

Abstract Depiction of Human and Animal Figures

Examples from Two Centuries of Art and Craft

Neil A. Dodgson

Victoria University of Wellington

Wellington, New Zealand

neil.dodgson@vuw.ac.nz




Figure 1: Abstracted human figures. Details taken from, left to right, *Paysage*, Paul Gauguin, 1901, oil on canvas; *Dans på stranden*, Edvard Munch, 1899–1900, oil on canvas; *Étude pour “Une baignade à Asnières”*, Georges Seurat, 1883, oil on wood; *Chemin montant dans les hautes herbes*, Pierre Auguste Renoir, around 1875, oil on canvas. All images are photographs of details of the original works taken by the author.

ABSTRACT

The human figure is important in art. I discuss examples of the abstract depiction of the human figure and the challenge faced in algorithmically mimicking what human artists can achieve. The challenge lies in the human brain having enormous knowledge about the world and an ability to make fine distinctions about other humans from posture, clothing and expression. This allows a human to make assumptions about human figures from a tiny amount of data, and allows a human artist to take advantage of this when creating art. We look at examples from impressionist painting, cross-stitch, knitting, pixelated renderings in early video games, and the stylisation used by the artists of children’s books.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Expressive '18, August 17–19, 2018, Victoria, BC, Canada

© 2018 Copyright held by the owner/author(s). Publication rights licensed to the Association for Computing Machinery.

ACM ISBN 978-1-4503-5892-7/18/08...\$15.00

<https://doi.org/10.1145/3229147.3229152>

CCS CONCEPTS

• **Computing methodologies** → **Non-photorealistic rendering**; *Perception*; • **Applied computing** → *Fine arts*;

KEYWORDS

impressionism, abstraction, representation, perception

ACM Reference Format:

Neil A. Dodgson. 2018. Abstract Depiction of Human and Animal Figures: Examples from Two Centuries of Art and Craft. In *Expressive '18: The Joint Symposium on Computational Aesthetics and Sketch Based Interfaces and Modeling and Non-Photorealistic Animation and Rendering*, August 17–19, 2018, Victoria, BC, Canada. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3229147.3229152>

1 INTRODUCTION

Impressionist and other painters are able to indicate the presence of a human figure using a few strokes of a brush. This contrasts with portraiture which uses many thousands of strokes to represent a human at a fine level of detail. We look at a variety of examples of abstracted figures to elucidate the common aspects.

How can a few strokes of paint be recognised as a human figure? The key mechanism is that the human brain is able to work with

incomplete representations. Our brains are well-tuned to spotting humans. Figure 1 shows details from four paintings created around the turn of the twentieth century. As an introduction to the challenge, we look at each of these in detail to see what information a human being adds to the image to impute far more meaning that might seem justifiable from the limited number of brush strokes.

Gauguin's figure of a priest with children has only sketchy outlines of the people. He provides no details of the faces but still we are able to see that the priest has a beard and that the children are looking to the left as we see the image. Understanding that these are children requires knowledge of the relative heights of children and adults, which we all have, whereas understanding that this is a priest requires knowledge of the dress common in France at the time. Even without that particular knowledge, we can still tell that this is an adult with a hat, a beard (hence male), and a long flowing robe that comes down to his feet.

Objectively, Munch's abstracted figures are nothing more than filled outlines in yellow and white with wisps of contrasting paint attached to their tops. Subjectively, these are women dancing on the beach. We know "dancing on the beach" from the title of the work. But what tells us that these are women? Both figures appear to be wearing dresses, both have long flowing hair, and the left figure has a waist that is stereotypical for a young woman. This *could* be interpreted as a red-headed young woman dancing with a priest but we tend to make the more likely interpretation that we are seeing two women, joyfully dancing. We are even able to impute the emotion "joy" from our understanding of what sorts of emotions would be expressed by dancing on the beach and from the position of the figures: their centres of gravity are not above their feet, so they must be balancing using their joined hands, which implies that they are moving rapidly.

Seurat's figure is a few strokes of pink paint topped by a blob of brown. This all set against a blue, pink and brown background. Despite this, we identify a young man, standing to mid-thigh in water, with his right hand on his hip and his right leg thrust forward. It takes a lot of implicit knowledge about the world to formulate this interpretation from the limited visual stimulus.

