

Lessons from research in color science on the bleeding edge

Giordano Bruno Beretta; Palo Alto, California/USA

Abstract

The end of life is a good time to look back on what I have learned in the past 50 years and share my lessons. Except for stints in engineering and marketing, I have worked primarily in research labs. Although I was mostly in an imaging lab, de facto my interest has been predominantly in color science. I have worked in industry, but on the side, I have volunteered for national science foundations, learned societies, and patent offices. The main takeaway is that life in research is not smooth, you have to be resilient to setbacks, and be well-connected.

This article does not describe an experiment, consequently, it has an unusual structure. In the first section, I will describe the evolution of the research ecosystem in Silicon Valley, then I describe the many vicissitudes of a specific color interest. Finally, I share a few words on the Color Imaging conference at the Electronic Imaging symposium.

The ecosystem Before 1990

I finished my doctorate during the Cold War, when, after the Sputnik crisis, one of the main drivers for research in Silicon Valley was to outbrain the Russians. In industrial research labs, we had as a goal to instigate paradigm shifts, *i.e.*, change the way people work and live by inventing radically new technologies and work models, following McLuhan's slogan that *instead of saving work, labor-saving devices permit everybody to do their own work* [18, p. 36].

While we worked primarily for the product divisions of the company owning the lab, we also received direct projects from the federal government to produce gizmos that were not available on the market. Such g-projects were *cost plus*, meaning that when the gizmo was delivered, the government paid the cost of producing the gizmo plus a percentage as profit to finance further research.

A researcher would work for the glory, not for the money. The typical salary was the arithmetic mean of the salary of an assistant professor at the local university and that of an engineer in a product division. While labor was quite cheap, the cost of running a lab was high: each researcher had a comfortable office and the typical workstation bill of materials was 20% more than a year's salary, and used 4000 watts of energy [21].

At that time, companies had pension plans and the executives had long-term plans, therefore investing in expensive research on disruptive technologies made sense. We thought that after we had a prototype for a new technology, it would take 5 years to commercialize it: 1 year to sell the idea and 4 years to productize it. In retrospect, we were far off: it would take 15 years to sell the idea to the executives and less than a year to productize it because our prototypes were shrink-wrap ready.

We were all inspired by Doug Engelbart's demo [12], in

which he demonstrated windows, mice, hypertext, networking, collaborative editing, versioning, etc. Because in the left hand he had a keyset (chorded keyboard) and in the right a mouse, he was said to be *dealing lightning with both hands*. Decades later the demo was renamed to *the mother of all demos*. In 1968, in research, personal computing was a well-known goal, advocated by J.C.R. «Lick» Licklider, who had a profound impact on how computer science research was conducted [28].

While we were afraid of being years behind, our executives had personal secretaries to whom they would dictate memos, which the secretaries would send via interoffice mail to the typist pool. For the managers, progress was using a copy center instead of making carbon copies (cc)—they would have never been caught dead typing a memo or standing at a copy machine. This mental gap was why it took well over a decade for a prototype to move out of the lab. While the company had a networked office system with department printers, scanners, email, and document servers, our executives decided to go in the typewriter business. Gary Starkweather, the inventor of the laser printer, called it *fumbling the future* [23]. While in the lab documents contained active images [17], almost a decade later executives still claimed nobody would ever use or process an image on an office computer [1, p. 281].

After 1990

When the Cold War ended and the Berlin Wall came down, there was a radical change in industrial research in Silicon Valley. The government's cost plus jobs almost disappeared, the cost of workstations dropped by an order of magnitude, and the energy consumption by 2 orders. Enterprises replaced pension plans with individual retirement accounts invested in mutual funds, which have a short-term horizon. Therefore, executives focused on quarterly results instead of decades out. There was no longer an incentive for shifting paradigms.

Companies started perceiving research as a cost instead of an investment, and costs are something that must be curtailed. Lab managers, who before were active researchers, were replaced by business managers with MBAs, who were expected to quickly monetize the research [26]. The new managers were rewarded for quick transfers from the lab to business divisions, so they recited mantras like *do not take initiative, listen to the divisions*.

