

Chapter 16

Don't Measure—Appreciate! NPR Seen Through the Prism of Art History

Peter Hall and Ann-Sophie Lehmann

16.1 Introduction

The question of how to judge NPR images and video is often raised, both within and without the field. The argument from outside is that NPR papers cannot be properly assessed, because there is no objective measure of quality; some even imply that NPR is somehow of lesser value for it. We do not accept the argument and reject the implication. Indeed in this book Isenberg ably argues that NPR images not only can be assessed but should be assessed. Importantly, he points out that the purpose of the NPR image (e.g. scientific visualisation) should have an impact upon the method of assessment (see Chap. 15). We focus on that class of NPR images which have the primary purpose of “being art” and consider how to judge these images based on methods used in Art History.

The problem of assessing NPR is relevant across the whole field, but is most pertinent for NPR pieces (images and video) that have been produced automatically. After all, an artist using a computer to produce art should rightly compete with artists using traditional media. Therefore it is algorithms for automatic NPR we focus upon, but to obtain a more rounded and robust view we do include interactive NPR too.

One answer, a seemingly obvious answer, to the question of evaluating NPR is “use the Turing Test”. There are many variants, but all of them ask whether a human has produced a piece or not. We argue that the Turing test is not suit-

P. Hall (✉)

Department of Computer Science, University of Bath, Bath, UK
e-mail: pmh@cs.bath.ac.uk

A.-S. Lehmann

Department of Media and Culture Studies, University of Utrecht, Utrecht,
The Netherlands
e-mail: a.s.lehmann@uu.nl

able for us because its aim is to identify the maker, not to evaluate the output: we are not interested in assessing forgery. For similar reasons we argue against experiments that claim to measure NPR value in some objective way. There is no reliable objective measure, instead NPR, like all art must be *appreciated*. This means its value derives first from the rest of NPR, and second from the culture in which NPR as whole sits. It is hardly a surprise to find a Western researcher producing NPR that resembles the Western tradition, and an Eastern researcher producing NPR that resembles the Eastern tradition [44]. The algorithms each produces can be assessed on equal terms, but evaluating the output may be more problematic because each of us inherits a cultural bias that we may not be aware of.

Consider this: suppose alien life were to make contact with humans here on Earth. It is a fair bet that creatures advanced enough to reach us would have a highly developed culture. The question is: would we recognise it as art? We would know it is not made by humans, so the Turing test is clearly of no relevance. The general problem is one of learning to appreciate art, be it from another planet, from another era or continent, or—as in our case—as made by a computer.

Given that NPR is analogous to human produced art it is relevant to look to Art History to help us *appreciate* its value. The question of why an art work is good or not, and whether art can be judged in such terms at all, has usually been easy to answer in relation to individual art works, yet difficult to clarify when it comes to developing general objective criteria. To develop such criteria has nonetheless been an on-going enterprise in the history of art, ever since Giorgio Vasari founded the discipline of Art History in the sixteenth century and Georg Wilhelm Friedrich Hegel argued in the nineteenth century that art moves towards the expression of perfection [22]. Yet since Modernism contradicted dominant teleological models that advocated an aesthetic development towards a particular standard, Art Historians and Art Philosophers have eschewed attempts to define aesthetic value in an objective sense. That is, however, no barrier to a discourse which evaluates works of art as being great, mediocre, or poor, within the realm of a particular school, style or period [30].

This chapter continues by first arguing in more detail against tests and experiment as a way to solve the problem of evaluating NPR. Next we identify norms used when assessing NPR, in particular the internal norms that might be used by a typical reviewer of an NPR paper, and we also point to external cultural references that assist us when understanding the place of NPR as a whole in the history of art. In this way we identify both an internal scale by which individual contributions can be appreciated, and we begin to calibrate that scale against wider alternatives. We conclude that at least some knowledge of Art History is required to satisfactorily appreciate NPR.

Before continuing we wish to note that NPR is a very active field, and we have been able to cite only a small fraction of the work; there is a great deal of excellent work we have been unable to refer to in this chapter.

16.2 The Unsuitability of the Turing Test and the Impossibility of Absolute Aesthetic Measure

It is not uncommon to hear the Turing test advanced as a solution to the problem of judging the aesthetic value of NPR pictures, for example [6, 73, 82], and specifically in NPR [64]. In this section we argue against both the Turing test and alternatives.

Our argument for the ineligibility of the Turing test is this: the issue at hand is the *aesthetic value* of an image, not whether a human or computer made it. In other words we are not asking “is it made by a human or machine?” but “is it justifiably a good picture?” We want a framework for assessing NPR, not a measure of the deceptive potential of a machine. This is a simple argument, but it is worth exploring a little more.

Before continuing, we should narrow our understanding of *aesthetic value* and in particular differentiate it from *aesthetic quality*. Aesthetic quality has a wide meaning, but essentially possesses a subjective character. Without reference to any agreed definition we can recognise the individual aesthetic qualities in the joy of a summer's day, the exhilaration of winning at sports, and the warmth of a lover's touch. The degree of feeling in each case is the aesthetic measure or, as we call it the “aesthetic value”, of that particular quality.

We are concerned with the visual aesthetics of pictures. In keeping with the range of aesthetic quality, there are many distinct flavours of visual aesthetic. Different cultures throughout history and across the world have produced, and continue to produce, distinctive forms of visual art. Indeed, cultures are in part delimited by their sense of aesthetic quality, as expressed in pictures. Therefore we must appreciate art, including NPR, relative to the cultural norms in which it was produced. It follows that there is more than one measure of *visual aesthetic value*. Nonetheless we can point to ‘good’, ‘bad’ and ‘indifferent’ examples within each cultural tradition, implying distinct but congruent scales. A corollary is the impossibility of assessing the aesthetic value of an NPR piece without reference to not just to other NPR pieces (so setting an internal scale) but also with reference to the wider culture in which the pieces sit (so calibrating the NPR scale against other scales).

We are not alone in rejecting the Turing Test as a valid aesthetic measure. Pease and Colton [61] give a detailed account of why the Turing test is not suitable for assessing creativity in general. Like us they point to differences in cultural heritage, but they make additional points too, including that the knowledge of the agent making a piece of art is important when assessing that piece—information that the Turing test deliberately excludes.

The question of measuring aesthetic value is much older than the Turing test it dates back to Antiquity and the notion of ideal proportions in humans, buildings, and nature, as well as the fascination with overarching principles guiding beauty, such as the number π . These discussions always concern issues of regularity and the ideal relation between the parts of a whole. For example, there is Pliny's famous account of the painter Zeuxis, who in looking for a model among the maidens of Croton for a painting of the most beautiful woman—Helena—ended up combining

the most beautiful part of five maidens, as no individual human being could embody perfect beauty alone. This anecdote would shape art theory for centuries after.