Renoir creates two figures using many strokes of oil paint. The painting is less orderly than the other three examples, but we can still distinguish certain features. In contrast to the other three examples, the rear figure here has been granted some facial features, but they are indicated using just three blobs of paint that suggest eyes and a mouth. As with Gauguin's example, from the relative sizes of the figures, we assume that this is an adult and child. Both are wearing hats. The adult is carrying a red umbrella. The umbrella and the style of hat, combined with our knowledge of nineteenth century European dress, lead to the assumption that this is a woman. From our knowledge of society, we can jump to assuming that this is a female relative or nursemaid looking after a small child. These assumptions have taken us a long way from the reality, which is a disordered, meagre image that, at first glance, could be no more than random blobs of paint.

Although abstract, these four examples have human figures that are relatively large (several centimetres high). As we will see, at smaller size, an artist can produce even more extreme abstraction, yet still reliably convey posture and action with only a couple of blobs of paint or a tiny number of pixels.




Figure 2: Detail (bottom) and its context (top) taken from *Nymphéas, reflets de saule*, Claude Monet, 1916–1919, oil on canvas. The lily flower is 25 strokes of paint. Both images are photographs of details of the original work taken by the author

2 THE CHALLENGE FOR COMPUTER GRAPHICS

The challenge is whether an automated computer graphics system can possibly produce *appropriate* abstractions of human figures. We are faced with an immense amount of prior knowledge that humans use to interpret these images [Schwarcz and Schwarcz 1991]. Even without the necessary cultural understanding of conventions for clothing, most humans can interpret all these images well.

We have had level-of-detail methods in computer graphics for decades [Luebke et al. 2003]. Cartography developed abstraction over centuries [Rees 1980] that led to early active research in appropriate abstractions for computerised cartography [Jones 2013; Shea and McMaster 1989], which we see now in a polished form in apps such as Google Maps. Renderers use level-of-detail to allow them to use accurate geometry for nearby objects and approximate geometry for far objects. There are methods for automatically reducing the polygon count in objects, using geometric and perceptual formulae to determine what reduction best maintains the geometry, connectivity and look of the object [Cignoni et al. 1998].



Figure 3: Abstracted human figures. Details taken from, left to right, *Haymaking*, Camille Pissarro, 1874, oil on canvas; *Le Semeur (The Sower)*, Vincent van Gogh, 1888, oil on canvas; *La Nuit étoilée (Starry Night Over the Rhône)*, Vincent van Gogh, 1888, oil on canvas; *Argenteuil*, Claude Monet, 1875, oil on canvas. All images are photographs of details of the original works taken by the author

Artists are also able to do this. Monet, famous for his waterlily paintings [Wildenstein 1978], became so good at representing the flower that he was able to use just a few strokes of paint to produce an impression of a waterlily. Figure 2 shows one of his later painting. There are approximately 25 brush strokes in each lily, using five paint colours. It is possible to take a photograph of a waterlily and appropriate image filters to produce a wide range of abstracted effects [DeCarlo and Santella 2002; Kyprianidis et al. 2013] but this is nothing like the stroke-based approach that Monet took. It is plausible that a computer algorithm could generate similar effects from an appropriate geometric representation of the flower's structure, perhaps with a little human help in getting that structure well defined [Hertzmann 2003; Kowalski et al. 1999].

For human figures, it is not so plausible that we can approximate what Monet does with his waterlilies. So much more is going on in the human brain when processing human figures: a human artist needs to be involved to ensure that the result is plausible. What hope is there for developing a computer graphics algorithm that can achieve the level of abstraction that human artists can achieve? We will look at more examples from art, then consider two areas where we might find inspiration for what could be achieved and how it might be tackled: children's storybook illustrations and pixelated abstractions.

3 EXAMPLES OF ABSTRACTION IN ART

To provide some context for the challenge, we consider the ways in which six impressionists (and post-impressionists) abstracted human figure, look at the particular example of artist L.S. Lowry who deliberately used unrealistic representations, and look at how our ability to recognise human figures extends to familiar animals.