Concomitantly, the industry started to transition from technology innovation to supply chain management, which led to cuts in engineering, and the researchers ended up becoming stop-gap implementers. As an aside, while before in international trade discussions the U.S. emphasized industrial products, after the end of the Cold War its emphasis shifted to agriculture, which initially caught trade partners by surprise.

Today, in Silicon Valley researchers are often embedded in product divisions, working elbow-to-elbow on long tables with the engineers, but they report to a central CTO or a VP of research

and are not part of the engineering teams. On one side, this allows them to glean important hard problems with which the engineers are grappling and get inspired for new technologies. On the other side, when engineers get stuck with a problem for which there is no clear solution on the web, they can informally ask the local researcher for a piece of advice.

Research vs. engineering

Why is there research? Technological progress is in the air and most people are on the same page. For example, the pharmaceutical industry is very secretive because the research is particularly risky and highly rewarding. However, when a new drug has been developed, all companies working on it file their patent applications within the same week.

Profit is maximized by taking risks, which means by inventing new paradigms that allow people to do their own work and increase worker productivity. This requires knowledge.

There are two types of knowledge, tacit and explicit. Tacit knowledge consists of mental models, beliefs, and perspectives so ingrained that we take them for granted and therefore cannot easily articulate them. Tacit knowledge is deeply rooted in action and an individual's commitment to a specific context. Explicit knowledge is formal and systematic; it can easily be communicated and shared in product specifications, a scientific formula, or a computer program.

In engineering, we learn explicit knowledge and use it to create less expensive or better products. In research, the scientist considers a work process and invents new technologies to perform that work in a new more efficient way. Research is accomplished by performing the work in the conventional way embedded in a team of masters in the art. This way, researchers acquire the master's tacit knowledge. Then they use the scientific process to make this knowledge explicit by describing it in a report and teaching their coworkers. The team then builds a prototype and creates the engineering instructions. By figuring out how to fabricate a better product, the team creates new tacit knowledge that is used to design the follow-up product. This is the knowledge-creating organization described by Ikujiro Nonaka [20] and illustrated in Fig. 1.




Figure 1. The knowledge creating organization.

How is this accomplished? In Lick's time, a new paradigm is an emerging property, *i.e.*, comes from all known knowledge augmented by tacit knowledge. The question is how to emerge a new property.

When personal computers first became widely used, software was shipped with user manuals and training materials. A research topic in human-computer interfaces (HCI) was how to

teach how to use the software. For example, some people learn better by example and some people learn better to clear step-by-step instructions. A diagram like that in Fig. 2 emerged, with the emphasis on the column of cognitive styles.

motor	cognitive style	move minds	move molecules	shape volume
inner (god)	field dependent (global context)	art	science	architecture
speculative design				
outer (user)	field-independent (analytical, detail oriented)	design	engineering	structural engineering
entrepreneurialism				
greed (self)	type A	art director	consultant	building contractor

Figure 2. Some people are better at thinking by analogies, others by logic.

Some people are better at grasping the big picture (global context) and thinking in analogies: it is a field-dependent cognitive style. Other people are very analytical and are better at developing how a detail follows from another detail: their cognitive style is field-independent. People who master both cognitive styles are called *speculative designers*, and these are the people research labs try to hire, picking them from a short list of the 100 best computer scientists in the world.

The cognitive styles are a matter of degree and a speculative designer can be driven more by an inner motor or an outer motor. In the first lab in which I worked, the manager had researchers work in pairs, one better at the big picture and one more analytical. Through critical dialogue, this also encourages serendipity.

Bringing forth a substantial paradigm change requires knowledge in many contexts. Forming the best researcher pairs is not sufficient. The lab manager has to create an environment in which the researchers work synergistically like in a hive: the sum of the pairs is a new entity that is a different and superior entity to its parts.

