to each point) using the conventional NDS algorithm. Next, all front labels for each point are counted to form the front frequencies (right). As an example, the point shown with a star has been assigned two times to the first front and two times to the third front. Notice that, for the same point, only the first two hue-wheels would not suffice as it would have been assigned only to the third front despite high potential for improving the gamut.

number of second fronts. This procedure is schematically shown in Figure 4. Note that multiple turns of the hue wheel, with different number of sectors and angular offsets, ensure all points are sorted in different configurations and are ranked in a proper way. In Figure 4, for example, it is easy to see that some points may not be sorted properly using a single hue-wheel configuration. The pseudo-code for the sorting algorithm, called Monte-Carlo multi-objective sorting (MCMOS), is shown in Algorithm 1.

3.2.2 Achromatic Exploration. In our chromatic exploration, the lightness values (CIE L^*) are discarded. We perform two separate explorations for black and white colors on the lightness axis. For the black colors, we *minimize* the criteria (1), thereby encouraging low chromaticity colors, and replace criterion (2) with lightness minimization. Exploring white colors is the same as the black colors except in (2) we maximize the lightness.

4 IMAGE REPRODUCTION

Once the performance space of our laser marking system is explored, it can be exploited for color image reproduction. We propose to adopt the principles of halftone-based color printing for color laser marking. For this, we find a number of *primary* colors (equivalent to printer's inks) that meet our resolution requirement and produce the largest color gamut (Section 4.1). We then build a color management workflow that takes input color images and marks the closest approximation using the selected color primaries (Section 4.2).

4.1 Primary Extraction

The primary extraction is concerned with selecting a set of colors that generates the maximum color gamut through halftoning. While, unlike printers, the number of primaries is not strictly limited, a

Algorithm 1: Monte-Carlo multi-objective sorting (MCMOS)

```

1 Input
2    $P$  : Population
3    $<$  : Partial order on  $P$ 
4    $t$  : Stopping threshold
5    $\alpha^{lim}$  : Offset limits for hue wheel
6    $C_{\#}^{lim}$  : Limits for number of circular sections
7 Output
8    $P^S$  : Sorted population  $P$ 
9 Procedure
10 begin
11   // Initialize front frequencies with 0
12    $\Theta \leftarrow 0$ ;
13   // Initialize frequencies difference
14    $\Delta_{\Theta} \leftarrow \infty$ ;
15   // Initialize prior front frequencies
16    $\Theta^{Pr} \leftarrow \Theta$ ;
17   while ( $\Delta_{\Theta} > t$ ) do
18     // Draw hue wheel parameters
19      $\alpha \leftarrow \text{random}(\alpha^{lim})$ ;
20      $C_{\#} \leftarrow \text{random}(C_{\#}^{lim})$ ;
21     // Split population into circular sections
22      $C \leftarrow \text{HueWheel}(P, \alpha, C_{\#})$ ;
23     // Iterate over circular sections
24     for  $c \in C$  do
25       // Compute front indices
26        $\delta \leftarrow \text{NDS}(c, <)$ ;
27       // Update front frequencies
28        $\Theta_{c, \delta}++$ ;
29     end
30     // Compute change in front frequencies
31      $\Delta_{\Theta} \leftarrow \|\Theta - \Theta^{Pr}\|_F$ ;
32     // Cache front frequencies state
33      $\Theta^{Pr} \leftarrow \Theta$ ;
34   end
35   // Sort points according to front frequencies
36    $P^S \leftarrow \text{Sort}(P, \Theta)$ ;
37 end

```

smaller number of primaries lead to improved marking time as they cause fewer switching delays (Section 5.1.1). Not all colors in the explored gamut can be considered for primary extraction. Thus, before primary extraction, we *prune* the gamut by excluding colors that: 1) don't satisfy our specified resolution requirement, 2) reveal low repeatability, and 3) exhibit non-uniformity.

Similar to the gamut exploration, we extract the achromatic and chromatic primaries separately. We start with the explored chromaticity gamut of the laser marking system composed of a discrete set of colors. We find the convex hull of this set in the CIE xy chromaticity space where $x = X/(X + Y + Z)$, $y = Y/(X + Y + Z)$ [Wyszecki and Stiles 1982]. The reason for applying the convex hull in the CIE xy is that, unlike CIE a^*b^* or CIECH, it is a linear space under halftoning. We experimentally confirm this linearity in Section 5.3, where we show that colors inside the convex hull can be reproduced through halftoning with high accuracy. The colors in the convex set give the largest area and therefore the largest chromatic gamut. In order to reduce the number of primaries, we can further exclude those members of the convex set that don't contribute to the gamut area significantly. The achromatic primary extraction selects the darkest and the brightest colors with negligible chromaticity from within the black and white explored gamuts, respectively.




