**Robert Hopkins – Workshop Paper**
**Factive Pictorial Experience:**
**Photography before and during the digital revolution**

*Note: readers more interested in photography than philosophy may choose to skip the sections Photography and perception; Other information-preserving systems, and Walton's view (pp.11-16).*

## Some questions and an answer

Here are some common thoughts about photography. Traditional photography, using photosensitive film and the accompanying developing techniques, is a way of making pictures that are special in two ways. First, photographs have an epistemic status that ordinary, 'handmade', pictures such as drawings, paintings, etchings and watercolours, do not. Photographs offer us a way of finding out about the world that is more secure than that offered by handmade pictures—when these last play that role at all. Hence photographs can act as evidence, in newspapers or courts of law, as other pictures cannot. This thought is sometimes elaborated by analogy with two other sources of knowledge. While other pictures make knowledge available in something like the way the (verbal) testimony of a reliable informant does, photographs yield knowledge somewhat as perception does. Second, this difference in epistemology is accompanied by one in phenomenology: we experience photographs differently from other pictures. They seem to put us in a relation to their objects that is somehow more intimate, more direct, than that in which we stand to the objects handmade pictures depict (Walton 1984). Hence if we find that what we took for a photograph is in fact a hyperrealistic painting, or vice versa, our experience shifts. But, it has also been suggested, both these features are under threat from the replacement of traditional methods of making photographs with digital ones. Digital photographs are closer to handmade pictures in certain key ways that prevent them being special as described. The ease with which they are manipulated aligns them, epistemically and experientially, with non-photographic pictures. It is sometimes suggested that their arrival will even undermine our attitudes to photographs made the traditional way (Savedoff 2000: 201). These too we will cease to treat as epistemically privileged, or to experience differently from handmade pictures.

Is any of this right? What exactly does the specialness of traditional photography consist in? I will defend the following view. Like other pictures, traditional photographs support pictorial experience—we see things in them. But unlike other pictures, that experience is factive: it is guaranteed to reflect the facts. What we see in traditional photographs is, of necessity, true to the facts, and in particular to how things were when the photograph was taken. At least, this is the experience traditional photography is designed to produce and which it does indeed produce, when everything works as it should. It is this that explains traditional photography's special epistemic status, and the special experience it instils. Whether digital photography also offers those benefits depends on whether it too forms a system of picture making that, when everything works as it should, supports factive pictorial experience.

## Pictorial experience and factivity

At the centre of this account lies the idea of factive pictorial experience. To understand that idea, we must grasp what pictorial experience is, and what factivity amounts to. Let's start

with the former. Pictorial experience, or seeing-in, as I will also call it, is a distinctive visual experience that representational pictures—*all* representational pictures—sustain. It is in having this experience that we grasp what a picture depicts. One way to familiarise yourself with the idea is to consider the situation in which you see a picture but can make no sense of it. It is clearly a surface marked in various ways. Perhaps you suspect that it is a picture of something. But you don't yet see what it depicts. Then, as sometimes happens, the penny drops. You continue to see a marked surface as before you, but now the marks make sense visually. You see them as organized in a particular way, and thereby as (in some sense) presenting you with an object or scene—a dog in a park, perhaps; a river iced over in winter; or a simple geometrical object, such as a cube. It is this second experience of the picture that is seeing-in. Every representational picture supports an experience of this kind, whether or not one comes to it by first experiencing the picture without seeing anything in it. And only pictures support it. We don't have that experience when we see objects face-to-face, since seeing-in involves, as seeing face-to-face does not, awareness that what is before one is marks on a surface. And we don't see things in non-pictorial symbols, such as words. Recognizing a (written) word does indeed involve seeing the marks that compose it as organized in a certain way. But they are not organized in such a way as to present you with whatever the word refers to. You see the marks as, say, forming the word 'dog', but not as presenting you with *a dog*.

So much for pictorial experience: what of factivity? That is a feature of some of our mental states. The factive mental states are those that necessarily reflect the facts. More precisely, a kind of mental state is factive if the following is true: for any particular state of that kind, if the state represents it as the case that *p*, then it is indeed the case that *p*. Factive states, then, capture how things are. Indeed, they cannot fail to do that—they cannot represent things to be other than how they really are. This may seem puzzling. How can any mental state be immune to error in this way? Some examples may help. Perceptual states, such as seeing and hearing, are factive. If I perceive that the bus is full, then the bus is full. I can make a mistake, of course. It might look to me as if the bus is full when in fact it's half empty (perhaps all the passengers are sitting on the side nearest me). But that just shows that I didn't really *perceive* that the bus is full—what I took to be a perception wasn't really one. Knowledge provides another example. If you know the atomic number of gold to be 42, then the atomic number of gold is 42. Again, you might be mistaken, about this as about any other matter. But, if you are, that just shows that your state of mind doesn't count as *knowledge*. As these examples suggest, factivity is often less a matter of our mental workings being magically free from error, and more a matter of how we choose to describe those workings. Sometimes it suits us to group mental states by their success. To call an experience a perception, or to describe a belief as knowledge, is to identify it as a state that gets things right.