This can be further accelerated by leveraging historical tacit knowledge in addition to going out and finding it as an apprentice. In the case of color reproduction, in the past, the company had researchers who invented color offset printing and color photography. By taking them out of retirement and having them work elbow-to-elbow with us young researchers, we could grasp the implicit knowledge much faster. Additionally, we apprenticed by working on offset print jobs, like printing color articles and proceedings [24, 25].

At every step of the way, both engineers and researchers have to ask themselves how they can do it better based on their unique experience and knowledge. Bill Hewlett used to attend project and lab reviews as often as he could. When a presentation was too rosy, he had a standard question: *Of all the avenues you tried, how many were successful?* In product divisions he expected 50%, in HP Labs he expected 15%; if the answer was higher, he would reply: *You did not take enough risk!* Engineers are expected to be on the cutting edge, but researchers must be on the bleeding edge.

A color research adventure

My first research was in computational geometry, and I was responsible for the maintenance of a novel hierarchical VLSI design rule checker (DRC) that somebody else wrote. At that time chips were full-custom and the workstations were hardly able to

process the huge layouts. When a designer finished a layout, he gave it to me to run the DRC, which could take a couple of days.

When the first layouts were ready, the DRC found hundreds of thousands of errors, which made it useless. My manager asked me to fix the program so it flags only the dozen expected errors. When I examined the layouts, I discovered that the flagged errors were real: there must be something else going on because those VLSI designers were the best in the world. After the author of the layout program did not find an error in his code, I decided to go and interview the chip designers to acquire tacit knowledge.

I asked them to display on their monitor a particular cell in the layout of their chip and asked them to explain it to me. To my surprise, nobody was able to just point to locations in the cell and describe how it worked. They had to remove some geometry, explain the rest, then undo the removal and iterate the process. One designer had a very hard time pointing to transistors. Transistors are the active devices in the circuit; thus, when VLSI designers look at the layout of a cell, they can quickly recognize the cell's function by matching the pattern formed by the gates.

With that particular semiconductor process, a transistor was represented with a green channel in the diffusion layer intersected by a red polysilicon rectangle [19]. Is that not something that is hard to interpret for chromatically challenged people? When I asked the designers to take a color vision test, they all refused (only 20 years later one of them admitted to being a full dichromate). To be on the safe side, I assumed all VLSI designers were chromatically challenged and set out to redesign the color palette used in the interactive layout editor.

Not knowing anything about color, I walked over to the color science lab and asked for advice. I was handed a copy of Wyszecki and Stiles [29] and advised, that all answers were in there. After reading a good portion of the book, using color map animation, I wrote a tool to interactively edit a color palette by dragging points in a chromaticity diagram while the layout editor was active, rendering the layout with the new color palette. When I plugged the new color palette into the layout tool and let the VLSI designers revise their cells, the number of design rule errors went down to the expected dozen per chip—after all, these VLSI designers were the best of the best!

Be persistent

To be honest, there was a lot of work writing the tools and integrating them [4] with the VLSI design system and the operating system. In our research lab, we believed that the first implementation has many rough edges because we were exploring and trying out alternatives. The praxis in the lab was to persist after the successful first release and rewrite the system from scratch with what we had learned. I persisted some more and devised an even better color palette [5], further speeding up the VLSI design process. The designers were now able to run the DRC by themselves and I no longer had to provide this service.

This is an example of improving an illustration by designing a palette in which the colors are more distinct and hence the illustration more readable. There are many visual impairments in addition to color vision deficiency, for example, the effects of aging with the lens becoming opaque (cataract). Color vision deficiency occurs mostly in males; in females, the opposite condition, called tetrachromacy, produces a completely different ordering of the colors in a palette, like the Farnsworth-Munsell 100 hue test.

When machine learning algorithms are used, it is important to note that a sensor is quite different from the human eye. For example, in automotive systems, the colors are typically RCCB instead of RGGGB, where C stands for clear. Also, street signals are designed for humans and for automotive systems they can be flickering when LEDs are used. An algorithm cannot be trained on the ground-truth generated by a naive human.