Fig. 5. Laser marking's color reproduction workflow. ① An input image is mapped to the gamut of the laser marking system. ② In the color separation step, for each mapped color, the corresponding area coverages of each primary is computed (creating the gamut and color separation are built upon a color prediction model). ③ The continuous area-coverages are binarized and placed next to each other using a juxtaposed halftoning method. ④ The raster halftone images are converted into vectors.

4.2 Color Management Workflow

A color management workflow ensures color reproducibility across different imaging devices. The classic example is printing where the input images, from a camera for example, are printed as accurately as possible. Figure 5 sketches our proposed color reproduction workflow for color laser marking. Given an input color in a given color space, e.g., sRGB, we first ensure its reproducibility by mapping it into the color gamut of laser marking. The color separation computes the coverage of different laser primary colors which, when placed next to each other through halftoning, reproduce the input color.

A typical printer's color reproduction workflow generates different colors by spatial blending and superposition of multiple inks. Should we imitate such a workflow for the laser marking process, we would need to ensure both laser primaries *and* their superpositions are optimal. Exploring the design space for such an unlikely combination is significantly more difficult. Instead, we propose to place different primaries strictly next to each other resulting in a considerably simpler exploration where we search only for a set of suitable primaries (and not their superpositions). We build on an existing *juxtaposed* halftoning reproduction workflow [Babaei and Hersch 2012] developed for the color reproduction of opaque inks. In order to establish such a workflow, we synthesize juxtaposed halftones of extracted primaries (Section 4.2.1), predict the integral color of multi-primary halftones (Section 4.2.2), and numerically invert this prediction model in order to map input colors into primary halftones (Section 4.2.3).

4.2.1 Juxtaposed Halftoning. Color halftoning converts a continuous tone color image into a set of binary images each corresponding to one of the printer's inks. The discrete-line juxtaposed halftoning [Babaei and Hersch 2012] synthesizes these binary images in the form of lines and places them next to each other without overlapping. In the original algorithm designed for bitmap printers, using *digital lines* [Reveilles 1995] allows for subpixel thickness, low computational complexity, and, importantly to us, continuity. A continuous laser path ensures less switching delays, and therefore, faster marking with lower graininess caused by the two ends of each marked vector. As the original juxtaposed halftoning is designed for

raster devices, we need to transform the resulting raster images into vector representation suitable for our laser device. For this purpose, we use a *naive line* (a discrete line with unit thickness) as a mask and slide it on each halftone *layer* corresponding to each laser primary (Figure 5). This produces a list of vectors of different primaries which span the image plane and are sent to the laser device for marking.

4.2.2 Forward Color Prediction. The color prediction model has two roles in our proposed color management workflow. First, it constructs the color gamut generated by halftoning a set of primaries. It predicts the color of several thousands of halftones spanning the space of the relative area of primaries in each halftone, known as *area coverages*. The gamut surface is then fitted to this volumetric point cloud and is later used for gamut mapping. Second, the forward model is used in the color separation step that computes the area coverages of the primaries for any input colors to be reproduced.

We use the Yule-Nielsen (YN) prediction model to predict the multi-color, juxtaposed halftones of laser primaries. The Yule-Nielsen equation [Yule and Nielsen 1951] predicts the CIEX color coordinate (X_t) of a juxtaposed halftone as

$$X_t = \left(\sum_{i=1}^q a_i (X_i)^{1/n} \right)^n \quad (4)$$

where X_i is the CIEX value of the i -th primary, and a_i is its area coverage. The same equation applies for predicting CIEX and CIEY color coordinates. The exponent n , called the Yule-Nielsen n -value is a tuning parameter which we further discuss in Section 5.3.

4.2.3 Color Separation. Color separation builds on the forward prediction model to compute the particular primaries and their area coverages that reproduce a given color (inside the color gamut). As the YN model is not analytically invertible, color separation is carried out by optimization:

$$\underset{\mathbf{a}}{\operatorname{argmin}} \Delta E_{00}(\operatorname{Lab}(\operatorname{YN}(\mathbf{a})), \mathbf{c}) \quad (5)$$

$$\|\mathbf{a}\|_1 = 1, \quad \mathbf{a} \in [0, 1]^q$$

where \mathbf{c} is the target color in the CIELAB color space and \mathbf{a} is the optimization variable, i.e. the vector of area coverages of q primaries. As the CIEDE2000 color-difference formula [Sharma et al. 2005] is used for the distance metric, the modeled color using the YN model ($\operatorname{YN}(\mathbf{a})$) should be converted to CIELAB from CIEXYZ (denoted by function Lab in Equation 5). Equation 5 searches for an area coverage vector that, after being marked, results in the minimum distance to the target color. As we juxtapose different primaries, their relative area coverages should sum up to 1 and be non-negative.

5 RESULTS AND ANALYSIS

In this section, we present different analyses and evaluations of both gamut exploration and image reproduction.