In the early twentieth century, the philosopher Birkoff defined aesthetic measure as the ratio of ‘orderliness’ over ‘complexity’ [5]. Psychologists have used this to propose a specific measure for colour harmony [55], later tested experimentally [33], and the relation between beauty and truth in scientific experiment continues to be discussed [74]. Birkoff’s measure has rarely been used in computational environments but did influence a study into the layout of graphical windows [57]. More recently, colour harmony has been studied [40] using a fuzzy logic framework. The measure has also influenced a study that concludes the aesthetic value of a scientific visualisation impacts upon understanding [26]. Whether both ‘orderliness’ and ‘complexity’ can be defined and measured in images and in video is an interesting open question. A particular issue arises because they do not seem to be independent and both relate to entropy: a system that is well ordered is not complex—it has low entropy; conversely a complex system is (almost by definition) is not well ordered and has high entropy.

We find the argument that aesthetic value has an explanation rooted in our evolutionary history to be appealing. If this is the case, then our sense of aesthetics should be embodied within image statistics. This argument has led to studies using photographs of natural scenes [9, 24], and of art [32]. Other studies show that natural scenes and paintings exhibit similar statistical properties, in particular the Fourier spectrum for photographs is about $(1/f)^{1.2}$ for natural scenes, which compares to $(1/f)^{1.4}$ for artwork [31, 75]. Sparseness of information has been found to be a useful cue [25] as is local contrast [27] which is consistent with an evolutionary explanation of beauty—although aesthetic value and beauty are not necessarily synonymous.

Whether such studies will lead to a measurable quantity that generalizes over all images is an open question, after all it has proven elusive for centuries. This is because quantifying the aesthetic value of an image is not possible based on the image alone. Rather, aesthetic value is intimately connected with the meaning of a picture. For example, knowing that Picasso’s *Guernica* is a comment on the Spanish Civil war will affect the aesthetic value one ascribes to it.

With these ideas in mind we turn away from the Turing test, Birkoff’s measure, and even image statistics as means of making progress when assessing the aesthetic value of an NPR picture. Similarly we are cautious of claims that eye-tracking data can be used to evaluate NPR images [65], not only because the attribute of aesthetic value lies beyond the picture itself but also because many pictures of high aesthetic value have no detail to focus upon (e.g., Rothko), while others are a myriad of detail (e.g., Pollock).

Instead we opt to assess the NPR in the same way that all art is assessed, which is to gauge it against the wider cultural background of existing art. In doing so we must take care, because human and computer aesthetic scales may not be coincident. Hence the problem of assessing NPR becomes similar to that faced by Art Historians when assessing any new movement in art. Our effort to establish assessment guidelines is not intended to revive the discussion as to whether computers can

generate art, but instead points to the relation NPR undoubtedly has with forms of visual, artistic expression [79].

16.3 Understanding NPR as Art

Our central claim that NPR should be assessed and understood—that is, appreciated—as art. Within the wide range of cultural image production we can consider NPR as a particular sub-genre of computer generated imagery; a point of view justified by the fact that NPR has its own recognisable norms. Currently, these norms are latent—that is they have not been explicitly articulated, but they exist nonetheless, otherwise consistent refereeing would be impossible. It is not at all obvious that these norms are uniformly agreed in all of their parts, but there is sufficient accord for NPR to be a recognisable sub-genre (hence characterisable by some set of norms). These norms not only help us delimit the field, but also to set up an internal scale. By comparing these norms with others from other fields we move towards calibrating the NPR internal scale against its wider cultural background.

16.3.1 *Internal Norms*

NPR's most defining norm is certainly the “Non-” that divides NPR from photorealism, the initial driving force in the development of computer graphics. Although the name and the inherent opposition of different kinds of realism has often been criticised, and different names have been proposed [67], it seems to best describe the tenet of the genre as being to create convincing, effective and scientifically valid images in other ways than complete visual verisimilitude, thereby employing the whole wide range of mediated visual expressions that are not bound to the photographic mechanism/technology. It has rightfully been argued that photo-realistic rendering and NPR both are equal members of the family of computer depiction [20, 23]. Yet the distinction is kept intact as images remain to be judged against the two overarching principles of photographic documentation of the visual world on the one hand and its pictorial representation on the other (notwithstanding the on-going hybridization of these principles [52]).

While photorealist rendering only has one defining medium to be judged against, NPR has many. It is this many-headed hydra that makes NPR interesting, yet difficult to assess. Its initial indistinctness therefore affords the affinity to existing, pictorial styles and hints at the immense project that NPR is: while an (ideal) image-machine that can generate photorealistic renderings is imaginable, an image-machine that cannot only generate all known pictorial styles but also new pictorial styles, may be a utopian project [66].

Art History offers an appropriate model for the assessment of a multitude of different styles because it developed criteria based on resemblance, style as well as

originality long before the invention of photography. Photography, in its early days perceived in art historical terms as the “brush of nature”, only briefly and in part served as a foil to judge art against. Mainly, by claiming visual realism as ‘good’, it stimulated painting to develop a new diversity of styles moving beyond the mimetic representation of the visual world [28, 29]: without photography it is questionable whether Cubism would ever have been, for example. NPR therefore is essentially connected to art historical principles of aesthetic evaluation.

This connection predicts the second norm of NPR: the fact that many NPR authors claim relationship between their NPR and some or other school of art: the output their algorithms produce is said to be like that of some or other school or artist. Not all authors of NPR algorithms make such a claim explicit, but even amongst those who are silent on the issue a likeness to some school remains. This provides us with an obvious link to wider culture, and invites us to make direct comparison between particular pieces on NPR and pieces produced by humans—but only in the sense of calibrating the cultural progression of NPR. (The Turing Test is appropriate only if an author claims that their algorithm is capable of forgery, most claim a likeness.) However, we should be careful to note that the styles of artists within a school offer considerable variation, and even a single artist varies over their lifetime. These matters complicate comparison between NPR and schools, and lead us to treat claims of likeness (explicit or otherwise) with caution. Even so, linking NPR to existing art is inevitable, and this link is a norm of NPR.

Closely related to the third norm of emulating pictorial style is the norm of emulating media: is it not possible to emulate oil painting without emulating oil paint. Success in this norm is much easier to assess, at least if we adopt a narrow point of view that judges the visual similarity between a synthetic and real medium. Whether NPR will ever develop its own unique medium or media is an interesting open question.

A fourth norm is the elimination of direct human input to create artistic images. Arguably, it is this norm that leads some to argue for some objective measure of success. We recognise that not all NPR shares this aim, some NPR algorithms aim to make it easier for human artists to create art, sometimes by the provision of synthetic media, other times smart tools. In this paper we are thinking almost exclusively on fully automated NPR, because it is there that the relevant issues have their sharpest profile.

A fifth norm is that algorithms should be as simple and as elegant as possible, particularly for automated NPR. This norm is inherited from the wider contexts of Computer Science and Mathematics. Additionally, it corresponds to Art Historians using knowledge about how a piece was made when coming to understand the aesthetic value of a piece. In some sense, the art of NPR is designing simple and elegant algorithms that produce output of high aesthetic value.

The final norm is novelty. In the case of human produced art this means finding a new way to express or communicate. Within NPR, novelty usually has a more restricted definition: it is the algorithm that must be novel when compared to relevant literature. This restricted view provides some evidence that NPR is not sufficiently mature for novelty to be judged more widely. We do not see this as a criticism, but as a challenge.