Keep your ears open: serendipity

One day I was working in my office when I heard a PR person giving a tour of the lab to an important guest. My office door was ajar, so I could not see, just hear. The guest told the host that the artist who painted these pictures went through an amazing transformation, similar to the painter William Turner when he crossed the Alps and arrived in Italy.

I was puzzled because there were no paintings in the area where my office was, so when they left, I went out to figure out what they were looking at. The only items on the walls were the layout check-plots the designers used to hang outside their doors for critique and discussions with their colleagues. Then it dawned on me: some of the check-plots were produced with the original color palette and any updates were pinned next to the older plots, but printed with the new color palette.

After some thinking, I realized that in a way VLSI layout check-plots look like complex Max Bill paintings. Then I remembered that a few years earlier, in 1981, I had taken a docent tour of a retrospective of Georges Vantongerloo, who selected the colors for his abstract paintings on trajectories in a color space. Maybe, the visitor was saying that the new color palette I built algorithmically led to more pleasing plots.

I went to visit the in-house graphic artist who was creating the slides for external presentations and asked him how he designed the color palette for a slide deck. He lamented that it was very difficult because we did the editors for him, but there was little know-how in how to design color palettes on computers. I returned to my office and implemented two tools for creating color palettes, one based on symmetries in a color space and one consisting of a database of palettes of famous artists, pigment available in a region at a given time, and emotions, based on Shigenobu Kobayashi's *Color Image Scale* [16]. I gave the new tools to the graphic artist to try out.

Shortly thereafter, two industrial designers from the product division came to visit me. They had created the desktop (icons, window decorations, etc.) for the next-generation workstation product, and just a couple of months before the shipping date the management had suddenly decided to use a color monitor instead of a black-and-white one. They were in panic mode and found out that I had some tools to design color palettes. I sat down with the designers and taught them how to use my tools while observing how they used them and updating the tools in real-time to improve them.

Unfortunately, their managers decided that color would be used only for the desktop, not in the editors for producing color documents. There was no transfer, but the company luckily decided to patent the tools. After a few years, when personal computers started having color monitors, software companies augmented their graphic software with color support and licensed my patents, which ended up being lucrative [3, 6, 11].

Do not assume others are as good as you are

In our lab, to do this research, we had to implement color management systems. We thought it was not a big deal because our retired colleagues who had invented color offset printing told us how to do it. When I transitioned to a new company, I realized that others were struggling with basic problems.

The product engineers sent me a few print samples from a new printer and asked me for advice to make the inks more vivid. They had a plot of a CIE 1931 chromaticity diagram with a triangle containing the spectral locus and wanted inks that had a gamut as big as the triangle. The inks were highly fluorescent and the images looked unreal.

The first problem was that using 24-bit color representations, there was insufficient precision and the images were contoured. Second, there is no need to be able to reproduce all perceivable colors. On my visit to Tokyo, I took some copies of Fig. 2(3.3.9) in the Wyszecki book [29, p. 167], a coffee filter for an office coffee maker, and a small potato. I held up the coffee filter with the potato inside and explained they were trying to reproduce the whole filter, while object colors were only in the potato. The colors produced by a digital scanner, a camera, or graphics software on a color display, are even only in a smaller solid inside the potato. The demo worked and a few weeks later I received printouts using new inks that looked much better.

A couple of years later the economic bubble in Japan burst, and I had to find a new job. With the potato in mind, I measured the relatively small gamut of the paints available to Flemish painters and compared it with that of contemporary printers. When you look at the original painting in a gallery, the perceived dynamic range is much larger than expected. My wife, who was an accomplished amateur painter, explained to me the techniques, her tacit knowledge. I prepared a job interview presentation proposing to invent an algorithm to segment an image, determine the semantics of each segment, and apply the painter's techniques.

At the first interview, I was chased out of the room and did not get the job. I realized that in other companies the engineers viewed color reproduction as finding the best mapping for an aperture color on a source device to that on a destination device. People were studying look-up table representations and had no concept of complex color images.