5.1 Experimental Setup

5.1.1 Laser Marking Hardware and Parameters. We assemble a laser marking device (Figure 6) whose main components are a ytterbium fiber laser system (IPG Photonics YLPM-1-4x200-20-20) and a galvanometric scanner (Scanlab IntelliScan III 10). The laser system




Fig. 6. Our assembled laser marking device with annotated modules.

(20 W, 1064 nm) generates a laser beam which is redirected by the scanner's Galvo Mirror system to any desired spot on the substrate. Equipped with an infrared F-Theta lens ($f=163$ mm), the scanner is capable of imaging a planar field of 116×116 mm. For safety reasons, the whole laser marking system is covered in an enclosure within which the laser source is separated by aluminum plates. An air filtering system blocks small particles from spreading in the room. In most of our experiments we use 1 mm thick stainless steel type 1.4301 V2A (AISI 304) as a substrate. We use thicker substrates and another stainless steel alloy (AISI 430) in some experiments in Section 5.4. Note that color laser marking is also possible on titanium.

In this work, we study 7 marking parameters (Figure 7), including *laser parameters*:

- (1) *Frequency*: Defines the number of laser shots per second (1.6-1000 kHz, 100 Hz steps),
- (2) *Power*: Adjusts the output power per shot (0-100%, 256 steps),
- (3) *Pulse width*: Defines the duration of a single shot (4, 8, 14, 20, 30, 50, 100, 200 ns),

and *scanning parameters* that forms a *line cluster* with properties:

- (4) *Speed*: Defines the travel speed along a vector while marking (0-2000 mm/s, 1 mm/s steps),
- (5) *Line count*: Defines the number of lines in a cluster (1-20 lines, 1 line steps),
- (6) *Hatching*: Defines the distance between lines within a cluster (1-15 μ m, 1 μ m steps),
- (7) *Pass count*: Indicates the number of times a vector is marked (1-10 passes, 1 pass steps).

It is worth noting that, due to technical limitations of the laser source, laser parameters cannot be used at arbitrary combinations. Furthermore, it is not possible to vary laser parameters *on the fly*. For example, switching frequency and power takes 0.6 ms and 3 ms respectively; changing the pulse width takes about 2 seconds as it requires reestablishing the connection between the controller board and the laser. In our path planning, we therefore allow *switching*




Fig. 7. Visualization of laser (left) and scanner (right) parameters. The multi-pass line cluster in the diagram on right forms the final color.

delays after changing these parameters ensuring the laser source can properly adapt to the new parameters.

5.1.2 Performance Space Measurement. For each point in the design space, we need to measure its performance in order to decide how to use that point in our exploration framework. All performance criteria, apart from the *design space diversity*, can be evaluated by measuring the thickness of a marked line cluster and the color of a marked *patch*. This is performed in two stages. First, for measuring the cluster's thickness, we use a hand-held digital microscope (Reflecta DigiMicroscope USB 200). In a second step, we use the thickness of a given cluster to mark its corresponding patch by juxtaposing multiple clusters within the desired area.

Both hue and chromaticity, the pillars of our performance space exploration, are computed from CIELAB, a perceptual color space. We therefore perform a *colorimetric calibration* [Hong et al. 2001] for measuring the color of marked patches. The colorimetric calibration connects camera RGB signals to CIEXYZ coordinates through a form of regression. The CIEXYZ values then can be converted to the CIELAB coordinates using a set of well-known, analytical transformations [Wyszecki and Stiles 1982]. For training the regression, we use 121 printed color patches, with known spectra measured with an X-Rite i7 spectrophotometer, and obtain the *ground-truth* CIEXYZ values assuming D65 illumination.

We capture the same printed patches with a Nikon D750 DSLR camera (with macro lens Tamron SP 90mm F/2.8 Di) obtaining raw RGB signals that have been corrected for spatial and temporal light fluctuations. Our colorimetric calibration shows high accuracy on a test set of 16 printed patches with an average $\Delta E_{00} = 2.26$ and maximum 5.00. This calibration is therefore used to estimate the CIELAB color of marked patches.

The structural nature of oxide colors causes a significant change in their appearance depending on the viewing and illumination geometry. We observe that laser-marked colors appear most saturated at specular and near-specular geometries. Therefore, inspired by previous work on metallic prints [Pjanic and Hersch 2013], we confine our color reproduction to *non-diffuse* geometries. To this end, the stainless steel substrate is illuminated with a large, diffuse area-light tilted approximately 45° from the substrate's normal and captured with the camera with a similar angle.

5.2 Gamut Exploration Evaluation

For evaluating the proposed gamut exploration algorithm, we perform multiple runs while discarding different objectives during




Fig. 8. Color gamut evolution of a *full exploration* (with f_C , f_{HS} , f_R , f_{PSD} , f_{DSD}) on AISI 304 stainless steel.