A sixth norm is “prettiness”, or a “wow factor”, although this is rarely stated, at least within NPR. Wider Computer Graphics, on the other hand is not quite so bashful: the idea that output should “look good” (that is, look like a photograph) is acknowledged as the driving force. NPR seems to have inherited the “look good” criterion but modified it away from using the photograph as a foil to using schools and artists. We argue that this is a norm that should not be used for NPR, for at least two reasons: first it is culturally specific, e.g. animation in Eastern Europe is very different from animation in the USA, so that, unless researchers and reviewers are cognisant of their own cultural bias, the “look good” criterion is a distorting prism. Second, and most important, even with a given tradition, art does not have to look good to qualify as great art, as we illustrate in the next subsection.

We would not be surprised if reviewers of NPR submissions make use of the above norms when arriving at a judgement regarding acceptance of a paper, indeed we would be surprised if these norms were not used. To be clear, we are not claiming these are the only issues used by reviewers—clarity of writing is not included in the above list but is a factor when assessing papers, for example. However, the norms correspond to large degree to the taxonomy of NPR that is used when assessing the aesthetic value of the NPR output—it is, after all, impossible to ignore the algorithm. These norms are useful in assessing NPR internally, that is in relation to itself. However, as mentioned already, NPR should be assessed against a wider cultural context if it is to be more roundly understood.

16.3.2 Cross Cultural Comparisons

Internal norms help us give an initial assessment of individual NPR pieces with respect to the rest of NPR. However, the progress in NPR as a whole can be assessed only by referencing it to art as a whole: if NPR is ever to claim artistic merit, then it must be judged in an equivalent way.

Art is much more than producing pretty objects. Of course it is true that art can be beautiful, but “prettiness” *per se* is no criteria at all by which to judge any work of art. We have already mentioned Guernica: the painting is impossible to appreciate without recognising that it represents the bombing of unarmed civilians—hardly a pretty subject—hardly a pretty painting. Many of Goya’s paintings are difficult to look at, unpleasant even, explicitly depicting, as they do, the horrors of war. The work of Joseph Beuys, made from fat and fur, is all the more compelling when one realizes the artist survived in fat and fur after being shot down as a fighter pilot. The photographs of Dorothea Lange depicting social deprivation in America in the 1930s are not mantle-piece objects. Yet all of these artists produce work of the highest aesthetic value; art is not imprisoned by “looking good”.

Art as practised by humans draws upon every conceivable experience, and appreciating any piece demands the viewer draws from that same well. Unfortunately the depth and breadth of the pooled, common experience is vastly wider and deeper than that of any particular individual, or even that of any particular culture. Thus,

a viewer versed in the Western tradition may find it difficult to appreciate Oriental art, for example; and vice versa. It takes effort and education and some humility to overcome these barriers.

By comparison, NPR is very limited in its scope. Most of NPR sets out to imitate art that already exists. More exactly, NPR sets out to imitate the appearance of art that already exists. Even there it is limited because the appearance is nearly always limited to output on a screen; much of the power of a Van Gogh, for example comes from paint so thickly applied it holds the passion of its maker. That passion is completely drained out when the work is presented flat, on a postcard or on a computer screen.

So far as we can tell, NPR is judged by the six norms explained above, especially on “looking good”, meaning similarity—assumed or claimed—to existing (culturally specific) genres; and by the elegance and novelty of the algorithm. None of the six norms reference the wider issues of concern to art, and because of that we cannot accept the full weight of Hertzman’s argument [38] that NPR has explanatory power with regard to art. For example, NPR does not aim to comment on social issues but art often does, and since such a commentary is essential for art appreciation NPR is the poorer for being mute. However, Hertzmann may be correct in that NPR may help to understand perceptual elements, such as why the patterns used by Van Gogh are so appealing, or how Cezanne composes his paintings. Some Art Historians argue that pictures possess a grammar [69]; NPR might be of assistance there.

When compared to the rest of art, NPR is very much found to be wanting. Works of NPR cannot currently be assessed in exactly the same way as works of art produced by humans. This is not to deny the area as a subject of study. On the contrary, acknowledging the scale and the nature of the NPR project is, we claim, the best way to drive the field. Moreover, remarkable progress has already been made, progress we now explore a little. We will find that there are three basic questions NPR algorithms must address, and that progress to date has answered one very well, partially answered the second, but barely considered the third.

16.4 The How, Where, and What of NPR

In this section we will discuss the technical issues that face NPR in general, which we name ‘how’, ‘where’, and ‘what’. These issues can be used not only to broadly classify NPR algorithms but to chart the progress of NPR against art in general.

There are three basic issues that any picture maker (human or computer) must address: how to make marks at all, where to make marks, and what to depict. *How* to make marks means deciding on media such a oil paint or pencil, as well as a particular rendering styles such as cross-hatch, impasto, palette knife, including also the choice of no perceptible mark. Having chosen how to make marks, the maker must next chose *where* to place marks on the image plane. At its most basic, marks could be placed at the edges of an object, or inside it. Support for this pair of issues comes

from [83], who separates *projective system* from the *denotational system* when discussing art. He argues that schools of art are delineated by the class of projection (from 3D to 2D) they use, rather than the marks they make. This dualism is of course directly tied to the question of whether the image is figurative and therefore refers to an object in the world, or if it is abstract and so refrains from an iconic relation to the visual world. In the following, we discuss images of the former kind, with various gradations of abstraction, but always with some relation to the visible world. A human artist may not be conscious of solving ‘how’ and ‘where’; but must solve them nonetheless. NPR algorithms are forced to explicitly solve these problems.

The third issue is *what* to depict. Humans will typically be primarily concerned with the semantic meaning of the picture; and presumably come equipped with a rich and complex internal representation capable of supporting pictures of the highest aesthetic value. This is currently beyond the scope of NPR: finding an equivalent is the grand challenge for the field.

16.4.1 How: Mark Making and Media Emulation

Any picture is an accumulation of marks: whether interactive or automatic, all approaches to NPR require an agent (human or computer) to make marks. The easiest way to make a mark using a computer is to distribute uniform colour around a point, line, or area. This approach yields marks that are flat and consistent—characteristics that are almost unique to early computer generated art. However, many people dislike such marks because it is difficult to be expressive with them, and in the early history of computer graphics they have described them as cold and unemotional (for some excellent examples see the SIGGRAPH documentary *The Story of Computer Graphics*, 1999). The response of the NPR has been to engage in research to emulate traditional artistic media.

Variations in media have long existed, with early interactive systems leading the way both in two dimensions from image [34] and in three dimensions by painting on models [36]. Media emulation has now become a staple of NPR. In fact, the affordance to simulate all artistic materials—from oil paint to charcoal, pen and ink, pastel, clay etc.—may be regarded as a distinctive material quality of NPR [49].