I ended up switching careers to image and video compression, commercial print workflow, and document management, all in black-and-white or multiple channels. However, color problems do not disappear when nobody looks at them. Problems of the readability of color foreground text of a colored background come up all the time; in document management, this requires the automatic correction of large numbers of documents.

Over time, I build a well-architected library allowing me to interchangeably model using physics (spectra), psychophysics (colorimetry), vision (absorption curves), and cognition (color terms). When I came across new results, I could incorporate them into the relevant module and all my tools inherited the new findings. For example, I could build models based on color terms in the Coloroid system [7] to assess readability.

Misoneism

Returning to mapping aperture colors for color reproduction, images reproduced without color management do not necessarily

look wrong, just inferior. The reason is that the color constancy mechanism in the human visual system can compensate to a certain degree for incorrect color rendering. I thought that for office and consumer applications if individual colors are not reproduced with high fidelity, at least the palette of colors in an image should retain its integrity when it is reproduced.

I thought that *color integrity* entailed two postulates:

1. foveal colors should not cross name boundaries
2. the error vectors should have a uniform flux

In the mid-1980s, Maureen Stone implemented a color selection tool in which she had a specification method inspired by the ISCC-NBS method of designating colors [15] that allowed multiple adjectives for any color term. The tool was very successful, and I observed how people used it. My new intuition was that after reproduction, colors should not change the name.

The second intuition was that very roughly, color constancy works by finding the brightest color, declaring it *white*, and normalizing to it. If a reproduction had multiple whites, the human visual system would not know on which one to normalize and the chromatic adaptation would fail.

I proposed to my manager a psychophysics experiment to assess if this intuition has value. He told me that the mantra of senior management is *do not take initiative, listen to the divisions*, therefore, as an entry-level worker, I should never propose new ideas. Furthermore, there were several vicious turf fights for color research, so I should stick to the compression of multi-channel images and never mention the word *color*.

In May 1997 I presented a paper on image compression at the AIC conference in Kyoto [8]. One afternoon, Robert Hunt asked me if I could spare an hour to do a deep dive into my paper. At the end, I asked him if he could comment on my idea of color integrity to assess if it would make sense. After listening, he replied that since my employer is not interested, I should present it in a panel discussion and let somebody else do the experiment. He arranged for such a panel discussion at the Color Imaging Conference taking place 6 months later in Scottsdale.

The panel consisted of distinguished color scientists, except for one product manager from a large company, who did not have a color science background. He interrupted every intervention with a loud voice, stating that his company would solve all color reproduction problems, and that we should not talk about issues. The moderator was not able to control this disruptive person and terminated the panel discussion.

A decade later I was working on simulators of commercial print workflows to discover errors before starting a job. One job type was very large prints for in-store displays. To reduce the print time, it was desirable to use the fastest halftoning algorithm not producing visible artifacts and if possible omit the black printer. The subjects were clothing, like denim, and cosmetics, like faces with heavy make-up. Remembering the color integrity intuition, I designed a process that would print a small color scale of critical colors, apply the print simulation, and use the color management system to calculate the rendered color scale. The print was acceptable when there were no transpositions [2, 9]. This method can be substantially improved by extending it to tetrachromacy.

In summary, when everybody is focused on reproducing single colors, it is difficult to introduce methods based on color palettes, but I persisted and many years later I got a patent and

even used the method to paint our house [10]. However, decades have passed and today this work is irrelevant, as much more sophisticated methods have been developed to study color constancy using deep neural networks for augmented and virtual reality applications [13].

The Color Conference at the EI Symposium

After the invention of photography, there has been intensive research on masking techniques for automatic color reproduction. After in the 1930s and 1940s color science developed the fundamental principles, color imaging evolved very rapidly with color offset, color photography, and color television. These technologies became so good, that the various companies stopped investing in color imaging research, relying on honing the new technologies.

It was only in the mid-eighties, that the computer industry was ready to embrace color science. Unfortunately, the color imaging scientists had retired by that time, and the art was lost. It was not easily possible to bootstrap the knowledge, because color science is not just about physics but also about psychophysics, statistics, cognitive sciences, etc., requiring both depth and breadth of expertise.