Fig. 9. Color gamut evolution of a random exploration on AISI 304 stainless steel.

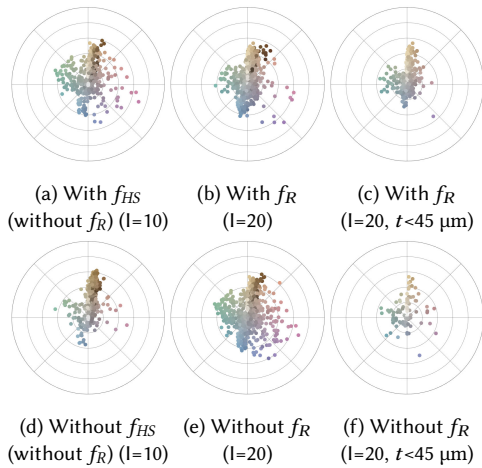


Fig. 10. Explored gamuts with different configurations on AISI 304 stainless steel.

different runs to show the objective's effect on the exploration behaviour. For a fair comparison, we always start with the same randomly generated initial population. All generations have the same population size of 100. For the hue wheel, we limit the random number of circular sections between 4 and 72 while the random angular offset α is between 0 and 2π . The stopping threshold for MCMOS is set to 0.001%.

Figure 8 demonstrates the evolution of our chromatic gamut, in hue-chroma polar diagram, when optimizing all performance criteria (referred to as the *full exploration*). Overall, we see a decent evolution of colors with a symmetric, dense color gamut. Interestingly, the purple to red regions are populated with a considerable delay, suggesting that some colors are more challenging to find than others.




Fig. 11. (a) Average thicknesses over iterations with (continuous line) and without (dashed line) f_R (only $t < 80 \mu\text{m}$ were considered). (b) Average reproduction error of the Yule-Nielsen model for different n -values. (c) Extraction of chromatic primaries.

Comparison with Random Marking. We compare the gamut of a set of randomly marked patches (Figure 9) with the results of our exploration. Compared to our full exploration (Figure 8), random marking does not lead to adequate gamut growth. As the random exploration does not include the resolution objective, it is more illustrative to compare Figure 9d to Figure 10e as they both feature the same number of samples and none of them includes the resolution objective. The stagnant behavior of random marking over time (Figure 9) and a lack of systematic resolution enhancement suggest that a very large number of samples is required to match the full gamut generated by our algorithm.

Effect of Monte-Carlo Hue Wheel. In order to evaluate the effectiveness of our Monte-Carlo hue wheel algorithm, we run two similar explorations where the only difference is that the hue-spread objective f_{HS} is enabled in one (Figure 10a) and disabled in the second (Figure 10d). We observe that the MC approach promotes the hue diversity resulting in a symmetric color gamut. Ignoring the Monte-Carlo method introduces a bias toward areas with high chromaticity. In Figure 10d, for example, since the initial population (shown in Figure 9a) has a large number of chromatic yellow members, this area is emphasized during the exploration.

Effect of Resolution Objective. Marking high-quality images requires a set of diverse, saturated colors which are placed next to each other at a high spatial resolution. This criterion is defined by f_R where we encourage design parameters that mark thin line clusters. A comparison of two explorations with equal number of iterations, one with (Figure 10b) and another without (Figure 10e) the thickness minimization reveals that this objective slows down the color gamut growth and the overall gamut area by disfavoring saturated but thick colors. Crucially, however, it generates a *denser* gamut at lower thicknesses, visible when comparing gamuts that include only colors with small thicknesses (Figures 10c and 10f). A dense color gamut is very important during primary *pruning* (Section 4.1). Furthermore, Figure 11a shows the average thicknesses of the whole population at each iteration. Unlike the exploration without thickness objective, the full exploration shows a steady decrease in the marked line thicknesses.

5.3 Color Image Reproduction Evaluation

We evaluate the proposed color reproduction pipeline by quantitative analysis and also a variety of full-color marked images.




Fig. 12. Multiple examples of original (left), gamut mapped (center) and marked images (right). The marked images, on AISI 304 stainless steel, are photographed in the non-diffuse mode. Images © Sebastian Cucerca, Steve Johnson, Albrecht Dürer, Vahid Babaei.

Primary Extraction. After the pruning step discussed in Section 4.1, we obtain a total of 6 primaries including a black and white primary and 4 chromatic primaries shown in Figure 11c. Their parameters are reported in Appendix A. In the pruning stage, we check the spatial and temporal repeatability by marking the candidate primaries at four different locations of the substrate and compare their colors pairwise. Primaries having an average ΔE_{00} higher than 4 among all comparisons are discarded. Also for the progressive primary discarding, we reduce the number of primaries until the gamut area drops by more than 10%. For resolution pruning, in this work, we keep colors with thickness $40 \pm 5 \mu\text{m}$. Additionally, we consider colors with thickness around $20 \mu\text{m}$ as juxtaposing two of them results in the target resolution.

