

OF MACHINES LEARNING TO SEE LEMON

JOHN FASS
ALISTAIR MCCLYMONT

In the seventeenth-century heyday of Dutch still life painting, exemplified by painters such as Jan Davidsz de Heem and Pieter Claesz, lemons feature heavily as opportunities for the display of matchless technique. They also fulfilled many complex symbolic functions. For the newly enriched bourgeoisie of the Dutch Golden Age, still life paintings were subtle displays of ownership and wealth derived from the Dutch colonies in Indonesia, Sri Lanka, and Taiwan, and the trade in commodities extracted from these territories to Europe. This trade was facilitated by the state-supported monopoly of the Dutch East India Company, an early example of a consolidated global corporation. Over its two-hundred-year history, the Dutch East India Company transformed from a trading company into a transglobal body with many of the characteristics of an independent state. It was able to implement its strategies through force of arms, the establishment of multiple independent markets from Amsterdam to Jakarta, and the determined exploitation of conquered lands and peoples.

Still life paintings are a product of this sociopolitical system, which accrued unprecedented levels of wealth for those with access to the opportunities provided by the company, as Toby Sonneman has pointed out in her history of the lemon.¹ The visual work performed by fruit in these images reminds the viewer that the owner can afford to purchase and consume fruit imported from far away (and, of course, that they have sufficient spare income to

engage an artist to mediate this message). Lemons, like all fruit in still life paintings, conjure impressions of decay and the transient nature of human life — but they also have a particular place in communicating acidity or bitterness, as well as the interior-exterior dynamic of hidden power exerting its influence seen in the contrast between the rough outer skin of a lemon and its glistening lustrous interior. Art historian Julie Berger Hochstrasser has emphasised the importance of pictures of lemons as especially extravagant ways of displaying affluence—“the ostentation of a whole lemon peeled and sitting at the ready, just for a little squeeze of juice”²



Still Life with a Silver Jug and Fruit, Jan Davidsz. de Heem, 1652.

Lemons were by no means unknown in Northern Europe, having been brought back from the Middle East by Crusaders in the eleventh century, and were certainly known in antiquity as recorded by Theophrastus’ *Historia Plantarum* of 300BC.³ Nevertheless, in the seventeenth century, only those with the financial means to purchase and consume them would know what they were and what they tasted like. Recognising a lemon by its shape, texture, and colour—for what it was and for its cultural and symbolic significance—was only possible for those who could see real lemons in the marketplace or identify them in visual representations.

In contrast to the social reading of how physical objects are recognised, in neuroscience and computer science object recognition in the human visual system is considered to be a function of complex brain processes that depend

on a cascade of reflexive operations. Di Carlo observes, “we effortlessly detect and classify objects from among tens of thousands of possibilities and we do so within a fraction of a second despite the tremendous variation in appearance that each object produces on our eyes.”⁴ The complexity of these neural operations are evident in the fact that “All visual cortical areas share a six-layered structure and the inputs and outputs to each visual area share characteristic patterns of connectivity”⁵ The task of producing a representation of sufficiently high fidelity to visually identify objects is poorly understood, and there is considerable debate about how it takes place in the brain. Computer scientists have concentrated on constructing computational models of perception, in order to produce explanations along the lines of a Turing machine, i.e. mathematical models that can simulate an infinite number of states. Neuroscientists, in contrast, have focused on the spatial distribution of the relevant brain activity and how these areas may be connected to each other — in other words operating at the cellular and molecular level of cortical circuitry.



An algorithm learns how to identify lemon, Photo: Alistair McClymont.

A computer vision system recognises a lemon in a very different way, although it is rooted in the search for the neural ‘algorithm’ of human visual object recognition. Computational object recognition requires ‘training’ using relevant data. The ability of a computer to correctly identify a lemon depends on the number, quality, and accuracy of examples in the data class ‘lemon’ the system has been exposed to. In addition, the human visual system is very good at processing interruptions to the visual field (we have no difficulty in, say, recognising a tennis racket that is resting on a chair in front of a window that looks onto a river). In the scenario above a computer would need to identify every object, perceive how they are arranged in space, assign the correct label to each one, and recognise the whole as a scene. As Luc van Gool points out, “The same object will look different depending on the viewpoint, the illumination, or the occlusions caused by other objects in front.”⁶ Furthermore, the wide variation between instances of the same object means that the “recognition of an object as belonging to a particular group is a harder problem for a computer than the recognition of a specific object.”⁷ One lemon does not look exactly like another, and therefore significant computational resources must be devoted to distinguish between different views, types, or examples of lemons as an object class. As a result, much attention is given to pre-categorisation of images of objects via tagging or other taxonomic labelling methods. This requires large amounts of individual images but, perhaps more significantly, a dominant logic of categorisation to work.

As more and more detailed models are developed based on the millions of images used as training data for object recognition learning algorithms, so system complexity increases. Antonio Torralba, associate professor of computer science and engineering at MIT states, “Deep learning works very well, but it’s very hard to understand why it works—what is the internal representation that the network is building.”⁸ The outcome is an opaque system, resistant to analysis, impervious to scrutiny. Often the data scientists behind this work have no idea why they obtain certain results, and have to commit resources to reverse-engineering them in order to gain a deeper understanding.⁹ This opacity of computational recognition systems means the steps involved in recognising a lemon are transformed from a set of specific associations, enculturated by human circumstance and experiential phenomena, into the output of an impenetrable probabilistic matching algorithm. Image recognition systems are of course, the result of a set of cultural assumptions about efficiency, accuracy, and performance, usually enacted through precisely defined operations in institutional environments.

The resulting set of representations, whether cellular or computational, remain mysterious in origin, prone to error, ambiguous in value, of erratic reliability and doubtful authenticity.

- 1 Toby Sonneman (2012) *Lemon: A Global History*, London: Reaktion Books.
- 2 Julie Berger Hochstrasser, (2007) *Still Life and Trade in the Dutch Golden Age*, New Haven, CT: Yale University Press. Pp 200.
- 3 Benedict Einarson (1976) 'The manuscripts of Theophrastus' *Historia plantarum*', *Classical Philology*, 7(1), pp. 67-76.
- 4 James J. DiCarlo, David E Zoccolan, Nicole C Rust (2012) 'How Does the Brain Solve Visual Object Recognition?', *Neuron*, 73(3), pp. 415-434 .
- 5 Daniel J. Felleman, David C. Van Essen (1991) 'Distributed hierarchical processing in the primate cerebral cortex', *Cerebral Cortex*, 1(1), pp. 1-47.
- 6 Quoted in ICT Results (2009) *Teaching Computers To Recognize Objects*, Available at: <https://www.sciencedaily.com/releases/2009/06/090601090029.htm> (Accessed: 19th August 2018).
- 7 *ibid.*
- 8 Quoted in Larry Hardesty (2015) *How computers can learn to recognise objects*, Available at: <https://www.weforum.org/agenda/2015/05/how-computers-can-learn-to-recognise-objects> (Accessed: 19th August 2018).
- 9 *ibid.*