In fact, it was only in the early days, that Computer Graphics had a “typical” appearance because it could not yet render specific material qualities in detail (accordingly, it was often compared to plastic, which also is often like, but not quite like, a material it imitates). Emulation involved modelling the application device (brush, pencil, etc.) the pigment, be it liquid or solid, and the receiving surface. Just as photorealism models the interaction of light with matter, so media emulation models physical media and their application devices. Hairy Brushes is one early example [72], followed by many others including physical simulations of brush hairs [48]. However, the transfer of paint from brush to surface depends on more than the physics of brushes; it depends too on the bi-directional flow of physical paint [16]. Today, many of the traditional media have been emulated, not limited to

oil paint [4], watercolour [16], pencil [71], charcoal [53], pen-and-ink [18], wood-engraving [59], copper-plate [50] and mixed-media [8]. This is a non-exhaustive list as a full list would occupy several pages, but NPR has made a few interesting omissions, including tempera and fresco.

Not all mark making requires brushes, pencils or such like. Mosaics use small coloured have been studied within NPR [21]. Larger scale cut from photographs showing different views of the same object have been used to simulate Cubist-like NPR [12]. Collage is a related genre in that it requires pieces cut from many pictures, but each of a different object. Automated collage has recently been addressed within NPR; Huang et al. [41] describe a sophisticated system for cut-and-paste from Internet images to produce Arcimboldo-like pictures.

16.4.2 Where: A Salient Question for NPR

We argue that where to place marks is more important than what marks are made. We have already noted that Willats [83] correlated schools of art with a projective rather than a denotational system. It is true that the nature of marks often defines the personal style of an artist (connoisseurship works that way). However, in an NPR context the location of a mark is the most important factor in the production of aesthetic value. We can conduct a thought experiment in which the marks making up a picture are all replaced with marks of another type: would we expect the aesthetic value of the picture to change? Our answer is ‘yes’, because although marks are important, the change would not be so much as if the same marks were moved to non-salient positions, for example.

Placing marks is a very difficult problem. It is hardly controversial and no accident that the highest quality NPR comes from interactive systems, in which the responsibility of *where* to place marks is passed to the user. Many interactive systems are now very sophisticated [68]. At its most complex, deciding where to place a mark depends upon the semantic meaning of the object that mark is related to, and to all other marks in the picture. Possibly apocryphal, Cezanne is reputed to have remarked to Ambroise Vollard that he could not possibly change one stroke on the hand of his portrait, for then he needed to change all others. For automated NPR, the ‘where’ question is a significant one.

In terms of 3D models, automation began in the early 1990s [63] that used edge maps computed by differentiating depth maps. More recently locating marks often comes down to deciding which parts of a model are salient, at least from a given point of view [7, 62], including when these change in time [78]. However, the focus of our attention is image based NPR.

When working from images, early attempts at automatically choosing mark locations depended upon edge detection [51] or local variance [76]. Both of these are *low level* operations, meaning they make only weak assumptions about image content and so are applicable to many input images. The output of these and similar algorithms is often related to the Impressionist school [51]. It remains one of

the most extensively studied approaches to producing NPR. Innovations include a coarse-to-fine placement of edge-approximating strokes rather than simple blobs of paint [37] and more recently the use of filters based on structure tensors [47] or edge-aware Laplacian pyramids [60].

However, edge detection maps look very different from human sketches, even when made from the same photographic source. In fact, the difference is measurable using precision-recall plots [54]. Typically humans make many fewer marks than computers, and place marks judiciously—so that the content of the picture is efficiently depicted. In any case, it is interesting to observe that many artists, typified by Cezanne, tried to move away from painting edges at all. However, even Cezanne failed to completely remove edges. Nonetheless, edge detection and other low-level approaches to answering the ‘where’ question tend to treat all detections with equal weight, whereas artists will usually give greater weight to some edges rather than others. We call this *differential emphasis*, and reaching for it has been a driver of NPR. Interestingly, NPR rarely considers marks that become invisible by over-painting, as is the case in oil paintings by Jan van Eyck or the Photorealists.

The easiest way to pursue differential emphasis in NPR is to make use of salience maps. A salience map highlights areas that are supposed to be important to understand an image. The production of salience maps without explicit reference to any image content has been discussed in the Computer Vision literature for some time, Itti, Koch and Niebur provide a well known example [43]. Salience was recognised as useful to NPR from around 2002. DeCarlo and Santella [17] use eye-track data; they assumed that peoples’ gaze dwells on important image regions for longer than it does less important regions. They build a hierarchical description of an image, which is rendered from top to bottom to maintain details. Such an interactive approach is of limited value to automatic NPR because it provides no model of salience that can be used generally. Recently a model that predicts where humans look has appeared [45], and it has been used to emulate the results of DeCarlo and Santella [17]. However, excellent results can be achieved in other ways. Bangham et al. [3] build a hierarchical image description, based on morphological operations centred on intensity extrema, which those authors assume to be salient. Collomosse and Hall [11] defined salience via rarity, the idea being that salient image regions are uncommon in a given picture; hence they assumed that salience is a global property in that the whole image must be taken into account. The same authors used their definition of salience maps to define an objective function that guided a genetic search that lay down brush strokes in an optimal way: to “smudge out” unwanted detail while maintaining acuity where necessary [13].

Despite these efforts the problem of salience has yet to be definitely solved. A general solution will, to borrow from probability, almost for sure have to be conditioned on task: the question to ask is *is this image element salient, given this image, and given this task?* For example: is this segmented region salient, given the task is to paint a portrait of Ambroise Vollard? An answer clearly requires identification of the region in question: an eye will be painted differently from a tree in the background, and a scar on one person may identify them but be removed for someone else. Knowledge of the subject should be used, where it is available. The down side

is that using prior knowledge limits the scope of things that can be rendered. For example DiPaola [19] produces excellent portraits, but at the cost of specialising in portraiture. DiPaola writes that “... *In general, artistic methodology attempts the following: from the photograph or live sitter, the painting must ‘simplify, compose and leave out what’s irrelevant, emphasizing what’s important’.* Since human painters have knowledge of the source imagery, we are limiting this approach to portraiture and therefore take advantage of portrait and facial knowledge in the NPR process.”

It remains the case that even if salience maps improve the “look good” criterion (should be accept that measure), the outputs will still fall short of the highest aesthetic value witnessed in human art. Hence we are irrevocably drawn towards the dual subjects of abstraction and meaning, from which art derives so much of its power.

16.4.3 What: Steps Towards Abstraction and Meaning

Salience is helpful, maybe even necessary, for automated NPR, but is not sufficient. Moreover, we should not confine ourselves to thinking about brushes or pencils but instead consider all forms of picture making, such as mosaics, paper cut-outs, marquetry, etc. Such methods may require a different definition of salience to those currently in use. Projective systems too should be accounted for—in art, the rules of linear perspective are honoured more in their breach than their observance.

Artists, young and old, good or bad, in all parts of the world and throughout history are characterised by the projective systems they use [83]. Many of these defy simple mathematical modelling, in particular, Willats includes composition in his definition of projective systems, but composition is rarely researched in NPR.