The idea I want to explore is that some pictorial experiences are also factive. In these cases, what is seen in the picture is indeed how things really are. However, if we're to apply the above definition of factivity to seeing-in, we need to clarify two matters. The first is that it makes sense to talk of seeing facts in pictures. Of course, it sounds odd to talk of seeing in the picture that *p* (for instance: she saw in the picture that the river was iced over). But the obstacle here is only one of idiom. For one thing, it is natural to talk of seeing in pictures such things as *an iced-over river*, or *the suffering of the Greeks*. If so, what is seen in the surface already contains the materials for a fact: an object (the river) and a property it posseses (being iced-over). For another, the unnaturalness associated with 'seeing-in' is lessened when we

turn to 'pictorial experience'. If we can *see* (face-to-face) that a river is frozen, and thus in a perfectly good sense *experience* that the river is so, why should we not *pictorially experience* that the river is frozen over, when confronted with a picture of the scene? We should allow, then, that pictorial experience can indeed be of facts, that facts can be seen in pictures.[1] The other matter is that, if pictorial experience is factive, then, usually at least, the facts it reflects will be in the past. They were the facts that obtained when the picture was made. To allow for this, we need to make the definition of factivity more flexible: a factive state is one that necessarily reflects how the facts *are or were*.

These clarifications in hand, we can apply the notion of factivity to pictorial experience. There is a kind of pictorial experience for which this is true: for any particular experience of that kind, if what is seen in the picture is that *p*, then it is or was the case that *p*. The main thought I wish to offer is that this is true of our experience of photographs. Traditional photography, when it works as it should, sustains factive pictorial experience. Handmade pictures, in contrast, do not. This is what is epistemically special about traditional photographs, and the source of our distinctive experience of them. A proper explanation of the special features of photography will, however, have to wait. First I must try to persuade you that the claim I want to use to do this work, that photographs support factive pictorial experience, is true.

## *Does* photography support factive pictorial experience?

### *(i) Preserving information*

How would things have to be, for our experience of photographs to be factive? That experience is the last stage in a long chain of causes and effects. The chain begins with light from the photographed object falling on the film; involves various stages of development and, for print photographs at least, printing; then light passing from the resulting print to the visual system of the viewer; and ends with visual processing in the viewer's brain bringing about pictorial experience of the photograph. If that experience is to be factive, it seems that every stage in this chain must itself exhibit something like factivity. Every stage must in some sense preserve information available at the previous stage, so that, if the later stage contains the information that *p*, so did the earlier stage.[2] For if information can be introduced at any stage after the first, that information will not reflect the facts that obtained when the picture was taken. And then, if the rest of the system goes on to preserve that rogue information, the end product—seeing-in—might wrongly present that as how things were in the photographed scene.

Any causal system that preserves information in this way can allow only a restricted role for human agency. In particular, it cannot at any stage give an essential role in shaping the succeeding stages in the chain to how some person takes things to be. If how things are at the umpteenth stage depends on how I or anyone else thinks they were at an earlier stage, then we lose our guarantee that the only information in the system originates from the initial input. The reason is simple: people make mistakes. They can mistake one thing for another,

[1] There are other reasons for resisting this than mere linguistic awkwardness. However, those reasons count just as heavily against the idea that we can *perceive* facts. The factivity of pictorial experience is thus no more problematic than the factivity of perception.

[2] This doesn't count as factivity proper for two reasons. First, factivity is a feature of mental states. Second, factivity requires that its bearer *represent* things as being a certain way. There can be information where there is no representation, and states that bear information need not be mental.

and even fabricate things that were not there at all. If a system relies on people's experiences, beliefs or judgements for preserving information, it cannot churn out factive mental states as its output.

This is what prevents handmade pictures from supporting factive seeing-in. Of course, handmade pictures can reflect the facts. An illustration in a maintenance manual, a portrait in oils, or a reporter's sketch of the courtroom can all accurately convey how things are in the scenes they portray. When they do, what we see in these pictures will be how things really are (or were). But factivity requires more than capturing the facts, it requires doing so as a matter of necessity. Handmade pictures never do *that*, and the reason not is that, when accurate, they are so because they reflect how someone—whoever made the picture—took things to be. As a result, handmade pictures are always vulnerable to error, whether or not they succumb to it. If the reporter is red-green colour blind, or mistakes a woman for a man, or hallucinates an extra stenographer in the courtroom, the resulting picture will not reflect how things were that day in court.[3] That is the situation with all handmade pictures: even when they aim at accuracy, how they show things to be depends on how their artists take them to be. As a result, the seeing-in they support is not factive.

Traditional photography, in contrast, involves a causal chain that is free from the influence of people's beliefs and experiences in this way. Of course, the photographic process involves people at various points. The most crucial are the beginning and end, in the form of photographers and the viewers of the finished photograph; but often people are also involved in the intervening stages of developing and printing. Generally the causal chain is sustained at key points by human intervention, and photographers in particular determine where the chain starts: what scene is photographed, at what moment, from what angle, and in what conditions. What folk do not do, however, is to play an essential role in the preservation of information. Once the shutter has been opened and the film exposed, the passage of information from one stage to another is essentially independent of how anyone thinks things were in the scene photographed. Indeed, at each stage the passage of information is independent of how anyone takes *any* earlier stage in the chain to be. Choices in developing may determine *which* information is preserved: for instance, if we choose to develop the image so as to reveal detail in areas of deep shadow, sacrificing that of areas in full sunlight. But what those choices do not do is to determine whether the information is there to be preserved in the first place. There is no equivalent here of mistaking someone's gender, thereby producing a picture that shows a man where, in reality, a woman stood.