3.1 Impressionist and post-impressionist abstractions of human figures

I chose the four examples in Figure 3 to show a range from something that has a reasonable amount of detail to something that is barely recognisable as a human figure. The left two have sufficient detail that they could be generated from simple 3D models. Consider the images from left to right. Pissarro's figures are sketchy, but could be generated by conventional rendering followed by an appropriate image filter. Van Gogh's sower is also broadly correct, the only things missing are facial features. The two tiny figures from van Gogh's *Starry Night* are more challenging to reproduce on a computer; the hands are single blobs of paint and the women's facial features are indicated by dots. How would you ensure that the computer rendering would produce something that a human would interpret as a hand? an eye? a mouth? Looked at close up, Van Gogh's man has no distinguishable facial features at all but, from a distance, the human brain can be convinced that there are eyes, nose, mouth and beard. Monet's figure from *Argenteuil* is the most challenging of these four examples: a region of shaded green, surmounted by a pink oval with some internal features, itself surrounded by a halo of dark paint. The whole is painted against a lighter green background, again with some internal features. A first glance at the whole painting, this figure is easy to dismiss as just part of the scenery. It is similar in colouration to a large tree nearby in the background. More detailed observation leads to noticing the pink oval could be a face and the position of the figure (on a boat or on the riverbank just behind a boat) indicates that it cannot be a small tree, so we impute that this must be a human. It is even possible to convince yourself that you perceive eyes and a nose on the face but close observation shows that those flecks of paint in the face do not constitute those facial features.




Figure 4: Extremely abstracted human figures. Details taken from, left to right, *La Mosquée*, Pierre Auguste Renoir, 1881, oil on canvas; *La Voile verte*, Paul Signac, 1904, oil on canvas; *Soleil couchant sur la Seine à Lavacourt, effet d'hiver*, Claude Monet, 1880, oil on canvas; *Le Chemin de Montbuisson à Louveciennes*, Alfred Sisley, 1875, oil on canvas. All images are photographs of details of the original works taken by the author




Figure 5: Details of mid-nineteenth century watercolours. From left to right: *Mahurangi*, Caroline Abraham, 1852; *Taken from the Pier at Auckland*, Caroline Abraham, 1852; *Near Mercer, Waikato*, Albin Mercer, ca. 1850. All images are photographs of details of the original works taken by the author

Figure 4 shows what happens when you push the abstraction to the limit. My experience is that an artist needs to do considerable experimentation to know which abstractions generate the appropriate response in a human viewer. Renoir's crowd of people in *La Mosquée* are vague patches of white and blue paint with small pink circles, yet we perceive a crowd of people dressed in white robes with white head coverings. Signac's pointillist rendering of Venice includes two tiny silhouettes, whose posture shows that they are working hard to propel their boats. One of the figures has only one arm but we assume that the other is hidden rather than missing. Monet's impressionist painting of the Seine includes sketchy outlines in dark blue that we interpret as two boats with at least one person in each (is the second figure in each boat another person or a small sail?). Finally, Sisley's scene at Louveciennes, includes five vertical strokes of pink paint on a blue background. Despite any detail, a human viewer can interpret this as bathers in the river. This interpretation also cascades a range of other assumptions, including that the painting must be of a warm day.

Use of such abstraction pre-dates impressionism. Watercolourists were producing similar abstractions of human figures in the mid-nineteenth century (Figure 5). Watercolourists today are trained in similar techniques. For example, Michael King demonstrates how to draw a simple figure in three or four strokes: a stroke for a head, one for a body, a third for a shadow [King 2016].

Our ability to extract such a wealth of meaning from such detail-less abstractions is one of the issues that makes computer vision so challenging: how can you replicate a person's decades of knowledge about the world? It also makes abstraction hard for computer graphics: can we automate what these artists have done?

3.2 Non-realistic figures: L.S.Lowry

L.S. Lowry (1887–1976) was an English artist whose main oeuvre is the depiction of life in the industrialised north-west of England [Watters 1999]. His urban landscapes (e.g., Figure 6) are peopled with human figures that are popularly and disparagingly referred to as “matchstick men.” Naïvely he could be considered to be a poor



Figure 6: Detail taken from *The Rush Hour*, L.S. Lowry, 1964. Original image from aboutmanchester.co.uk.

painter, because the matchstick men are unrealistic, but this is a misunderstanding. The matchstick men are his deliberate choice.

Lowry's early work shows that he could easily handle human portraiture, perspective, light and shade but these are largely missing in his later paintings. His later painting style is a deliberately adopted impressionism. Lowry's is not the attractive impressionism of lily ponds and cathedrals at sunset. His is a grim, grey, grimy impressionism of urban Lancashire, devastated by the depredations of the industrial revolution. It is the distillation of the industrialised Lancashire townscape: dirty, sad, poor, and hard. It is an impression not a true-to-life telling. The people are stylised so that they do not draw the eye. Lowry needed people to populate his towns, otherwise the townscape would be sterile. But he also knew that the viewer's eye would be drawn inexorably to the people: it is human nature to be most interested in other humans. He says "Natural figures would have broken the spell of my vision, so I made them half unreal."