The need for artistic projection in NPR has been recognised for some time now [1]. An array of non-linear camera models are now available to NPR, which allows users to create pictures with more than one focal point; sometimes with finitely many, other times with an infinite number. Most of these cameras are designed to operate on 3D models; RYAN [10] is one such example that combines projections from a finite number of linear cameras into a unified whole. The General Linear Camera (GLC) differs in that being a single camera capable of non-linear projection [84] GLCs are specified by a user defining three vectors, which deforms a plane into a bilinear surface; whereas every linear camera has a plane of points it cannot image—those points are zero homogeneous depth, called *points at infinity*. GLCs possess a bilinear surface of points at infinity. The RTcam (Rational Tensor Camera) is more general still [35], they have tri-quadratic surfaces of such points, and is designed to operate over photographic input.

Before moving on to abstraction, it is worth mentioning real-world non-linear cameras. These comprise fish-eye lens, strip cameras on satellites, and many forms of mirror. In practice no real camera is perfectly linear, but defects in lens and/or mirrors show up as artefacts such as pin-cushion distortion. Artist David Hockney suggests that many artists, such as Vermeer, made use of the equivalent of cameras in their day (e.g., camera obscurer) which were non-linear, either because of

imperfections or because of the need to re-focus on different parts of a real-world scene [39]. It is interesting to reflect on whether any such imperfections might be recovered from paintings.

Perhaps the easiest way to move toward abstraction is image segmentation. The aim of segmentation is to partition an image into semantic regions such as “face” and “tree”—a problem that not only remains open within Computer Vision but also is arguably the most difficult of all problems in that field. Helpfully for NPR, segmentation can be hierarchical, so that an eye is segmented as a part within a face. We have already seen that heuristics such as eye-tracking and image morphology can be used to build salience maps. However, Computer Vision does furnish us with a battery of alternative techniques that are beginning to match human performance.

Voronoi regions are a popular choice in NPR research, typically defined via morphology to make renderings similar to stained glass [56]. Similar segmentation techniques, coupled with some interaction, lie behind the sketches produced by [81]. Scale-space hierarchies have also been exploited by the NPR literature to abstract images, see [58] for example. Others make use of hierarchical segmentations, such as N-cuts coupled with shape classification to produce paper cuts image that shadow Matisse and others [70].

Motion may also be segmented with positive results on NPR output. The early methods of painting video would paint strokes in frame one, and then use optical flow to push those strokes across frames [51]. Unfortunately, this leads to ‘flicker’, which is caused by a combination of several effects: optical flow is not defined in the interior of regions of uniform colour; it is noisy, it is interrupted by occlusion, and so on. Segmenting the video as a spatio-temporal volume is a first step to solving this issue [15, 80] because the eigenframe of segments (which are assumed to correspond to objects) act as object-centric reference frames in which to place strokes. Furthermore, tracking objects throughout a video sequences yields a trajectory that can be used to create typical cartoon effects such as streak-lines, squash-and-stretch, and anticipation [14].

There is little work in NPR beyond segmentation; *The Painting Fool* is a rare exception, see e.g. [46]. The Painting Fool is a computer program, but is perhaps better understood as a research programme investigating artificial intelligence and creativity.¹ It is one of the few computer programs to have exhibited in real galleries.

16.4.4 NPR As Perceptually Acceptable Photorealism

At first glance *NPR as Photorealism* is an oxymoron, but closer inspection reveals this is not the case, at least photorealism as perceived is a potential output of NPR.

Consider a feature film that requires some special effect. A good example appears in *The Mummy* comedy adventure in which a sand storm—raised by the power-hungry, newly resurrected mummy—is sent to devour the heroes and heroines. The

¹<http://www.thepaintingfool.com/index.html>.

face of the mummy gapes out of the storm wall, all built of moving sand. The motion cannot be real, but the visual appearance has to convince the viewer that the storm is real. This is not normally called NPR, and if NPR related only to rendering appearance then it should not be; yet if we allow NPR to refer to motions too, then it is NPR.

The question of photo-retouching is similar. It has long been common practice for artists to ‘retouch’ photographs, especially in advertising, to add highlights to a car, to remove skin blemishes from a model, or to replace objects altogether. This practice continues, but now using a computer instead of an airbrush. Since a photograph is the source, it is undeniable that retouching moves the image away from veridical photorealism and into what might be called *perceptually acceptable photorealism*. Whether this shift is sufficient such that the output can be declared as examples of NPR is a question that depends on the tightness of the definition of NPR. Since that definition is not tight at all, it is arguable that NPR includes examples of images that look photorealistic, but which cannot be photographs.

Additionally animation and film often combine different modes of visual styles—for instance photorealism, cartoon-style and impressionist pastel style—into new hybrids of NPR, which may even constitute the most original style in NPR to date. The result can be described as functional realism, i.e. a style geared towards creating particular effects for and in the viewer [23]. Film studios and advertising houses engage with this form of NPR on an everyday basis. To a large degree it is the singular most successful branch of NPR. It is true that it relies almost exclusively on human input, but it does suggest that NPR, understood in its most general sense, has a bright future.

16.5 Conclusion

We have considered the question *how should NPR be evaluated?* More exactly, how should NPR be evaluated, when ‘art’ is its ‘task’. It is a question that arises very often, both within and without the NPR community of researchers.

We agree with [38] that NPR cannot be assessed by experiments, and in particular agree with Pease and Colton [61] that the Turing test is not a valid prescription for NPR. We disagree that NPR has an explanatory power regarding art [38], instead it seems the other way around: art informs NPR. We argue that NPR cannot be evaluated in any objective measurable way, rather it is to be appreciated by reference first to internal norms, thereby distinguishing a scale from ‘bad, to ‘good’ for comparable work; and second to external norms that give reference to a wider cultural background. The second set of norms are mutable in that they depend on culture, on the intention of the artists, and so on; it is for this reason a single objective definition of aesthetic value has evaded both historians and philosophers of art, and it is likely to evade NPR too.

When NPR is compared to a wider culture we see it focuses on technical matters, such as *how* to make a mark and (more complex) *where* to mark. Addressing the

issue of *what* to make marks about is in its infancy in NPR, but is undoubtedly the central question asked by human artists. To be clear, technique is important to humans and to art history, but only in so far as it produces art of high aesthetic value—in this case judged by cultural norms that at present are beyond the reach of NPR.

NPR research is likely to progress in the mid-term by attending to thorny issues such as object identity and function, and in the longer term by the integration of deeper cultural knowledge into its output. It has yet to find its own distinctive style; the high-water mark for NPR at this moment is—arguably—Perceptually Acceptable Photorealism because that appears so often in photographs, films etc.