Many have noted this point of contrast between traditional photographs and handmade pictures (Walton 1984; ??). What has not been noticed is that this opens the way to the claim that photographs support factive pictorial experience. The explanation may be that the claim appears hopeless. Surely it is simply not true that our experience of photographs is factive. Looking at a photograph, I might see a royal rubbing shoulders with ordinary folk even though the picture is in fact of some lookalike. If the photograph was overexposed, I will see in it objects with colours less saturated than those really displayed. If the camera has

[3] The point holds even if the artist plays a more localised role in the causal chain. Consider, for instance, a picture made by copying a photograph by hand. Even if the photo captures the facts, the copy may fail to do so—if the copyist mistakes what the photo shows.

developed shutter shake, the scene visible in the photo will be blurred, and perhaps seem to be in movement, as the real scene was not. If the film has been double exposed, I will see in the picture elements of two different scenes somehow occupying the same space. Perhaps someone has been airbrushed out of the Politburo's appearance on the balcony, or little fairies have been added by hand to the image of the garden. In all these cases, the photographic 'system' is not information-preserving in the way described above. In all, the result is that what is seen in the picture is not how things were. If our experience of photographs can be erroneous in these ways, how can it be factive? Independence from how people take things to be may be necessary for a system to support factive experience, but is hardly, it seems, sufficient.

*(ii) Design and proper working*

We should not give up so easily. It is true that some of our experiences of photographs are erroneous, and hence the experiences we actually have do not form a factive kind. Nonetheless, traditional photography is *designed* to produce factive seeing-in. Moreover, that is precisely what we get, when everything works as it should. Provided we can legitimately help ourselves to these notions, we can save the claim that photography supports factive experience. It does so in the sense that that is what traditional photography aims at and, when it works, delivers.

Is appeal to these notions legitimate? The main reason for doubt lies in the danger of trivialising our claims. We're interested in factive pictorial experience in order to develop contrasts between photography and handmade pictures. But suppose we allow ourselves these new resources in the latter case. Much of our experience of handmade pictures is not factive. But concentrate on pictures *intended to capture the facts*, and limit attention to those cases in which *everything works as it should*, and there too we have a kind of pictorial experience that is factive. We've simply written their accuracy into our way of identifying the relevant experiences. If factivity comes this cheaply, we might wonder whether it's worth having.

To see if this danger can be avoided, let us proceed more cautiously. Factivity is a feature, not of individual mental states, but of the kinds to which those states belong. Individual states can get things right or wrong. The idea of factivity is that some kinds of state are such that *every* individual state belonging to that kind accurately reflects the facts. The claim that a given kind of state is factive will be trivial if we have no way to identify the members of that kind except as states that get things right. And the claim will be false if we pick out a group of experiences not all of which are accurate. Thus it would obviously be trivial to identify as factive those pictorial experiences <u>that get the facts right</u>. We do no better if we identify them as those <u>successfully intended</u> to do that. For the idea that the intention is successful just *is* the idea that the experiences in question are those that get the facts right. But nor can we simply drop talk of success and try to identify the experiences as those <u>intended (successfully or not) to get the facts right</u>. That avoids triviality only by leaving open the possibility that the experiences don't, in fact, fulfil that aim. It admits experiences that *do not* get things right, and so identifies a kind that is not factive after all. To avoid the twin perils of triviality and falsity we need a way to pick out the relevant kind of experience other than by accuracy, but which nonetheless leaves us only with experiences that reflect the facts. Is there such a way, for traditional photography?

There is, if we adopt a two-stage approach. The first stage is to acknowledge that traditional photography is indeed designed to sustain accurate seeing-in. That goal has guided the design and manufacture of photographic and development equipment, and our practices of taking and viewing photographs. Of course, if we are to avoid triviality, we cannot identify the factive pictorial experiences directly in terms of that goal. But the second stage is to recognize that, as a consequence of that overarching goal, each element in the complex causal chain in traditional photography is governed by a notion of *things working properly*. For instance, semi-automatic cameras are set up in such a way that, given the prevailing light conditions and the photographer's choice of aperture, the shutter speed is to be regulated so as to let in the right amount of light to form a differentiated pattern on the film. Too little light, and the pattern goes missing through underexposure. Too much, and it suffers the opposite fate. If an appropriate amount of light does not fall on the film, the mechanism has malfunctioned. This principle of the mechanism's functioning can be stated independently of the idea that photography aims at accurate seeing-in. It gains its point, however, from that overarching aim. Since nothing can be seen in film that is so underexposed as to yield an entirely dark image, or so overexposed as to yield an entirely light one, the principle of this part of the camera's working makes perfect sense, in light of that overarching goal.

The idea is that similar principles of proper functioning govern every element in the causal chain that constitutes traditional photography. At each stage, there are norms defining what it is for that element to be working properly. Those norms gain their point from the overarching goal of the design (to sustain accurate seeing-in), but they can be stated without reference to that goal. The proposal is to identify the factive pictorial experiences in these terms: they are those the photographic system produces when, in terms of the norms of proper functioning, every element works as it should. The claim that the experiences so identified are factive is not trivial. Moreover, I submit, it is true.

*(iii) Examples of norms*

To state any of these norms precisely would require mastering the details of how photography's various tools and methods have been designed. Fortunately, we only need to be persuaded that such norms exist and that they govern every aspect of the photographic process. Here, then, in outline, are some further examples. Consider the norms governing the development process. These prescribe that the developer—a person or an automated machine—should expose the film to the developing chemicals for long enough to best bring out the differentiated pattern formed on the film. There may, of course, be room for individual choice here. As we noted, given the varying light levels within the scene and the corresponding degree to which parts of the film have been exposed, the developer (at least when it's a person) may choose to exhibit detail in some part of the pattern at the price of doing so in another. That shows that sometimes the norms governing a process admit two different outcomes as equally acceptable. As before, the norms here gain their point from the overarching goal behind the system's design (accurate seeing-in); but they can be stated (insofar as they can be formulated at all) without mention of that goal. Other norms govern the workings of the lenses and focussing mechanisms, and concern the resolution of the image on the film itself. We can define proper working here in terms of purely optical notions, such as the convergence of light rays coming from the object at which the camera is pointed. Again, then, while the point of these norms is in the end to secure accurate seeing-in, they can be formulated independently.