Prediction Model Accuracy. For testing the accuracy of the Yule-Nielsen model, we mark the primaries and also 92 test patches with diverse area coverages of primaries. The resulting average ΔE_{00} error is 2.25 (Std=0.96, Min=0.50, Max=4.26) that demonstrates the high accuracy of our forward model. As we show in Figure 11b, the n -value equal to 1 works very well for our configuration, reducing our model to the widely known Neugebauer model [Rolleston and Balasubramanian 1993]. There are a handful of physical and empirical interpretations of the Yule-Nielsen n -value in literature [Lewandowski et al. 2006]. In the classic ink-on-paper prints, it accounts for the optical dot gain due to the lateral propagation of light inside the substrate [Hébert 2014]. Babaei and Hersch [2015]

showed that this parameter is responsible for shadowing and masking in metallic-ink halftones. From the optimal n -value for our setup we can infer that, as expected, the subsurface scattering in metal is very negligible. Furthermore, the marked primaries are very well leveled on the surface and cause no shadowing or masking.

Marking Color Images. Figure 12 presents different results generated by the proposed image reproduction pipeline with the chosen primaries. Comparing the marked images with their gamut-mapped counterparts, we can observe that the colors are reproduced faithfully. Furthermore, no significant artifacts are introduced in the laser-marked images. The gamut of the primary set allows marking diverse, vivid and relatively saturated colors as pointed by the image in the second row of Figure 12. Thanks to the high-resolution primaries, high spatial frequencies are preserved. This allows marking images with a vast level of details as shown in the bottom-right row of Figure 12.

5.4 Repeatability and Generalizability

In this section, we study the question of whether the marking parameters are transferable when using different marking settings or substrates. We start by marking a set of parameters reported in literature [Antończak et al. 2013] and show the results in Figure 13. Despite using highly similar hardware and materials, the reported colors are not reproducible on our setup. Also, color thicknesses vary significantly making them unsuitable for halftoning and therefore image marking.




Fig. 13. Marking parameters from Antończak et al. [2013], resulting in colors shown in the middle row (as reported in the original paper). Same parameters marked on the same material (AISI 304) using our device (bottom) lead to significant color differences (mean $\Delta E_{00} = 15.3$) and a huge thickness variation. We also show (top) the closest colors in our gamut to the reported colors in the middle.

In Table 1, we report color differences when marking on our setup a *general set* of parameters in different circumstances. The general set, consisting of 89 design points, is chosen to represent different colors in our explored gamut. We can see a pure repeatability test, on AISI 304 alloy, using the same marking settings leads to acceptable but not satisfactory accuracy. When we change the marking settings by using a thicker substrate (2 mm), and therefore exiting the focal plane, the repeatability worsens. Finally, using a different alloy of stainless steel (AISI 430) results in the worst repeatability.

For comparison, in Table 1, we also show the result of the same experiments performed using our 6 extracted primaries. We observe significantly higher accuracy in the pure repeatability experiment as the primaries have been pruned against this circumstance. Interestingly, the primaries show acceptable repeatability when marked out of focus indicating that the extracted primaries are robust against some perturbations. However, the larger deviation when using a new substrate suggests that the primaries cannot be used for marking images on new substrates accurately.

In Figure 14, we mark an image on a new substrate (stainless steel AISI 430) using primaries explored and extracted on our default substrate (AISI 304). While the image on the new substrate preserves the spatial details, most colors are significantly paler compared to the image marked on the default substrate. Thus, we perform a complete gamut exploration, primary extraction and color reproduction on the new substrate and show the results in Figure 14 (bottom row). Note that we cannot expect the same reproduction as on the default substrate since the color gamuts on two substrates are different. The gamut of the new substrate, calculated using a full exploration, is shown in Figure 15 (gamut obtained on the default substrate is shown in Figure 8).

Table 1. Repeatability errors of color laser marking in form of ΔE_{00} mean (and standard deviation).

Substrate A	AISI 304	AISI 304	AISI 304
Substrate B	AISI 304	2 mm thick AISI 304	AISI 430
General set	5.42 (5.63)	9.30 (6.27)	16.37 (7.50)
Primaries	1.96 (2.06)	6.46 (4.15)	12.33 (7.25)

5.5 Performance

An iteration of the gamut exploration with a population size of 100 takes around 30 minutes. This includes marking the single clusters,



Fig. 14. Marked images on AISI 304 (top left) and AISI 430 (top right) with the same primaries explored and extracted on AISI 304. A newly explored and extracted set of primaries on AISI 430 are used for marking an image: gamut-mapped (bottom left) and photograph (bottom right). Image © Piotr Didyk.