Some of us were lucky, and we were able to hire as consultants our companies' retired color scientists. The bounty was sufficiently rich that the Inter-Society Color Council (ISCC) organized a series of workshops in Williamsburg to spread the knowledge. Yet, colonial Williamsburg was very far from Silicon Valley, where the computer industry was.

The Society for Imaging Sciences and Technology (IS&T)—which was mostly about photography at the time—had a digital color imaging session at their Annual Meeting. SPIE, the international society for optics and photonics, was founded in 1955 to advance light-based technologies by organizing events in which members can exchange notes and knowledge. The SPIE had a mega-event called Photonics West and with the IS&T launched a Symposium on Electronic Imaging—Science and Technology (EI).

In 1993, a group that had met in Williamsburg started an EI conference on *Color Imaging: Device-Independent Color, Color Hard Copy and Graphic Arts*. The conference chair was Jan Bares and I helped him behind the curtain to study the submitted manuscripts to determine the emergent properties, create sessions, and order the papers in each session to form that linear thread that Gottfried Wilhelm Leibniz called a *filum Ariadne*.

By 1997, we had developed a very successful formula for the conference. One of the challenges of Photonics West was that there were about 15,000 attendees, many of them presenters, and it was difficult to keep them in the room, therefore, the *filum Ariadne* was key to our success.

We maintained a database of about 1,500 color imaging researchers, in which we were not only tracking their contact information but also their interests and current research topics. This allowed us to perform mailings of the call for papers with personalized cover letters, thus obtaining a higher submission rate than usual for conferences.

We invested significant time in carefully studying the submitted manuscripts and determine the emergent property and a path there. This allowed us to compile sessions that were telling stories about the future, thus keeping the audience in the room

for the entire session, mitigating conference hopping. At the beginning of each session, we put an invited paper, possibly with the world's top expert in the session's topic, and with the task to present an overview of the new results in the field.

We also motivated the session chairs to work hard. They were not allowed to just read the names of the presenters in the session: everybody in the audience can do that for themselves. Instead, the session chairs had to start each session with a brief overview of what was to come. At the beginning of each paper presentation, they had to bridge to the previous paper, to create a continuity that allows the audience to grasp the big picture. This was reinforced at the end of the session with a recapitulation of all presented papers and their connection.

The reason I worked behind the curtain was that I worked as an entry-level technical worker debugging firmware in facsimile machines. I was allowed to attend some conferences, but not chair them. Later, after chairing the additional 1998 EI conference in Zurich, Jan Bares had to step down, and I was pressured to be the chair. Since I was not allowed to, I convinced Reiner Eschbach to become a co-chair, so I could tell my manager that I was just a sidekick and Reiner was doing all the work. However, I kept being under very strong pressure to drop out of the conference.

One of the features of the conference was that each year we held a panel discussion on a controversial topic. The moderator was John Michaelis and the panel members were the foremost authorities in the field, so they could not be shot down. Although it took place in the evening after dinner, the room was always packed full and the discussions were very passionate. For example, one year the topic was electronic publishing, which was just emerging. In the discussion, we agreed that *electronic publishing* referred to taking the publishing workflow and replacing at every step a mechanical process with a digital step. In contrast, *digital publishing* referred to completely new workflows made possible by digital technologies. Because of the stature of the participants and the conference, this terminology was widely adopted.

At the beginning of the new millennium, for the 2000 symposium at the Color Imaging conference (3963), a management transition from Beretta/Eschbach to Eschbach/Marcu occurred. Later, Alessandro Rizzi and Shoji Tominaga joined. At some point, the conference changed its name to *Color Imaging: Displaying, Processing, Hardcopy, and Applications*.

One of the hallmarks of SPIE conferences is that the emphasis is on the timeliness of communication, not on scholarly perfection. In this spirit, we encouraged authors to present their most current work, even if was not yet finished. This was further cemented with a session on *The Dark Side of Color* introduced by Alessandro Rizzi, which would have been a much better venue for my proposal on color integrity than the panel discussion at the Color Imaging Conference in Scottsdale.