As our discussion of developing shows, some of the relevant norms govern the conduct of people. For instance, the photographer who works with an entirely manual camera should aim for the same balance of light as the semi-automatic camera's light-governing mechnanism, and to focus the lens so as to resolve, in the way just described, the image cast by the object on the film. (Anyone doubting that there really are norms here should note the language in which development labs describe failures to do as one should: 'too little light', 'poor focus' and so forth.) Of course, any photographer is free to ignore these norms. The claim is not that photographs are never taken deliberately out of of focus or with deviant exposures. Quite the reverse: a skilled photographer may exploit these deviations to great effect. The claim is only that there are norms that are here ignored. It is our experiences of photographs that conform to those norms that will be factive. (Parallel points hold for the developing process.)

Perhaps the most interesting norms govern our experience of photographs. It might seem that here at least there really is no room for the normative. One sees what one sees in a picture; can anything more be said? In what sense can one's experience of a picture be wrong or malfunctioning? Certainly viewers generally have less control over their experiences than do, say, developers over their actions in the dark room. But for all that, norms make perfect sense in this context. Indeed, anyone who finds the idea of seeing-in attractive is already committed to this idea. For, as others have noted (Wollheim 1987), it is always possible to see in pictures objects other than those they depict. One can see hunched figures in the silhouette depiction of a bat, a smooth fruit in a photograph of a dolphin's head, and faces in stained walls that do not depict anything at all. If seeing-in has anything to do with grasping what pictures represent, it must therefore admit of right and wrong.

Thus not all pictorial experience is right or appropriate. In the case of traditional photographs, appropriate seeing-in is a matter of experiencing the surface in a way that reflects the causal processes by which it has been generated. We need to be careful here, however. Does deference to the causal processes mean that the relevant norm is that we should see in the photograph the facts actually photographed? If so, triviality would, after all, cripple our attempt to identify the factive experiences by reference (*inter alia*) to this norm. It is as well, then, that that is not what such deference amounts to. Suppose the image in question is mildly overexposed, so that the colours are washed out, or that the camera has shutter shake and the image is blurred. Should we see the original colours and crisp contours in these photos, or less saturated hues and blurred outlines? Of course, if we want to reconstruct how things really were, we should conclude that the colours and lines were normal, and that the photograph fails to capture them. But our topic here is not the conclusions we should draw, but the pictorial experiences we should have. There is at the very least *a* sense in which what we ought to see in these pictures is the distorted colours and outlines the malfunction has put there. After all, only thus can we reconstruct the full facts—that a normal scene has been captured poorly. But if this is what we should see, the norm governing seeing-in is not one that would trivialise our position. Proper functioning at this stage is determined not by the facts photographed, but by the rest of the causal process. Failures further upstream are not corrected for, but accommodated. The norms governing seeing-in itself, like those governing other aspects of the system, while taking their point from the overarching goal of accurate pictorial experience, can be stated independently of it.

*(iv) The contrast gained*

I hope these examples make plausible that every stage in the photographic process is governed by norms that fit the template described. These norms in hand, we can define the kind of pictorial experience that is our concern as that generated by traditional photography when every element in the system works as it should. The suggestion is that the kind so defined is factive. Although I have not offered conclusive grounds for this claim, what I have said is at least enough to handle the threats to it above. In all the supposed counter-examples, one or more of the norms has been broken. This can be deliberate, as (perhaps) with double exposure or airbrushing out the comrade who has fallen from favour. It can be accidental, as with overexposure through incompetence and blurring through shutter shake. It might be a matter of seeing the wrong thing in the image, as with the royal and her lookalike. The general point is that every one of the cases seeming to show that our experience of photographs is not factive involves some infringement of some norm of proper functioning. If we can successfully accommodate all these examples in this way, the onus is on anyone sceptical of my claim that photographs support factive seeing-in to explain why we should doubt it.

Before moving on, let's briefly return to handmade pictures. My goal was to show that photographs support factivity as handmade images do not. Does my way to save the first part of that contrast also preserve the second? It does. Suppose we attempt to make sense of the idea that some handmade pictures support factive seeing-in. If we're to avoid trivialisation, we need to identify norms governing the process of making these pictures that fit the template above: they do not need specifying by reference to the idea of accurately depicting the facts, but nonetheless secure that the resulting picture allows us to see the facts in it. But the the artist's mental states play a key role in the process of making such pictures: information from the object portrayed can only reach the viewer via the artist's take on the world. How can we describe the norms governing this process so that they guarantee accurate output? We can do so only, I think, by invoking norms that explicity require the artist to grasp <u>how things really are</u>. For what other norms could do the job? In general, if an artist chooses to caricature her subject, or simply to draw some scene that is a product of her imagination, in what sense has the picture-making system (supposing there to be such a thing) gone wrong? Of course, there are contexts in which this idea makes sense. The artist reporting from court is in such a context. But here the norms precisely take the form we're trying to avoid. They are, for instance, that he should draw things <u>as he sees them to be</u>. The facts are invoked in specifying the norms, thereby rendering triviality inevitable. No interesting notion of factive pictorial experience can emerge.