Fig. 15. Color gamut evolution of a full exploration (with $f_C, f_{HS}, f_R, f_{PSD}, f_{DSD}$) on AISI 430 stainless steel.

measuring their thicknesses with a handheld microscope, marking the corresponding patches with proper distances of clusters, and finally capturing them with our colorimetric camera. The manual measurement of cluster thicknesses is the bottleneck as it takes approximately 20 minutes. Computing a new generation using the MCMOGA takes only a few seconds in Matlab.

The marking time of an image is a function of the number of vectors (after halftone vectorization) and the marking speed of different primaries. A larger number of vector causes more switching delays, making the marking time highly dependent on the image content in addition to its size. For example, the two marked images in the top row, and right side of the bottom row of Figure 12, despite a comparable image size (7 by 11 cm), required around 18 and 30 million vectors, and roughly 3 and 5.5 hours of marking time, respectively.

6 CONCLUSION

We presented a computational framework that enables a novel application of laser marking: color image reproduction. Our method first characterizes the device using an evolutionary exploration of its performance space and then exploits that space for marking high-resolution, colorful images.

Pitfalls, Limitations and Outlook. Proper imaging was a challenging part of this research. We started the colorimetric calibration of the camera using the spectra of laser-marked patches (and not



Fig. 16. Image marked on AISI 304 and captured in non-diffuse (left) and diffuse (right) modes. Painting © Alessandro Allori.

printed colors on paper). Although our spectrophotometer can perform measurements that include the specular component, the measured spectra were inconsistent with reality. Fortunately, switching to the printed colors for calibrating the colorimetric characterization proved to be robust. Another imaging-related issue was measuring the thickness of the clusters using a camera. The efforts were not successful and led to the use of a handheld microscope that significantly slowed down our exploration.

A clear limitation of our method is the significant change of appearance from diffuse to non-diffuse configurations as shown in Figure 16. While amusing, producing plausible images in a wider range of viewing/illumination geometries is an important research direction and a formidable challenge. In the future, we plan to analyze the performance of our method in the case the pruning stage is integrated into the exploration using new objectives. We intentionally did not consider the physical phenomena behind the laser marking and stayed with a black-box approach. Yet we believe a combination of both physical and data-driven methods can make the color image marking task more scalable.

A PRIMARY PARAMETERS

The design and performance space parameters of the extracted set of primaries are shown in Table 2. In the project’s website, we also publish the intermediate exploration data, i.e., design parameters and performance measurements for all marked patches at all iterations of our full exploration.

Table 2. Design and performance space parameters of extracted primaries on AISI 304 (top segment) and AISI 430 (bottom segment).

Frequency	Power	Pulse width	Speed	Line count	Hatching	Pass count	CIELAB	t
kHz	%	ns	mm/s	#	μm	#	L^*, a^*, b^*	μm
650.7	29.5	20	1897	7	2	6	64.4, -15.9, -0.3	21
887.7	44.0	4	1964	3	7	4	66.1, -4.5, -13.0	41
973.6	36.0	100	280	4	15	1	64.6, 13.9, 5.7	43
973.6	38.5	200	280	4	3	1	73.5, 3.9, 36.8	43
597.3	24.0	100	129	8	3	2	55.9, 0.5, 1.7	41
160.8	42.5	100	1820	1	15	2	100.0, 0.0, 0.0	40
973.6	35.0	30	1248	20	2	5	66.0, 12.5, -8.6	43
586.6	47.0	100	609	1	12	2	71.8, 3.5, 11.2	22
798.4	36.0	30	1693	7	3	5	74.6, -9.5, 12.2	21
980.0	38.0	8	1815	7	3	5	64.3, -13.5, -13.9	22
580.2	37.0	4	2000	7	5	6	51.3, 3.3, -20.4	39
152.7	47.0	100	768	1	15	10	65.2, 1.3, -0.57	36
388.8	36.0	20	1488	1	5	2	100.0, 0.0, 0.0	42

B NON-DOMINATED SORTING ALGORITHM

The non-dominated sorting (NDS), depicted in Algorithm 2, computes the front indices of population members. We refer the reader to Deb et al. [2002] for further details.