We have here a deliberate choice by a serious artist to use unreal, caricatured figures. Nevertheless, is straightforward to distinguish what *type* of person he is representing while simultaneously acknowledging that this is not a realistic depiction of a human. Despite the "half-unreal" depiction of the people, we can identify some as either woman or man, we can get a sense of mood, and we can even judge the time of day and the weather. How could an algorithm replicate this careful balance between the realistic and the unrealistic?

3.3 Animals: dogs, sheep and cows

We have seen stylised dogs in Lowry's work (Figure 6), which we interpret as dogs, rather than as foxes, wolves, cats or sheep, because their size and shape distinguishes them from cats (smaller) or sheep (fatter), and because dogs are expected to be a routine part of an urban environment, whereas wild animals are not.

Our learnt knowledge about familiar animals allows us to make these judgements and also allows artists to abstract away detail, almost to nothing, yet still produce something that is recognisable. Pissarro's cows at Éragny (Figure 7, left) have about the same level of detail as Renoir's woman and child (Figure 1, right) allowing anyone with some knowledge of cows to identify them as such. There is enough detail that an expert would be able to guess at the breed, possibly Maine-Anjou given the locale and their colouring.

Mulvey's painting (Figure 7, centre) has small blobs of white paint on a green background. My interpretation is that these are sheep grazing on a hillside, yet Mulvey's title is "Cows Grazing Grassmere." Are they sheep or cows? If your experience, like mine, is that cows are brown or black, as in Pissarro's painting, then your response will be "sheep" because sheep are the obvious white animal. If you, however, grew up in an area that farmed Charolais cattle, which are white, then you might interpret these blobs as the artist suggests, as cows. Even after being told by the artist that he intends these to be cows, I still see them as sheep.

Volpedo's painting is unquestionably of sheep (Figure 7, right). The interest here is that the figure has almost no detail. It is a fine example of light and shade, with only the faintest hints of eyes, nostrils and mouth.

4 CHILDREN'S BOOK ILLUSTRATIONS

We can go elsewhere for inspiration for what is and is not possible. One source is the illustrations in children's books [Nodelman 2005]. We look here, rather than at more conventional artwork, because a children's book illustrator will generally have to paint the same characters many times in any given book. The sequence of illustrations tells the story: "We can no more look at a single illustration in [a] book... than we can view 5 minutes of a 2-hour film... and say we have experienced the whole" [Kiefer 1995]. This need to have multiple illustrations of the same character gives us some idea of how a human achieves the different levels-of-detail required when having to render the same character at different sizes.

Having said that, consider some of the most famous illustrations, those of Beatrix Potter (1866–1943). Her *watercolours for her first book, Peter Rabbit* have Peter or his equally-sized sisters painted with heights ranging from just under 60 mm to just over 20 mm. This is a difference in scale of only three to one (the rabbits are painted in full in 21 illustrations, mean 38 mm, std. dev. 11 mm). In *Mrs Tiggywinkle* the eponymous hero is painted at a height from 40 to 75 mm, a range of less than two to one, and the little girl, Lucie, in that story is painted at a height of 40 to 55 mm, all the depictions are essentially the same size (Tiggywinkle is painted 15 times in full, mean 52 mm, std. dev. 11 mm; Lucie is painted 13 times in full, mean 51 mm, std. dev. 9 mm). These small ranges are evidence that Potter steered away from multiple levels of details: her characters are painted at roughly the same size in all instantiations. Informally, this observation generalises to many children's book illustrators.