## The epistemology and phenomenology of traditional photography

So, I claim, traditional photography, when it works as it should, sustains factive seeing-in. How does this help clarify what is special about photography, both as a source of cognitive engagement with the world and in terms of the way we experience it?

When photography works as it should, our pictorial experience is factive. This yields two immediate benefits, epistemologically speaking. First, factive experience, of any kind, is itself an epistemically valuable state. It is experience that is guaranteed to show one how the world is or was. It offers one secure engagement with the facts. If that's not an epistemic benefit, what is? Of course, epistemology is also, perhaps centrally, about knowledge. It would be nice

to be able to say something about how photographs yield that. The account allows us to do this, thus revealing the second benefit photography offers. Knowledge, it is traditionally thought, is true belief plus some further factor. The further factor is needed because beliefs may be true and yet fail to be connected to the facts in the way that knowledge requires. If I'm to know that *p*, then the fact that I believe that *p*, and the fact that that belief is true, must not be related by mere accident. There is a good deal of disagreement in epistemology about what the missing factor might be. Without becoming embroiled in those debates, we can see how photography provides the extra that true belief requires to qualify as knowledge. When photography works as it should, our pictorial experience is factive. Factive experience is guaranteed to get the facts right. Assuming I go on to form a belief that reflects what I see in the photograph, that belief will not be true by accident. It is true because it reflects the content of an experience that is itself guaranteed to get the facts right.[4]

Handmade pictures offer neither of these benefits. Since experience of them is not factive, it is not itself epistemologically valuable as experience of photographs is. And nor is it the case that any belief formed about the objects seen in a handmade picture can secure its claim to be knowledge by the means above. This is not to say that handmade pictures cannot offer us knowledge—they certainly can. The point is that, when they do, the beliefs they lead us to form do not count as knowledge in virtue of being formed in response to factive seeing-in.

So far so good. However, one may feel that an important epistemological issue has yet to be addressed. We have explained how beliefs formed on looking at photographs might count as knowledge. What we have not yet done is to explain how the viewer comes to form those beliefs in the first place. It is natural to assume that the reasoning he does or might (if he were to put his mind to it) undertake makes a difference to whether the resulting belief is rational. After all, in general we do allow photographs to guide our beliefs about how things are (or were), and do not allow handmade pictures to do so. For instance, we take a photograph in a newspaper to show the nature of the damage really done by a bomb, as a drawing in a war comic does not. So far, however, we have concentrated on a difference between the two cases—whether pictorial experience is factive—that is not, for all we've said, a difference of which the viewer is aware. What is the difference between the two cases *from his point of view*?

The short answer is that the differences for the viewer are more or less the differences already noted. Of course, many viewers will not have the concept of factive seeing-in. They would certainly be hard pushed to describe the norms governing each element in the photographic system. But viewers of photographs do, I suggest, grasp the outlines of the account above. They know that photography aims to produce pictures in which we see things as they really were, i.e. that it aims at accurate seeing-in. They know that, since it has this aim, things will have been designed in such a way as to secure it. And they know that, if things work as they are supposed to, that will be the result. Anyone who knows all this will indeed, when presented with something that looks to be a photograph, take himself to be licensed to believe things to be as they appear in the photo to be. Of course, things can go wrong, at almost every

[4] Note I do not say that the belief is itself guaranteed to be right. I may make a mistake, in the move from seeing-in to belief. Nor can we introduce a further norm, governing this transition, to identify, in a non-trivial way, the successful case: see the discussion below about the norms that govern belief.

stage in the photographic process. Sometimes it will be apparent that they have. If the snow looks greenish, no doubt the photograph has not been correctly exposed. But not every failure of the system to work properly will always be obvious. So there is often an element of gambling involved in the viewer's reasoning. He will work on the assumption that things have worked properly, without any guarantee that they have. However, gambles like this are a feature of epistemic life. Viewers can no more attain certainty in their beliefs when looking at photographs than they can in any other epistemic situation. Certainty is not necessary for either knowledge or rational belief.

These thoughts also enable us to address this section's other topic, the special phenomenology of looking at traditional photographs. If the viewer takes himself to be looking at a photograph, and takes that to be the product of a photographic system in which everything has worked as it should, then he takes his pictorial experience to be accurate as a matter of necessity. It is not just that he happens to be seeing in the picture things as they really were. That his experience is accurate in this way is secured by the fact that it is a photograph, and the fact that everything involved in its making worked properly. With a handmade picture, in contrast, while what we see in it may well be how things are or were, this accuracy is never guaranteed simply by the kind of picture it is. What is special about our experience of photographs, I suggest, is just that we take them to be in this way guaranteed to support accurate seeing-in. That is the source of our sense that they place us in a relation to the photographed events which is specially intimate and direct.

Perhaps it will seem puzzling how this could be the phenomenological difference between our experience of photographs and that of handmade pictures. Can our beliefs about the kind to which our mental state belongs really affect its phenomenology in this way? But consider an analogous case. Compare memories and imaginings. Often these can be very similar. It is possible to be unsure whether the image one calls to mind is a memory of some distant event or merely an imagining—perhaps one induced by having in the past being told about that event. In either case, the image might be accurate: it might represent that event as being as, in fact, is was. In neither case need this be an accident. But, even if one already has grounds to take the image to be accurate, if one settles the issue in favour of memory, one's experience may shift (Judson?). Again, we might say that we seem more intimately, more directly connected to the events our mental state represents. And again, we can give an account of this intimacy in terms of presumed factivity. To take one's state to be a memory, and not just an imagining, is to take it to be factive, to reflect the facts simply in virtue of the kind of state it is. If that is the source of the phenomenological difference in this case, why should it not be in the case of photography versus handmade pictures?