Algorithm 2: Non-dominated Sorting Algorithm

```

1 Input
2    $P$  : Population
3    $<$  : Partial order on  $P$ 
4 Output
5    $\delta$  : Front indices
6 Procedure
7 begin
8   for ( $p \in P$ ) do
9      $S_p \leftarrow \emptyset$ ;
10     $n_p \leftarrow 0$ ;
11    for ( $q \in P$ ) do
12      // Check if  $p$  dominates  $q$ 
13      if ( $p < q$ ) then
14        // Add  $q$  to points dominated by  $p$ 
15         $S_p \leftarrow S_p \cup \{q\}$ ;
16      else if ( $q < p$ ) then
17        // Increase domination counter of  $p$ 
18         $n_p \leftarrow n_p + 1$ ;
19      end
20    end
21    if ( $n_p = 0$ ) then
22      //  $p$  is in front 1
23       $\delta \leftarrow 1$ ;
24       $F_1 \leftarrow F_1 \cup \{p\}$ ;
25    end
26  end
27  // Initialize front counter
28   $i \leftarrow 1$ ;
29  while ( $F_i \neq \emptyset$ ) do
30    // Set storing next front's members
31     $Q \leftarrow \emptyset$ ;
32    for ( $p \in F_i$ ) do
33      for ( $q \in S_p$ ) do
34         $n_q \leftarrow n_q - 1$ ;
35        if ( $n_q = 0$ ) then
36          // Add  $q$  to next front
37           $\delta_q \leftarrow i + 1$ ;
38           $Q \leftarrow Q \cup \{q\}$ ;
39        end
40      end
41    end
42     $i \leftarrow i + 1$ ;
43     $F_i \leftarrow Q$ ;
44  end
45 end

```

ACKNOWLEDGMENTS

The idea of this work started while the last author was at Computational Fabrication group of MIT. We especially thank Michael Foshey and Javier Ramos, and also Jonathan S. Gibbs and John Hart of MIT Mechanosynthesis group, for their valuable help with initial experiments. We thank Hossein Najafi from Oerilkon and Pitchaya Sitthi-Amorn for their helpful pointers to proper hardware and software, and Kristina Scherbaum for her continuous support during hardware assembly.