A fine example where this is not true is Sarah Garland's *illustrations for Margaret Mahy's story, Dashing Dog* (2002). The principal characters are a family of five and their dog. Garland painted twenty-five illustrations for the book. The older girl, in a striped jacket, appears in eleven of the illustration. In most of these she is painted at about 100 mm high (eight versions, 70–138 mm, mean 102 mm, std. dev. 19 mm). In one illustration, she is 35 mm high with six red stripes on her jacket instead of the nine in the others. This demonstrates that the abstraction does not need to be accurate: we remember her having a striped jacket not how many stripes are on the jacket. The most interesting abstractions are in the book's climax, where the toddler in the family falls off a pier to be rescued by the dog. This is illustrated by four panels, rendered by Garland in the minimum of detail. In these the human figures are less than



Figure 7: Details taken from, left to right, *La Maison Delafolie à Éragny*, Camille Pissarro, 1889, oil on canvas; *Cows Grazing Grassmere (Summer)*, Jeremy Mulvey, 2002, acrylic on canvas; *Lo specchio della vita (E ciò che l'una fa, e l'altre fanno)*, Giuseppe Pellizza da Volpedo, 1895–1898, oil on canvas. All images are photographs of details of the original works taken by the author

a centimetre high and clearly demonstrate the limitations of water-colour as a rendering medium. For example, the older girl is now a few blobs of coloured paint, her jacket reduced to just two stripes. Yet, having followed these characters through six previous paintings, the reader is readily able to identify them, what they are doing, and something of what they are feeling. Remember that this book is aimed at children aged four to seven, so we are acknowledging that humans have developed the ability to understand illustrative abstraction by that age [DeLoache et al. 2003].

5 PIXELATED ABSTRACTIONS

All of the examples discussed so far create abstractions in a medium that could plausibly provide a higher resolution rendition, albeit in some cases requiring the artist to use smaller brushes to get that higher resolution. Pixelated abstractions, by contrast, occur when the rendering is constrained by the resolution of the medium, for example with the number of pixels on a computer screen.

There is a long history of creating pixelated representations of humans and animals stretching back through the history of cross-stitch to at least the start of the nineteenth century. Collectors are still able to buy cross-stitch samplers made over two hundred years ago, some of which have beautifully stylised figures on them. Figure 8 (left) shows a child's sampler from the mid-nineteenth century with clearly identifiable deer at left and a less-identifiable animal at right (perhaps a squirrel). Figure 8 (middle and right) shows modern examples, demonstrating that it is possible to render identifiable human beings and animals with a few carefully chosen stitches. The centre example is from a kit purporting to allow you to immortalise your family in thread. Careful choice of hair colour, eye colour, and favourite clothing allows creation of an artwork that represents a particular family in a way that is easily recognisable by members of that family. We also see that the family pet can be represented in an area of just 7×9 pixels and in just two colours: background and foreground. We can guess from the colour and shape that this is a dog and can, as with the cows earlier, make a guess as to its breed (perhaps a Miniature Schnauzer). To the family, however, this tiny figure would evoke a cascade of memories identifying the exact animal. At right, we see that it is possible to distinguish clearly between dog breeds despite the representations being fewer than ten stitches high.

The challenges for a computer algorithm are many. So much of the interpretation of these figures depends on human knowledge. For example, amongst the eight dogs in Figure 8, right, I recognise a Dachshund and a Dalmatian, but I cannot name the other six breeds. However, an expert identifies all eight (top row: Dachshund, Samoyed, Husky, Labrador; bottom row: Boston Terrier, Dalmatian, Pekingese, Bull Terrier). She comments that proportion is an interesting aspect here: we identify the Dachshund and Dalmatian even though their relative proportions make no sense. It seems we rely on colour and pattern rather than relative size. A further confounding factor for a computer algorithm is personal knowledge of context. The family in Figure 8 (centre) is unknown to me but, to friends and relations, they would be easily identifiable. I can, however, use my knowledge to guess that this is, from left to right, the family pet, a son about 8 years old, a mother, a father with a newborn baby in a sling, and a daughter about 6 years old. The way they are dressed, the way their hair is styled, and the fact that they are the exemplars for selling this particular type of craftwork, all conspire to allow us to impute many extra details about their lifestyle, household income, and location.

Further examples come in knitting (Figure 9) where the pixels are not square. Notice how the shape of the stitches affects the look of the figures. The sheep in the centre, in particular, appears to have a serious gash in its back as a result of the V-shape of the stitches. The stick figure is still recognisable as a human with wild hair, despite the paucity of representation.