**Photography and perception**

We are now in a position to turn to the last of the commonplaces concerning traditional photography with which we began. This is that, while the epistemology of handmade pictures resembles that of (verbal) testimony, the epistemology of photographs is closer to that of perception itself. How does this claim fare, given our conclusions? Since handmade pictures are only our topic insofar as they contrast with photography, I will not pursue the analogy between them and testimony. Suffice to say, first, that in both cases accuracy essentially depends on how someone (the artist or informant) takes things to be; and second, that

handmade pictures certainly are not analogous to perception in the ways I am about to describe.

On the account I am offering, photography does indeed bear some interesting analogy to perception. Like our experience of photography, perceptual experience is factive. And in both cases the factivity is a result of a system designed to support accurate experience, one that is working as it should. For perceptual experience is also the output of a complex system each element of which is governed by norms of proper working. (Think, for instance of what it is for the lens of one's eye to work properly—it should focus on the retina light coming from selected areas of the environment.)[5] Those norms do not reduce to the idea that the system as a whole should produce experiences that reflect how the world is, but that overarching goal is their ultimate source.

However, there are key disanalogies too. First, the source of the norms governing the photographic system is human design. The light-control mechanism in a camera should work in suchandsuch a way only because it has been made so as to obey certain principles; and those principles were devised only in pursuit of the wider goal, on the part of the designers of the system, of bringing about accurate seeing-in. In the case of perception, in contrast, the source of any normativity lies in evolution. The eye is governed by norms of proper functioning only because it has evolved so as to perform a certain task, one that yields a selective advantage.

Second, the concept of perception stands in a different relation to factivity than does the concept of photographic experience. Being factive is just part of the definition of perception, as it is not part of the definition of our experience of photographs. In the latter case, we avoided stipulating that the experiences that interested us were accurate, so as to not to trivialise the claim that those experiences are factive. Instead, we picked out the experiences by the idea of the system working properly, and then made the substantive claim that the experiences so identified will all be accurate. It is at least unclear that we could do the same in the perceptual case while operating with our ordinary notion of perception. For, as I noted when first introducing factivity, it seems to be stipulative that perceptions are accurate. (The same is true of memory and knowledge.)

Does this second difference threaten our project? That some of our concepts of mental states involve factivity by definition does not show that every concept of a mental state that is factive must do so. At the least, then, our claims about photography are not under threat. What may seem in danger is the parallel between its epistemology and that for perception. In both cases there is the same epistemic good, factivity. But if it comes by stipulation in the one case, and is earned in the other, aren't the two in all other respects radically different? This is too quick. The fact that our ordinary concept of perceptual states has factivity built into it does not show that every way of thinking of those states must do so. Thus we might hope to make a non-trivial claim about factivity framed in terms other than the everyday notion. We might try to

[5] Here the norm governing this part of the perceptual system is very similar to that governing an element in photography, the focussing of the camera lens. It does not follow, however, that in general the two sets of norms are similar. Indeed, nothing I say here or elsewhere amounts to accepting the long tradition that tries to see photography as somehow reproducing the workings of our visual system. On my view, the link between them is limited to the rather abstract point of analogy above.

make sense of the norms governing the operation of the perceptual system, identifying a class of mental states as the output of the system when those norms are followed, and go on to make the substantive claim that every one of those states is accurate. If we did all this, we would offer an epistemology of perception closely parallel to that here suggested for photography. It is not our business now to assess whether that can be done. The point is that, until we have tried, there is no reason to doubt that the parallels between photography and perception go further than merely being factive.[6]

**Other information-preserving systems**

In effect, traditional photography combines two features. It is pictorial, in that it allow us to see in it what it represents. And it is information-preserving, in that how it represents things as being does not depend directly on how anyone took things to be. It is the intersection of these two features that is responsible for the phenomenon I take to be central to its epistemic and phenomenological distinctiveness, factive pictorial experience. But how exactly do the two features work together to secure that special status? One might worry that the epistemic work is really done by photography's preserving information, the fact that it does so via seeing-in only really contributing to phenomenology. To examine this worry, consider other forms of information-preserving system.

Traditional photography is not unique in deploying causal processes that preserve information to produce accurate representations of the world. Many other devices do just that: think of scientific equipment such as geiger counters and blood pressure monitors; or devices that produce pictures, although not ones in which we see the facts they convey, such as infra-red images of heat distribution. (Snyder and Allen 1974 call cases of the latter sort 'scientific photography'.) What is true of traditional photography is true here too: these devices are designed to produce accurate representations and, when they work properly, that is precisely what they do. Are any of these devices epistemically special in the way that photography is?