REFERENCES

- David Price Adams, Ryan D Murphy, David J Saiz, DA Hirschfeld, MA Rodriguez, Paul Gabriel Kotula, and BH Jared. 2014. Nanosecond pulsed laser irradiation of titanium: Oxide growth and effects on underlying metal. *Surface and Coatings Technology* 248 (2014), 38–45.
- Arkadiusz J Antończak, Dariusz Kocoń, Maciej Nowak, Paweł Kozioł, and Krzysztof M Abramski. 2013. Laser-induced colour marking - Sensitivity scaling for a stainless steel. In *Applied Surface Science*.
- Arkadiusz J Antończak, Bogusz Stępak, Paweł E Kozioł, and Krzysztof M Abramski. 2014. The influence of process parameters on the laser-induced coloring of titanium. In *Applied Physics A*.
- Thomas Auzinger, Wolfgang Heidrich, and Bernd Bickel. 2018. Computational design of nanostructural color for additive manufacturing. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 159.
- Vahid Babaei and Roger D Hersch. 2012. Juxtaposed color halftoning relying on discrete lines. *IEEE Transactions on Image Processing* 22, 2 (2012), 679–686.
- Vahid Babaei and Roger D Hersch. 2015. Yule-Nielsen based multi-angle reflectance prediction of metallic halftones. In *Color Imaging XX: Displaying, Processing, Hardcopy, and Applications*, Vol. 9395. International Society for Optics and Photonics, 93950H.
- Vahid Babaei and Roger D Hersch. 2016a. N-Ink Printer Characterization With Barycentric Subdivision. *IEEE Transactions on Image Processing* 25, 7 (2016), 3023–3031.
- Vahid Babaei and Roger D Hersch. 2016b. Color reproduction of metallic-ink images. *Journal of Imaging Science and Technology* 60, 3 (2016), 30503–1.
- Milton Birnbaum. 1965. Modulation of the reflectivity of semiconductors. *Journal of Applied Physics* 36, 2 (1965), 657–658.
- Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and TAMT Meyarivan. 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE transactions on evolutionary computation* 6, 2 (2002), 182–197.
- A Perez Del Pino, JM Fernández-Pradas, P Serra, and JL Morenza. 2004. Coloring of titanium through laser oxidation: comparative study with anodizing. *Surface and coatings technology* 187, 1 (2004), 106–112.
- Carlos M Fonseca, Peter J Fleming, et al. 1993. Genetic Algorithms for Multiobjective Optimization: Formulation, Discussion and Generalization.. In *Icga*, Vol. 93. Citeseer, 416–423.
- Mathieu Hébert. 2014. Yule-Nielsen effect in halftone prints: graphical analysis method and improvement of the Yule-Nielsen transform. In *Color Imaging XIX: Displaying, Processing, Hardcopy, and Applications*, Vol. 9015. International Society for Optics and Photonics, 90150R.
- Roger D Hersch, Philipp Donzé, and Sylvain Chosson. 2007. Color images visible under UV light. In *ACM Transactions on Graphics (TOG)*, Vol. 26. ACM, 75.
- Guowei Hong, M Ronnier Luo, and Peter A Rhodes. 2001. A study of digital camera colorimetric characterization based on polynomial modeling. *Color Research & Application: Endorsed by Inter-Society Color Council, The Colour Group (Great Britain), Canadian Society for Color, Color Science Association of Japan, Dutch Society for the Study of Color, The Swedish Colour Centre Foundation, Colour Society of Australia, Centre Français de la Couleur* 26, 1 (2001), 76–84.
- P Laakso, S Ruotsalainen, H Panssar, and R Penttilä. 2009. Relation of laser parameters in color marking of stainless steel. In *12th Conference on Laser Processing of Materials in the Nordic Countries*.
- C Langlade, AB Vannes, JM Krafft, and JR Martin. 1998. Surface modification and tribological behaviour of titanium and titanium alloys after YAG-laser treatments. *Surface and Coatings Technology* 100 (1998), 383–387.
- Samantha K Lawrence, David P Adams, David F Bahr, and Neville R Moody. 2013. Mechanical and electromechanical behavior of oxide coatings grown on stainless steel 304L by nanosecond pulsed laser irradiation. *Surface and Coatings Technology* 235 (2013), 860–866.
- KM ŁeRcka, AJ Antonczak, B Szubda, MR Wójcik, BD SteRpak, P Szymczyk, M Trzcinski, M Ozimek, and KM Abramski. 2016. Effects of laser-induced oxidation on the corrosion resistance of AISI 304 stainless steel. *Journal of Laser Applications* 28, 3 (2016), 032009.
- Achim Lewandowski, Marcus Ludl, Gerald Byrne, and Georg Dorffner. 2006. Applying the Yule-Nielsen equation with negative n. *JOSA A* 23, 8 (2006), 1827–1834.
- Huangang Liu, Wenxiong Lin, and Minghui Hong. 2019. Surface coloring by laser irradiation of solid substrates. *APL Photonics* 4, 5 (2019), 051101.
- László Nánai, Róbert Vajtai, and Thomas F George. 1997. Laser-induced oxidation of metals: state of the art. *Thin Solid Films* 298, 1-2 (1997), 160–164.
- Petar Pjanic and Roger D Hersch. 2013. Specular color imaging on a metallic substrate. In *Color and Imaging Conference*, Vol. 2013. Society for Imaging Science and Technology, 61–68.
- Jean-Pierre Reveilles. 1995. Combinatorial pieces in digital lines and planes. In *Vision geometry IV*, Vol. 2573. International Society for Optics and Photonics, 23–34.
- Robert Rolleston and Raja Balasubramanian. 1993. Accuracy of various types of Neugebauer model. In *Color and Imaging Conference*, Vol. 1993. Society for Imaging Science and Technology, 32–37.
- János Schanda. 2007. *Colorimetry: Understanding the CIE System*. John Wiley & Sons.
- Adriana Schulz, Harrison Wang, Eitan Crinspun, Justin Solomon, and Wojciech Matusik. 2018. Interactive exploration of design trade-offs. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 131.
- Christian Schumacher, Bernd Bickel, Jan Rys, Steve Marschner, Chiara Daraio, and Markus Gross. 2015. Microstructures to control elasticity in 3D printing. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 136.
- Gaurav Sharma, Wencheng Wu, and Edul N Dalal. 2005. The CIEDE2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations. *Color research and application* 30, 1 (2005), 21–30.
- Pitchaya Sitthi-Amorn, Nicholas Modly, Westley Weimer, and Jason Lawrence. 2011. Genetic programming for shader simplification. *ACM Transactions on Graphics (TOG)* 30, 6 (2011), 1–12.
- Eric J Stollnitz, Victor Ostromoukhov, and David H Salesin. 1998. Reproducing color images using custom inks. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*. ACM, 267–274.
- Vadim Veiko, Galina Odintsova, Elena Gorbunova, Eduard Ageev, Alexandr Shimko, Yulia Karlagina, and Yaroslava Andreeva. 2016. Development of complete color palette based on spectrophotometric measurements of steel oxidation results for enhancement of color laser marking technology. *Materials & Design* 89 (2016), 684–688.
- VP Veiko, AA Slobodov, and GV Odintsova. 2013. Availability of methods of chemical thermodynamics and kinetics for the analysis of chemical transformations on metal surfaces under pulsed laser action. *Laser Physics* 23, 6 (2013), 066001.
- Rui Wang, Bowen Yu, Julio Marco, Tianlei Hu, Diego Gutierrez, and Hujun Bao. 2016. Real-time rendering on a power budget. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 1–11.
- Gunter Wyszecki and Walter Stanley Stiles. 1982. *Color Science*. Vol. 8. Wiley New York.
- JAC Yule and WJ Nielsen. 1951. The penetration of light into paper and its effect on halftone reproduction. In *Proc. TAGA*, Vol. 3. 65–76.
- Bo Zhu, Mélina Skouras, Desai Chen, and Wojciech Matusik. 2017. Two-scale topology optimization with microstructures. *ACM Transactions on Graphics (TOG)* 36, 5 (2017), 164.