The low-resolution pixelated displays used in the 1980s and 1990s required similar levels of abstraction. Early computer games are a prime example of using a low number of pixels to represent characters with whom the player identifies [Loguidice and Barton 2009]. For example, *Mario has evolved from a 16-pixel high character*, where each pixel was individually crafted by an artist, to a fully 3D rendered character where the artist makes the 3D model but leaves it to the renderer to decide what colour each pixel should be. Most human observers would recognise that the pixelated version is plausibly the same character as the detailed 3D model. Although the 3D model is already a stylised human being, it is a different type of abstraction from that which created the pixelated version. There are multiple challenges here for an algorithm designer: how do you take what is known about a human and create the caricature that is



Figure 8: Left: Detail of a Victorian child’s sampler, by Mary Ann Crack, showing stags that are 20 stitches high (original image from liveauctioneers.com). Middle: Detail of a modern cross stitch sampler of abstracted people (original image from marthastewart.com). Right: Detail of a modern cross stitch sampler of abstracted dogs by Kyoko Maruoka (original image from creativepoppy.com).



Figure 9: Details from two knitted jerseys. Photographs taken by the author.

Mario? Given the detailed 3D model, how could you determine the correct colour and arrangement of pixels to create the 16-pixel high abstracted version? A human artist was clearly capable of doing this, but could this process be automated?

A second example is from the 1991 video game, *Lemmings*. The “lemmings” are illustrated in the cover art for the game (Figure 10): a stylised humanoid figure with pink skin, green hair, and a blue robe. Despite never having seen such a character before playing the game, a human can impute these properties into the incredibly low resolution versions that appeared on the screen. Lemmings on screen were nine pixels high, their feet comprised two pink pixels each, their hair was six or seven green pixels. The animation loop was so well done that it gave the impression of the hair bouncing or flapping in synchronisation with the walk cycle. Players empathised with these tiny pixelated characters. How is it possible to create such engaging characters with such a paucity of representation?

6 DISCUSSION

David Salesin’s keynote at NPAR 2002 described seven grand challenges for non-photorealistic rendering. Gooch et al. followed up with a discussion on progress on each of these at NPAR 2010 [Gooch et al. 2010]. Producing good abstractions of human and animal figures falls under Salesin’s second challenge: “abstraction to capture the essence of an image.”

From the above examples we can extract several lessons. First, the need for and representation of abstraction depends on the medium. Pixels, cross-stitch, watercolour, and oils all have particular properties that mean that, for rendering figures at small size, abstraction is necessary. Further, they each limit, in their own way, what that abstraction can be.

Figure 10: Details of Lemmings cover art and actual game play (original images from retro-daze.org and classic-retro-games.com).

Second, understanding an abstraction critically depends on human knowledge. The examples from painting show that human beings are phenomenally good at extracting meaning from a tiny number of brushstrokes. The cross-stitch examples illustrate that the amount of meaning a given person can extract depends on particular prior knowledge.

The act of depiction can be thought of as an optimisation problem: the artist aims to *maximise* a desired visual response in the viewer [Durand 2002; Hertzmann 2010]. In many of our examples, the artist appears to be aiming for a *minimum* stimulus that provokes the appropriate response.

There has already been good progress on abstracting human figures, but not down to the levels of abstraction we see in some of the examples here. Methods exist to provide abstraction of photographs to pixelated art. Gerstner et al. provide a method that can take a high resolution image of a face and produce a plausible pixelated versions [Gerstner et al. 2013]. Their best examples reduce a photograph of a human face to 14×20 pixels. Compare this with the 8×8 pixels in Figures 8 and 10. The difference is that Gerstner et al. are using the image’s pixel data to get their reduction, whereas the more dramatic reduction shown here requires going further than abstracting an image of the face to instead creating a caricature that represents the essence of the person.

There are methods for vectorizing pixel art [Kopf and Lischinski 2011] and converting vectored art to pixels [Inglis et al. 2013]. These could provide insights into appropriate ways to convert 3D models to these more extreme abstractions.

In the realm of painterly art, there is excellent existing work on abstracting a human face to a sketch [Berger et al. 2013] but, as with Gerstner et al.’s work on pixelation, Berger et al.’s work only

goes so far: to an abstraction of the face but not to the outlines we see in, for example, Figures 1 (left) and 3 (right).