Finally, and pragmatically, we argue that NPR might widen its internal norms to include terms of reference that more closely resemble those used by art historians. We suggest 8 points to consider:

1. When choosing a certain artistic school be aware of the historical background and the artists. For example, the emergence of Impressionism depended two developments: (1) Leaving the studio and the academy as a restrictive environment that had adapted a photorealistic style where the invisibility of the brush stroke was the highest ideal and (2) painting outside, trying to capture natural lighting phenomena directly with paint and without preliminary sketches, thereby developing a quicker, literally patchy manner of painting.
2. Within schools, individual artists' styles vary greatly, so that claims such as 'this paper provides an algorithm that produces art in the style of the Impressionists' needs significant qualification to have real meaning.
3. Individual artists' style varies in time (early and late style, the most famous being Picasso); again qualification is needed to be precise.
4. Materials afford certain processes and movements (brush *strokes*, pen and ink *hatching*). It could be that breaching these rules leads to non-physical media unique to NPR.
5. Media is more than physics! Materials have a distinct impact on style. Get familiar with material accordances, i.e. get the stuff and try it out in order to understand behaviour of oil paint, pastel, tempera, etc. NPR may develop new ways to apply media.
6. Art is not an accident: study, record, analyse artists' movements at work to understand salient choices.
7. Do not work from reproductions, but from the originals if at all possible. For instance, to understand why and how Claude Monet's representation of water, clouds or leaves works so well, one must view his work 'in the flesh'. Only originals allow perception of the surface structure, impasto, and texture of pictures and their materials. More than that, the originals often have a real and compelling power that can never be reproduced on a computer screen or printed on paper: maybe NPR should use media more often than it does [77].
8. Familiarity with basic principles of Art History will help when assessing NPR. Texts relate directly to the problems NPR face include Rudolf Arnheim [2], Ernst H. Gombrich [28, 29], and John Hyman [42].

NPR is in its infancy, and will no doubt flourish in the coming years. We predict it will move to become accepted as an art form in its own right. We suggest NPR should be appreciated in that way.

References

1. Agrawala, M., Zorin, D., Munzner, T.: Artistic multiprojection rendering. In: Peroche, B., Rushmeier, H.E. (eds.) *Proceedings of the Eurographics Workshop on Rendering 2000*, Brno, Czech Republic, June 2000, pp. 125–136. Springer, Berlin (2000). doi:[10.1145/647652.732127](https://doi.org/10.1145/647652.732127)
2. Arnheim, R.: *Art and Visual Perception: A Psychology of the Creative Eye*, 2nd edn. University of California Press, Berkeley (1974)
3. Bangham, J.A., Gibson, S.E., Harvey, R.: The art of scale-space. In: *British Machine Vision Conference* (2003)
4. Baxter, W., Wendt, J., Lin, M.C.: IMPaSTo: a realistic, interactive model for paint. In: Hertzmann, A., Kaplan, C. (eds.) *Proceedings of the Third International Symposium on Non-Photorealistic Animation and Rendering (NPAR 2004)*, pp. 45–56 (2004). doi:[10.1145/987657.987665](https://doi.org/10.1145/987657.987665)
5. Birkoff, G.: *Aesthetic Measure*. Harvard University Press, Harvard (1933)
6. Boden, M.: The Turing test and artistic creativity. *Kybernetes* **39**(3), 409–413 (2010)
7. Breslav, S., Szerszen, K., Markosian, L., Barla, P., Thollot, J.: Dynamic 2D patterns for shading 3D scenes. *ACM Trans. Graph.* **26**(3), 20 (2007). doi:[10.1145/1275808.1276402](https://doi.org/10.1145/1275808.1276402)
8. Brooks, S.: Mixed media painting and portraiture. *IEEE Trans. Vis. Comput. Graph.* **13**(5), 1041–10,540 (2007). doi:[10.1109/TVCG.2007.1025](https://doi.org/10.1109/TVCG.2007.1025)
9. Burton, G.J., Moorhead, I.R.: Color and spatial structure in natural scenes. *Appl. Opt.* **26**(1), 157 (1987). doi:[10.1364/AO.26.000157](https://doi.org/10.1364/AO.26.000157)
10. Coleman, P., Singh, K.R.: Rendering your animation nonlinearly projected. In: Hertzmann, A., Kaplan, C. (eds.) *Proceedings of the Third International Symposium on Non-Photorealistic Animation and Rendering (NPAR 2004)*, Annecy, pp. 129–138. ACM, New York (2004). doi:[10.1145/987657.987678](https://doi.org/10.1145/987657.987678)
11. Collomosse, J.P., Hall, P.M.: Painterly rendering using image salience. In: *EGUK '02: Proceedings of the 20th UK Conference on Eurographics*, p. 122. IEEE Comput. Soc., Los Alamitos (2002)
12. Collomosse, J.P., Hall, P.M.: Cubist style rendering from photographs. *IEEE Trans. Vis. Comput. Graph.* **4**(9), 443–453 (2003). doi:[10.1109/TVCG.2003.1260739](https://doi.org/10.1109/TVCG.2003.1260739)
13. Collomosse, J.P., Hall, P.M.: Genetic paint: a search for salient paintings. In: *Proceedings of EvoMUSART (LNCS)*. Lecture Notes in Computer Science, vol. 3449, pp. 437–447. Springer, Berlin (2005). doi:[10.1007/978-3-540-32003-6_44](https://doi.org/10.1007/978-3-540-32003-6_44)
14. Collomosse, J.P., Rowntree, D., Hall, P.M.: Video analysis for cartoon-style special effects. In: *Proceedings 14th British Machine Vision Conference (BMVC)*, vol. 2, pp. 749–758 (2003)
15. Collomosse, J.P., Rowntree, D., Hall, P.M.: Stroke surfaces: temporally coherent non-photorealistic animations from video. *IEEE Trans. Vis. Comput. Graph.* **11**(5), 540–549 (2005). doi:[10.1109/TVCG.2005.85](https://doi.org/10.1109/TVCG.2005.85)
16. Curtis, C.J., Anderson, S.E., Seims, J.E., Fleischer, K.W., Salesin, D.H.: Computer-generated watercolor. In: Whitted, T. (ed.) *Proceedings of ACM SIGGRAPH*, vol. 97, pp. 421–430 (1997). doi:[10.1145/258734.258896](https://doi.org/10.1145/258734.258896)
17. DeCarlo, D., Santella, A.: Stylization and abstraction of photographs. *ACM Trans. Graph.* **21**(3), 769–776 (2002). doi:[10.1145/566654.566650](https://doi.org/10.1145/566654.566650)
18. Deussen, O., Strothotte, T.: Computer-generated pen-and-ink illustration of trees. In: *Proceedings of ACM SIGGRAPH 2000*, New Orleans, LA, July 23–28, 2000, pp. 23–28 (2000). doi:[10.1145/344779.344792](https://doi.org/10.1145/344779.344792)