There is one crucial difference between these various systems and traditional photography. Although all these systems have some kind of representation as their output (a photograph, infrared image, or reading on the geiger counter), only in the case of traditional photography do we see the facts in that output. True, in all these cases there is something like factivity: the output is, when things have worked well, guaranteed to be accurate. This fails to count as factivity proper only because that, as we defined it, is a property of mental states. True too, there are mental states provoked by these systems—those the subject forms in engaging with the output, and thereby grasping the facts it describes. Nonetheless, something important is

[6] Timothy Williamson has recently (2000) articulated a view on which factive mental states lie at the heart of epistemology. The most distinctive feature of his position is the claim that knowledge is a *mental* state, and the claim that it is the most general factive mental state. In effect, Williamson takes knowledge as primitive, and uses the notion to do work in epistemology. Like Williamson, I have made central use of the notion of a factive mental state, pictorial experience of photographs. Like him, I think that factivity in a mental state is a, perhaps the main, way in which it can be of epistemic value to us. And I agree with him that knowledge is the most general factive mental state. (Perhaps for that reason, no account of how it can be factive, along the lines just suggested for perception, looks at all promising.) However, there end the similarities between our views. I do not suggest we treat factive mental states a epistemologically primitive. On the contrary, some of what I have said suggests at least the beginnings of an account of how our experience of photographs might come to be factive. I do not assume that knowledge is unanalysable. Nor does anything I say depend on the view that knowledge is a mental state.

missing: factivity in those mental states themselves. More precisely, there is no way to identify those states so that the claim that they are factive is both non-trivial and true. We cannot simply assume that the mental states formed in response to accurate representations will themselves be accurate, for at that stage (as at every other involving people) mistakes are always possible. So we cannot identify the factive states as those formed in response to the output, even when the system producing that output has worked properly. To identify the relevant states as those that get the facts right would be trivial. What we need is norms governing the formation of the mental states that secure accuracy without stipulating it.[7] In the case of seeing-in, we saw, suitable norms were available. In all other cases, I argue, they will not be.

There are two sorts of case to consider. In the first, such as the geiger counter, we do not have pictorial experience at all. We see the reading, and, using our knowledge of how the system works, form the belief that the radioactivity is at suchandsuch a level. The second sort of case is scientific photography, such as the infrared image. Here we recover the facts via seeing-in, but we don't see the facts in the image. We don't see heat here and cold there in the picture, if only because heat and cold are not the kinds of thing that *can* be presented visually (not directly, at least). Rather, what we see in the image is the distribution of colours across the imaged terrain, on which basis we form beliefs about heat. That doesn't amount to factive pictorial experience, since what is seen in is not even true to the facts. In all likelihood, the objects imaged will be completely different colours from those used to indicate their temperature. Here too, then, if any mental state is factive, it is belief.

Thus in either sort of case belief is the only candidate for factivity. But belief is never governed by norms of the sort we seek. The norm governing belief is not, for instance, *form your beliefs so as to accord with whatever the output suggests*. The central, perhaps only, norm governing belief is *believe what is true*. To see this, compare someone using a geiger counter he has reason to think dicky with our earlier viewer of a blurry photograph. Both may conclude that their devices have malfunctioned, and both may come to the right conclusion about how, in fact, things were with the objects that form the input to the system. But while the viewer of the photograph ought to reach this conclusion by seeing blurred objects in the picture, there is no sense in which the user of the geiger counter ought to do anything similar with belief. He does not come to the correct view of the facts by first following norms that yield an incorrect view of them, and then correcting for malfunction in the system. He should, of course, read the meter correctly. Even if what the meter says is wrong, he needs to form a belief that accurately reflects the reading. But that is a belief *about what the meter says*, not a belief about the radioactivity in the environment. When it comes to the latter, he will not simply mould belief to the reading on the meter. He will instead be guided by his suspicions about the meter's bias, adjusting his beliefs away from what the reading suggests. In forming that belief he aims to get the facts right. It is not even clear that he could aim to believe the facts to be the way the malfunctioning meter suggests. Thus the only mental states of his that represent the input (radioactivity in the environment) are guided simply by the norm of getting the facts right. Similarly, then, in the case where everything *has* worked as it should.

[7] In effect, this is to draw the states themselves into the system, so that they are its final output. Indeed, I described our experience of traditional photographs in just those terms.

Thus there are here no norms here by which we might identify without triviality a set of factive mental states.

None of this is to deny that devices such as geiger counters and the various forms of scientific photography are epistemically valuable. Of course they are. Indeed, they offer us things traditional photography does not: access to invisible features of the world, and knowledge of them that is precise in ways a photograph generally is not. However, there are also, we can now see, epistemic benefits traditional photography offers that they do not. Only it supports factive experience that represents how things are with the input to the system. Earlier I explained that this is epistemically valuable in two ways. Whatever the epistemic value of these other systems, they cannot offer quite what traditional photography does.[8]

**Walton's view**

Among the various accounts of what is special about traditional photography, one theory has provoked far more discussion than the rest. This is Kendal Walton's view that photographs are transparent, that when one sees an object in a photograph one literally sees that thing. Photographs are thus, like mirrors, spectacles and telescopes, aids to vision: they extend the range of what we can see. Comparing this position with mine may help clarify the latter.

Walton bases his claim in part on the fact that photography is information-preserving. Thus our positions are motivated by at least one consideration in common. Of course, the positions thus motivated are distinct. To say, as I have, that photographs support factive pictorial experience of their objects is not to say that they allow us to see those things. Seeing and photographic seeing-in are both factive, but it hardly follows that they are the same state. Nonetheless, one may wonder whether the two aren't related. Doesn't Walton's position entail mine? Seeing is a form of perception, and thus a factive state. If photographs allow us to see their objects, as Walton claims, doesn't it follow that they support factive experience of those objects, as I propose?

This is a mistake. Walton's account and mine are, in fact, logically independent. To grasp this, we need to distinguish perceiving an object from perceiving a fact. If I see you, I see an object. If I see that you've put on weight, I see that some state of affairs obtains, i.e. I see a fact. Seeing facts always involves seeing objects. It is by seeing you, or perhaps the deep impression you've left in the sofa, that I see that you've put on weight. However, it is possible to see objects without seeing facts. That is the situation with animals who have vision, but lack concepts to grasp the facts vision puts before them. Parallel claims hold for pictorial experience. There is a difference between seeing an object in a picture and seeing a fact therein. Again, the latter requires the former but, arguably at least, not vice versa.