One way forward would be to concentrate on ways to handle this extreme abstraction as a partially manual process. Considering the human figures in our examples, with the exception of Sisley's bathers in the Seine (Figure 4, right), all have a distinguishable head and body. The simplest abstraction is therefore a roughly circular blob for the head and a longer blob for the body, as in Abraham's watercolours (Figure 5). The ability to create abstractions with this paucity of detail is a significant challenge for an automated system. A halfway house is thus to use human input, similar to that used in font hinting [Stamm 1998], where a detailed 3D or 2D model is annotated in a way that allows the extreme abstraction. Following this path, we would produce systems that assist in creating of art rather than fully automating it [Winemöller 2013]; where the artist's intent is embedded in the representation from the start [Goldstein 1999].

It appears that the real challenge for computer graphics algorithmic design is that the same human knowledge required to understand the abstractions is required to create them. For example, it is difficult to imagine any algorithm that could take an unannotated 3D model of a lemming and create the pixelated animation cycles. It may be that this is an AI-complete problem: that is, a problem that requires the synthesis of human-level intelligence [Yampolskiy 2013]. On the other hand, we have already many good algorithms for making non-photorealistic and abstracted renderings. There is therefore hope that we can invent ways to make good, effective, abstractions that play well on the understanding of the human viewer.

ACKNOWLEDGMENTS

Thanks to the anonymous referees whose advice much improved this paper. Thanks to Sarah Dillon for identifying the dog breeds in Figure 8 (right). Photographs of details of the artwork were taken in the following galleries: the Fitzwilliam Museum (Cambridge, England), Auckland Art Gallery (Auckland, New Zealand), Musée d'Orsay, Musée de l'Orangerie, Musée du Luxembourg, and Petit Palais (all Paris, France).

LINKS TO IMAGE SOURCES

The images in this paper are of *details* taken from the artwork, not of the whole work. Hyperlinks are provided in the figure captions to online records of those works that are in public collections, which include images of the complete artwork. You can also find links in the main text to the children's picture books, for which images could not be included in the paper.