As a *segue*, despite years of battle for me to drop out of EI, a senior manager not in my food chain, in 1999 had me start a new EI conference on Internet Imaging for a principal researcher who was supposed to take it over the following year. The reason I was pressured was that two years earlier I had co-authored a paper on multiple categorization for collaborative annotation of images on the internet [27]. I gathered two dozen eminent researchers attending EI and invited them to a room for a few hours to establish the emergent properties in internet imaging. The result was benchmarking image retrieval algorithms [14]. Unfortunately, my

manager called it unfettered research, the conference chair never materialized, and the conference died.

This happened again in 2010, after I wrote an EI paper on GPU processing [22]: I was tasked to organize a conference on *Parallel Processing For Imaging Applications*. This time I asked who the chairs would be, so I would not again be left holding the bag. Indeed, after one year, I was able to quietly drop out. I also co-wrote my last paper [30].

References

- [1] G. Beretta, M. Gaiani, and A. Rizzi, “La riproduzione digitale del colore: una storia da quattro bit,” *Kritiké*, vol. 2, pp. 263–305, October 2021.
- [2] G. Beretta, E. Hoarau, I.-J. Lin, and J. Zeng, “Method to establish color reproduction tolerances for commercial print workflow,” HP Laboratories, Tech. Rep. HPL-2010-169, 2010.
- [3] G. B. Beretta, “Meta-palette,” Xerox PARC Video, June 12 1988. [Online]. Available: <http://youtu.be/sR8UzOQF9VQ>
- [4] —, “A new approach to imaging IC layout and schematics,” Xerox PARC, External EDL-88-3, April 1988.
- [5] —, “Selecting colors for representing vlsi layout,” Xerox PARC, External EDL-88-7, December 1988.
- [6] —, “Color palette selection tools,” in *Advanced Printing of Conference Summaries, SPSE’s 43rd Annual Conference*. Springfield: The Society for Imaging Science and Technology, May 20–25 1990, pp. 94–96.
- [7] —, “Pre-flighting variable data print jobs using the coloroid system,” in *The 10th Congress of the International Colour Association*, vol. 2, Granada (Andalusia, Spain), May 2005, pp. 1231–1234.
- [8] G. B. Beretta, V. Bhaskaran, K. Konstantinides, and B. K. Natarajan, “Encoding color images for the world wide web,” in *The 8th Congress of the International Colour Association*, vol. 2. Kyoto (Japan): AIC, May 1997, pp. 666–669.
- [9] G. B. Beretta, E. Hoarau, I.-J. Lin, and J. Zeng, “System and method for color reproduction tolerances,” Patent 9,355,339, May, 2016.
- [10] G. B. Beretta and Y. Nonaka, “House painting with NCS in the USA,” *Atti della Fondazione Giorgio Ronchi*, vol. LV, no. 1, pp. 93–118, 2000.
- [11] G. B. Beretta and M. C. Stone, “Color palette selection tools (video),” in *CHI’90 Technical Video Program—New Techniques*, ser. ACM SIGGRAPH Video Review, vol. Issue 55, Itasca, 1990. [Online]. Available: http://youtu.be/1rTacr7_H6l
- [12] D. E. Engelbart and W. R. English, “Augmenting human intellect,” Demo at the Computer Society’s Fall Joint Computer Conference in San Francisco, December 9 1968. [Online]. Available: <https://youtu.be/yJDv-zdHzMY>
- [13] A. Flachot, A. Akbarinia, H. H. Schütt, R. W. Fleming, F. A. Wichmann, and K. R. Gegenfurtner, “Deep neural models for color classification and color constancy,” *Journal of Vision*, vol. 22, no. 4, pp. 17–17, 03 2022.