19. DiPaola, S.: Painterly rendered portraits from photographs using a knowledge-based approach. *Proc. SPIE* **6492**, 33–43 (2007). doi:[10.1117/12.706594](https://doi.org/10.1117/12.706594)
20. Durand, F.: An invitation to discuss computer depiction. In: Finkelstein, A. (ed.) *Proceedings of the Second International Symposium on Non-Photorealistic Animation and Rendering (NPAR 2002)*, pp. 111–124. ACM, New York (2002). doi:[10.1145/508530.508550](https://doi.org/10.1145/508530.508550)
21. Elber, G., Wolberg, G.: Rendering traditional mosaics. *Vis. Comput.* **19**(1), 67–78 (2003). doi:[10.1007/s00371-002-0175-x](https://doi.org/10.1007/s00371-002-0175-x)
22. Fernie, E.: *Art History and Its Methods: A Critical Anthology*. Phaidon, Oxford (2011)
23. Ferwerda, J.A.: Three varieties of realism in computer graphics. In: Rogowitz, B.E., Pappas, T.N. (eds.) *Proceedings of Human Vision and Electronic Imaging VIII*, Santa Clara, California, USA, January 21, 2003. *SPIE Proceedings Series*, vol. 5007, pp. 290–297. SPIE/IS&T, Springfield (2003). doi:[10.1117/12.473899](https://doi.org/10.1117/12.473899)
24. Field, D.: Relations between the statistics of natural images and the response profiles of cortical cells. *J. Opt. Soc. Am. A* **4**, 2379–2394 (1987)
25. Field, D.: What is the goal of sensory coding? *Neural Comput.* **6**, 559–601 (1994)
26. Filonik, D., Baur, D.: Measuring aesthetics for information visualization. In: *International Conference on Information Visualization*, pp. 579–584 (2009)
27. Frazor, R., Geisler, W.: Local luminance contrast in natural images. *Vis. Res.* **46**, 1585–1598 (2006)
28. Gombrich, E.: *Art and Illusion: A Study in the Psychology of the Pictorial Representation*. Phaidon, Oxford (1983)
29. Gombrich, E.: *The Story of Art*. Phaidon, Oxford (1995)
30. Gombrich, E.H.: The claims of excellence. In: Gombrich, E. (ed.) *Reflections on the History of Art*, pp. 179–185. Phaidon, Oxford (1987)
31. Graham, D., Field, D.: Statistical regularities of art image and natural scenes: spectra, sparseness and nonlinearities. *Spat. Vis.* **21**, 149–164 (2007)
32. Graham, D., Field, D.: Variations in intensity statistics for representational and abstract art from the Eastern and Western hemispheres. *Perception* **37**, 1341–1352 (2008)
33. Granger, G.: Aesthetic measure applied to color harmony: an experimental test. *J. Gen. Psychol.* **52**(2), 205–212 (1955)
34. Haerberli, P.E.: Paint by numbers: abstract image representations. *Comput. Graph.* **24**(4) 207–214 (1990)
35. Hall, P.M., Collomosse, J.P., Song, Y.Z., Shen, P., Li, C.: RTcams: a new perspective on non-photorealistic rendering from photographs. *IEEE Trans. Vis. Comput. Graph.* **13**(5), 966–979 (2007). doi:[10.1109/TVCG.2007.1047](https://doi.org/10.1109/TVCG.2007.1047)
36. Hanrahan, P., Haerberli, P.: Direct WYSIWYG painting and texturing on 3D shapes. *Comput. Graph.* **24**(3), 215–223 (1990). doi:[10.1145/97880.97903](https://doi.org/10.1145/97880.97903)
37. Hertzmann, A.: Painterly rendering with curved brush strokes of multiple sizes. In: Cohen, M. (ed.) *SIGGRAPH '98: Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques*, pp. 453–460. ACM/ACM SIGGRAPH, New York (1998). doi:[10.1145/280814.280951](https://doi.org/10.1145/280814.280951)
38. Hertzmann, A.: Non-photorealistic rendering and the science of art. In: Collomosse, J., McGuire, M. (eds.) *Proceedings of the Eighth International Symposium on Non-Photorealistic Animation and Rendering (NPAR 2010)*, Annecy, France, June 7–10, 2010, pp. 147–157. ACM, New York (2010). doi:[10.1145/1809939.1809957](https://doi.org/10.1145/1809939.1809957)
39. Hockney, D.: *Secret Knowledge: Rediscovering the Lost Techniques of the Old Masters*. Thames and Hudson, London (2001)
40. Hsiao, S.W., Chiu, F.Y., Hsu, H.Y.: A computer-assisted colour selection system based on aesthetic measure for colour harmony and fuzzy logic theory. *Color Res. Appl.* **33**, 411–423 (2008)
41. Huang, H., Zhang, L., Zhang, H.C.: Arcimboldo-like collage using internet images. *ACM Trans. Graph.* **30**(6), 155 (2011). doi:[10.1145/2070781.2024189](https://doi.org/10.1145/2070781.2024189)
42. Hyman, J.: *The Objective Eye: Color, Form, and Reality in the Theory of Art*. University of Chicago, Chicago (2006)