REFERENCES

- Itamar Berger, Ariel Shamir, Moshe Mahler, Elizabeth Carter, and Jessica Hodgins. 2013. Style and abstraction in portrait sketching. *ACM Transactions on Graphics* 32, 4, Article 55 (July 2013), 12 pages. <https://doi.org/10.1145/2461912.2461964>
- Paolo Cignoni, Claudio Montani, and Roberto Scopigno. 1998. A comparison of mesh simplification algorithms. *Computers & Graphics* 22, 1 (1998), 37–54. [https://doi.org/10.1016/S0097-8493\(97\)00082-4](https://doi.org/10.1016/S0097-8493(97)00082-4)
- Doug DeCarlo and Anthony Santella. 2002. Stylization and abstraction of photographs. *ACM Transactions on Graphics* 21, 3 (July 2002), 769–776. <https://doi.org/10.1145/566654.566650>
- Judy S DeLoache, Sophia L Pierroutsakos, and David H Uttal. 2003. The origins of pictorial competence. *Current Directions in Psychological Science* 12, 4 (2003), 114–118. <https://doi.org/10.1111/1467-8721.01244>
- Frédo Durand. 2002. An invitation to discuss computer depiction. In *Proceedings of the 2nd International Symposium on Non-Photorealistic Animation and Rendering (NPAR)*. ACM, 111–124. <https://doi.org/10.1145/508530.508550>
- Timothy Gerstner, Doug DeCarlo, Marc Alexa, Adam Finkelstein, Yotam Gingold, and Andrew Nealen. 2013. Pixelated image abstraction with integrated user constraints. *Computers & Graphics* 37, 5 (2013), 333–347. <https://doi.org/10.1016/j.cag.2012.12.007>
- Dan Goldstein. 1999. Intentional non-photorealistic rendering. *ACM SIGGRAPH Computer Graphics* 33, 1 (1999), 62–63. <https://doi.org/10.1145/563666.563689>
- Amy A. Gooch, Jeremy Long, Li Ji, Anthony Estey, and Bruce S. Gooch. 2010. Viewing progress in non-photorealistic rendering through Heine's lens. In *Proceedings of the 8th International Symposium on Non-Photorealistic Animation and Rendering (NPAR '10)*. ACM, New York, NY, USA, 165–171. <https://doi.org/10.1145/1809939.1809959>
- Aaron Hertzmann. 2003. A survey of stroke-based rendering. *IEEE Computer Graphics and Applications* 23, 4 (July 2003), 70–81. <https://doi.org/10.1109/MCG.2003.1210867>
- Aaron Hertzmann. 2010. Non-photorealistic rendering and the science of art. In *Proceedings of the 8th International Symposium on Non-Photorealistic Animation and Rendering (NPAR)*. ACM, 147–157. <https://doi.org/10.1145/1809939.1809957>
- Tiffany C. Inglis, Daniel Vogel, and Craig S. Kaplan. 2013. Rasterizing and antialiasing vector line art in the pixel art style. In *Proceedings of the Symposium on Non-Photorealistic Animation and Rendering (NPAR '13)*. ACM, New York, NY, USA, 25–32. <http://doi.acm.org/10.1145/2486042.2486044>
- Chris B Jones. 2013. *Geographical Information Systems and Computer Cartography*. Routledge, Oxford, UK.
- Barbara Zulantz Kiefer. 1995. *The Potential of Picturebooks: From visual literacy to aesthetic understanding*. Merrill, Englewood Cliffs, NJ, USA.
- Michael King. 2016. Simple Figure Drawing in 3–4 Strokes. (2016). <https://www.youtube.com/watch?v=qPufWHc6r28> video, accessed 12 April 2018.
- Johannes Kopf and Dani Lischinski. 2011. Depixelizing pixel art. *ACM Transactions on Graphics* 30, 4, Article 99 (July 2011), 8 pages. <https://doi.org/10.1145/2010324.1964994>
- Michael A. Kowalski, Lee Markosian, J. D. Northrup, Lubomir Bourdev, Ronen Barzel, Loring S. Holden, and John F. Hughes. 1999. Art-based rendering of fur, grass, and trees. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '99)*. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 433–438. <https://doi.org/10.1145/311535.311607>
- Jan Eric Kyprianidis, John Collomosse, Tinghuai Wang, and Tobias Isenber. 2013. State of the “art”: a taxonomy of artistic stylization techniques for images and video. *IEEE Transactions on Visualization and Computer Graphics* 19, 5 (2013), 866–885. <https://doi.org/10.1109/TVCG.2012.160>
- Bill Loguidice and Matt Barton. 2009. *Vintage Games: An insider look at the history of Grand Theft Auto, Super Mario, and the most influential games of all time*. Focal Press, Burlington, MA, USA.
- David Luebke, Martin Reddy, Jonathan D. Cohen, Amitabh Varshney, Benjamin Watson, and Robert Huebner. 2003. *Level of Detail for 3D Graphics*. Morgan Kaufmann, San Francisco, CA, USA.
- Perry Nodelman. 2005. Decoding the images: How picture books work. In *Understanding children's literature*, Peter Hunt (Ed.). Routledge, Chapter 9, 138–149.
- Ronald Rees. 1980. Historical links between cartography and art. *Geographical Review* 70, 1 (1980), 61–78. <http://www.jstor.org/stable/214368>
- Joseph H Schwarcz and Chava Schwarcz. 1991. *The Picture Book Comes of Age: Looking at childhood through the art of illustration*. American Library Association, Chicago, IL, USA.
- K Stuart Shea and Robert B McMaster. 1989. Cartographic generalization in a digital environment: when and how to generalize. In *Auto-Carto 9: Proceedings of the Ninth International Symposium on Computer-Assisted Cartography*, Vol. 9. Baltimore, MD, USA, 56–67.
- Beat Stamm. 1998. Visual TrueType: A graphical method for authoring font intelligence. In *Electronic Publishing, Artistic Imaging, and Digital Typography*, Roger D. Hersch, Jacques André, and Heather Brown (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 77–92. <https://doi.org/10.1007/BFb0053264>
- Chris Waters. 1999. Representations of everyday life: LS Lowry and the landscape of memory in postwar Britain. *Representations* 65 (1999), 121–150. <http://rep.ucpress.edu/content/ucprep/65/121.full.pdf>
- Daniel Wildenstein. 1978. Monet's Giverny. In *Monet's years at Giverny: beyond impressionism*. Metropolitan Museum of Art, New York, NY, USA, 15–40.
- Holger Winnemöller. 2013. NPR in the wild. In *Image and Video-Based Artistic Stylisation*. Springer, 353–374. https://doi.org/10.1007/978-1-4471-4519-6_17
- Roman V. Yampolskiy. 2013. Turing test as a defining feature of AI-completeness. In *Artificial Intelligence, Evolutionary Computing and Metaheuristics: In the Footsteps of Alan Turing*, Xin-She Yang (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 3–17. https://doi.org/10.1007/978-3-642-29694-9_1