- [14] N. J. Gunther and G. B. Beretta, “Benchmark for image retrieval using distributed systems over the Internet: BIRDS-I,” in *Internet Imaging II*, ser. IS&T/SPIE Symposium on Electronic Imaging Science & Technology, G. B. Beretta and R. Schettini, Eds., vol. 4311, IS&T/SPIE. San Jose (California, USA): SPIE, January 2001, pp. 252–267.
- [15] K. L. Kelly and D. B. Judd, *The ISCC–NBS Method of Designating Colors and a Dictionary of Color Names*. Washington: National Bureau of Standards, November 1955, vol. 553.
- [16] S. Kobayashi, *Color Image Scale*. Kodansha International, 1992.
- [17] M. Lamming, “Digital darkroom,” Xerox PARC Video, November 1988. [Online]. Available: <https://youtu.be/ka3QiTOJPuk>
- [18] M. McLuhan, *Understanding media: the extensions of man*, 1st ed. MIT Press, 1994.
- [19] C. Mead and L. Conway, *Introduction to VLSI systems*. Addison-Wesley Reading, MA, 1980.
- [20] I. Nonaka and H. Takeuchi, *The knowledge-creating company: how Japanese companies create the dynamics of innovation*. Oxford University Press, 1995.
- [21] K. A. Pier, “A retrospective on the Dorado, a high-performance personal computer,” Xerox Palo Alto Research Center, Tech. Rep. ISL-83-1, August 1983.
- [22] J. L. Recker, G. B. Beretta, and I.-J. Lin, “Font rendering on a GPU-based raster image processor,” in *Color Imaging XV: Displaying, Processing, Hardcopy, and Applications*, vol. 7528. IS&T/SPIE, January 2010, pp. 75 280A–75 280A–15.
- [23] D. K. Smith and R. D. Alexander, *Fumbling the future: how Xerox invented, then ignored, the first personal computer*. William Morrow & Co., Inc., 1988.
- [24] M. C. Stone, “Color printing for computer graphics,” Xerox Palo Alto Research Center, Tech. Rep. EDL-88-5, December 1988.
- [25] M. C. Stone, W. B. Cowan, and J. C. Beatty, “A description of the color-reproduction methods used for this issue of Color Research & Application,” *Color Research & Application*, vol. 11, no. Supplement 1986, pp. S83–S88, 1986.
- [26] P. H. Sullivan, *Profiting from intellectual capital: extracting value from innovation*. John Wiley & Sons, 1998.
- [27] J. Tillinghast and G. B. Beretta, “Structure and navigation for electronic publishing,” in *Color Imaging: Device-Independent Color, Color Hardcopy, and Graphic Arts III*, ser. IS&T/SPIE Symposium on Electronic Imaging Science & Technology, G. B. Beretta and R. Eschbach, Eds., vol. 3300, IS&T/SPIE. San Jose (California, USA): SPIE, January 1998, pp. 38–45.
- [28] M. M. Waldrop, *The Dream Machine: J.C.R. Licklider and the revolution that made computing personal*. Viking Penguin, 2001.
- [29] G. Wyszecki and W. Stiles, *Color Science: Concepts and Methods, Quantitative Data and Formulae*, 2nd ed. New York: John Wiley & Sons, 1982.
- [30] Y. Zhang, J. L. Recker, R. Ulichney, G. B. Beretta, I. Tastl, I.-J. Lin, and J. D. Owens, “A parallel error diffusion implementation on a gpu,” in *Parallel Processing for Imaging Applications*, ser. Proceedings of SPIE, J. D. Owens, I.-J. Lin, Y.-J. Zhang, and G. B. Beretta, Eds., vol. 7872, no. 78720K. San Francisco Airport, California: SPIE and IS&T, January 2011.

Author Biography

Giordano Beretta received a diploma in mathematics from the ETH in Zurich (1977). After working for Burroughs, he returned to the ETH and obtained a doctorate in computer science (1984). He is a Fellow of the IS&T (2001) and SPIE (2002) and is recipient of the IS&T Raymond C. Bowman Award (2009). He worked at Xerox PARC, Canon Information Systems, HP Labs, Peaxy, and Samsung Semiconductor. His work has evolved from computational geometry to color science, imaging science, and data analysis.