43. Itti, L., Koch, C., Niebur, E.: A model of saliency based visual attention for rapid scene analysis. *IEEE Trans. Pattern Anal. Mach. Intell.* **20**(11), 1254–1259 (1998)
44. Jones, B.: Computer imagery: imitation and representation of realities. Leonardo. Supplemental Issue 31–38 (1989)
45. Judd, T., Ehinger, K., Durand, F., Torralba, A.: Learning to predict where humans look. In: *IEEE International Conference on Computer Vision (ICCV)* (2009)
46. Kreczkowska, A., El-Hage, J., Colton, S., Clark, S.: Automated collage generation with intent. In: *International Joint Conference on Computational Creativity* (2010)
47. Kyprianidis, J.E., Kang, H.: Image and video abstraction by coherence-enhancing filtering. *Comput. Graph. Forum* **30**(2), 593–602 (2011)
48. Lee, J.: Physically-based modeling of brush painting. In: *Computer Networks and ISDN Systems*, pp. 1571–1756 (1997)
49. Lehmann, A.: Taking the lid off the Utah teapot. Towards a material analysis of computer graphics. *Z. Medien Kult.-forsch.* **1**, 157–172 (2012)
50. Leister, W.: Computer generated copper plates. *Comput. Graph. Forum* **13**(1), 69–77 (1994). doi:[10.1111/1467-8659.1310069](https://doi.org/10.1111/1467-8659.1310069)
51. Litwinowicz, P.: Processing images and video for an impressionist effect. In: Whitted, T. (ed.) *Proceedings of ACM SIGGRAPH 97*, Los Angeles, CA, August 3–8, 1997, pp. 407–414. ACM, New York (1997). doi:[10.1145/258734.258893](https://doi.org/10.1145/258734.258893)
52. Manovich, L.: Image future. *Animation* **1**(1), 25–44 (2006)
53. Markosian, L., Kowalski, M.A., Trychin, S.J., Bourdev, L.D., Goldstein, D., Hughes, J.F.: Real-time nonphotorealistic rendering. In: *Proceedings of ACM SIGGRAPH 97*, pp. 415–420 (1997). doi:[10.1145/258734.258894](https://doi.org/10.1145/258734.258894)
54. Martin, D.R., Fowlkes, C.C., Malik, J.: Learning to detect natural image boundaries using local brightness, color, and texture cues. *IEEE Trans. Pattern Anal. Mach. Intell.* **26**(5), 530–549 (2004). doi:[10.1109/TPAMI.2004.1273918](https://doi.org/10.1109/TPAMI.2004.1273918)
55. Moon, P., Spencer, D.E.: Aesthetic measure applied to color harmony. *J. Opt. Soc. Am.* **34**(4), 234–242 (1944)
56. Mould, D.: A stained glass image filter. In: *Eurographics Symposium on Rendering: 14th Eurographics Workshop on Rendering*, pp. 20–25 (2003)
57. Ngo, D.C.L., Samsudin, A., Abdullah, R.: Aesthetic measures for assessing graphic screens. *J. Inf. Sci. Eng.* **16**(1), 97–116 (2000)
58. Orzan, A., Bousseau, A., Barla, P., Thollot, J.: Structure-preserving manipulation of photographs. In: Agrawala, M., Deussen, O. (eds.) *NPAP '07: Proceedings of the 5th International Symposium on Non-photorealistic Animation and Rendering*, pp. 103–110. ACM, New York (2007). doi:[10.1145/1274871.1274888](https://doi.org/10.1145/1274871.1274888)
59. Ostromoukhov, V.: Digital facial engraving. In: *Proceedings of ACM SIGGRAPH 99*, Los Angeles, CA, August 8–13, 1999, pp. 417–424 (1999). doi:[10.1145/311535.311604](https://doi.org/10.1145/311535.311604)
60. Paris, S., Hasinoff, S.W., Kautz, J.: Local Laplacian filters: edge-aware image processing with a Laplacian pyramid. *ACM Trans. Graph.* **30**(4), 68 (2011). doi:[10.1145/2010324.1964963](https://doi.org/10.1145/2010324.1964963)
61. Pease, A., Colton, S.: On impact and evaluation in computational creativity: a discussion of the Turing test and an alternative proposal. In: *AISB Symposium on AI and Philosophy* (2011)
62. Rusinkiewicz, S., Cole, F., DeCarlo, D., Finkelstein, A.: Line drawings from 3D models. In: *ACM SIGGRAPH 2008 Classes*, vol. 39, pp. 1–356 (2008). doi:[10.1145/1401132.1401188](https://doi.org/10.1145/1401132.1401188)
63. Saito, T., Takahashi, T.: Comprehensible rendering of 3-D shapes. *Comput. Graph.* **24**(3), 197–206 (1990). doi:[10.1145/97880.97901](https://doi.org/10.1145/97880.97901)
64. Salesin, D.: Non-photorealistic animation and rendering: 7 grand challenges. In: *Keynote talk at NPAP* (2002)
65. Santella, A., DeCarlo, D.: Visual interest and NPR: an evaluation and manifesto. In: Hertzmann, A., Kaplan, C. (eds.) *Proceedings of the Third International Symposium on Non-*

- Photorealistic Animation and Rendering (NPAR 2004), Annecy, France, June 7–9, 2004, pp. 71–78. ACM, New York (2004). doi:[10.1145/987657.987669](https://doi.org/10.1145/987657.987669)
66. Schelske, A.: Zur Sozialität des nicht-fotorealistischen Renderings. Eine zu kurze, soziologische Skizze für zeitgenössische Bildmaschinen. *Image: J. Interdiscip. Image Sci.* **6**, 47–58 (2007)
 67. Schirra, J.R.J., Scholz, M.: Abstraction versus realism: not the real question. In: Strothotte, T., Deussen, O. (eds.) *Computer Visualization—Graphics, Abstraction, and Interactivity*, pp. 379–401. Springer, Berlin (1998)
 68. Schwarz, M., Isenbarg, T., Mason, K., Carpendale, S.: Modeling with rendering primitives: an interactive non-photorealistic canvas. In: *Proc. NPAR*, pp. 15–22 (2007). doi:[10.1145/1274871.1274874](https://doi.org/10.1145/1274871.1274874)
 69. Smith, P.: Pictorial grammar: Chomsky, John Willats, and the rules of representation. *Art Hist.* 562–593 (2011)
 70. Song, Y., Hall, P., Rosin, P.L., Collomosse, J.: Arty shapes. In: *Proc. Comp. Aesthetics*, pp. 65–73 (2008)
 71. Sousa, M.C., Buchanan, J.W.: Computer-generated graphite pencil rendering of 3D polygonal models. *Comput. Graph. Forum* **18**(3), 195–207 (1999). doi:[10.1111/1467-8659.00340](https://doi.org/10.1111/1467-8659.00340)
 72. Strassmann, S.H.: Hairy brushes. *Comput. Graph.* **20**(4), 225–232 (1986). doi:[10.1145/15922.15911](https://doi.org/10.1145/15922.15911)
 73. Strothotte, T., Schlechtweg, S.: *Non-Photorealistic Computer Graphics: Modeling, Rendering, and Animation*. Morgan Kaufmann, San Mateo (2002)
 74. Tauber, A.: *The Elusive Synthesis Aesthetics and Science*. Kluwer Academic, Dordrecht (1996)
 75. Tolhurst, D., Tadmor, Y., Chao, T.: The amplitude spectra of natural images. *Ophthalmic Physiol. Opt.* **12**, 229–232 (1992)
 76. Treavett, S.M.F., Chen, M.: Statistical techniques for the automatic generation of non-photorealistic images. In: *Proceedings of the 15th Eurographics UK Conference* (1997)
 77. Tresset, P., Leymarie, F.: Generative portrait sketching. In: *Proceedings of VSMM* (2005)
 78. Umenhoffer, T., Szécsi, L., Szirmay-Kalos, L.: Hatching for motion picture production. *Comput. Graph. Forum* **30**(2), 533–542 (2011)
 79. Verlaek, P.: Non-photorealistic rendering as epistemic images. In: *Workshop on Abstract Images in Art and Science* (2009)
 80. Wang, J., Xu, Y., Shum, H.Y., Cohen, M.F.: Video tooning. *ACM Trans. Graph.* **23**(3), 574–583 (2004). doi:[10.1145/1015706.1015763](https://doi.org/10.1145/1015706.1015763)
 81. Wen, F., Luan, Q., Liang, L., Xu, Y.Q., Shum, H.Y.: Color sketch generation. In: DeCarlo, D., Markosian, L. (eds.) *Proceedings of the Fourth International Symposium on Non-Photorealistic Animation and Rendering (NPAR 2006)*, Annecy, France, June 5–7, 2006, pp. 47–54. ACM, New York (2006). doi:[10.1145/1124728.1124737](https://doi.org/10.1145/1124728.1124737)
 82. Wiggins, G.: A preliminary framework for description, analysis and comparison of creative systems. *Knowl.-Based Syst.* **19**(7), 449–458 (2006)
 83. Willats, J.: *Art and Representation: New Principles in the Analysis of Pictures*. Princeton University Press, Princeton (1997)
 84. Yu, J., McMillan, L.: A framework for multiperspective rendering. In: Keller, A., Jensen, H.W. (eds.) *Rendering Techniques 2004, Proceedings of Eurographics Symposium on Rendering 2004*, pp. 61–68. Eurographics Association, Annecy (2004)