Walton's claim about photography is, in effect, that it allows us to see objects. My claim concerns, not seeing, but pictorial experience, seeing-in. But more importantly for present purposes, it bears on pictorial experience of facts, not of objects. It is only insofar as pictorial experience can be of facts that it can be factive. For factivity amounts to the guarantee that

[8] This is not to say that no other device, possible or actual, could offer these benefits. In particular, there might be analogues in sensory modes other than vision that offer the equivalent in that mode of factive seeing-in. Sound recordings might provide an example.

what the experience presents as fact, is so. Clearly the notion makes no sense as applied to experience (pictorial or perceptual) that is purely of objects.

Still, when pictorial experience is both of objects and of facts, what then? Viewers of photographs surely almost always do see both objects and facts in them. They see, for instance, not just an old lady and a teacup, but that the woman is holding the cup. Since Walton claims that seeing things in photographs is seeing, it is at least open to him to claim that seeing facts in photographs is literally seeing those facts. Developed in this way, his position would indeed entail mine. When photographs allow us literally to see the facts, they support a kind of factive experience—just as I suggest.[9]

However, I doubt Walton would accept this development of his view. He wants to allow that when we see objects in photographs, that counts as seeing them even if we *mis*perceive their properties. [Ref] He wants to say, for instance, that when I see the overexposed photograph, I see the snow even though my experience misrepresents it as greenish. Since ordinary object seeing certainly can allow us to misperceive objects' properties, it is very natural for Walton to adopt this view about seeing mediated by photographs. Moreover, it spares him the work I had to undertake to handle such cases. In order to preserve the claim that photographic experience is factive, I had to set such cases aside. They are ones in which the photographic system has not worked properly. Walton can treat them as just as central as photographs that get things right. He ought, then, to resist the development above. If he does, our views remain independent.

If our views are distinct, which is correct? Perhaps it is an advantage of Walton's position that it can treat the overexposure case as perfectly standard. It clearly is a case in which *something* has gone wrong, but it's a matter for judgement whether the best way to accommodate that fact is by banishing it from the central instances of photography, as our theory sees them. But even if Walton's position has an advantage over mine in that respect, in general mine offers benefits his does not. Both give an account of the special epistemology and phenomenology of traditional photography. But Walton's view does so by making implausibly strong claims. He defends the claim that photography is transparent by appeal to the claim that it meets two conditions on seeing: it is information-preserving, and it invites certain kinds of mistake in grasping the photographed scene. But these conditions, as well as lacking the justification only intergration into a proper theory of vision could provide, are not sufficient for seeing. They are met in cases in which it is clear that one does not see the object from which information is transmitted (Currie ?) And anything missing from these cases that might explain why they are not transparent seems equally likely to be missing from photography. (For a candidate, see Meskin & Cohen?)

My account avoids these pitfalls. By refusing to claim that photographs are transparent, I undertake none of the burden of defending that striking claim. In that sense, the materials I

[9] Of course, Walton thinks the state of mind in question is experience (seeing), I think it is pictorial experience (seeing-in). But that doesn't affect the logical situation. Walton thinks some pictorial experiences,those of photographs, are *identical with* some experiences of the photographed objects. So the entailment goes through.

offer come cheaper—indeed, I hope to have persuaded the reader, at a perfectly sensible price. Yet they do all the work Walton wanted. Only a spendthrift would choose to pay more.[10]

**Digital photography at last**

I have attempted to say what is special about traditional photography, in terms both of epistemology and phenomenology. Photography aims at factive pictorial experience and, when all works as it should, that is precisely what it produces. How, if at all, does the arrival of digital techniques for making photographs affect this situation? Can digital photographs too be special in these ways? And if they cannot, does their presence also threaten to undermine the specialness of photographs made by traditional means?

*Prima facie* digital photography can be an information-preserving system. Many digital photographs show how things were, and do so without essentially depending, for the transmission of information from the scene photographed, on how anyone took things to be. And certainly digital photographs support seeing-in: we see things in them as surely as we do in any other kind of picture. At first glance, then, the materials for factive pictorial experience are present. The question is whether these appearances are borne out.

Why might they not be? There are two possible sources of difficulty here, and two broad ways in which those difficulties might make a difference. Difficulty might lie in the very constitution of digital photography: perhaps something about *any* digital photographic system prevents it from sustaining factive pictorial experience. Or it might lie in the possibilities digital photography opens up, even though those possibilities are not always exploited: possibilities for manipulating the image so that it no longer presents the facts (or no longer does so because it preserves information about them). And these difficulties might play out in preventing experience of digital photography counting as factive, or in preventing our *recognising* such experience as factive, even though it is.

Consider first, then, the possibilities that digital photography opens up, and in particular the ease with which it allows photographic images to be manipulated. Insofar as this possibility is exercised, its impact will be limited, at least from a theoretical point of view. After all, it has always been possible, even with traditional photography, to manipulate what is seen in the picture. (Think of the airbrushed cadre and the painted fairies.) If it is now easier to do this, this will presumably matter because it affects how often what seems to be an information-preserving picture is really one. By itself, that suggests nothing about the cases in which manipulation has not occurred. If some digital images support pictorial experience that is factive, the existence of other images that do not hardly undermines the position of the former. What might be affected is our attitude to images, both manipulated and otherwise. If we are unable to tell the pure cases from the manipulated ones, we may cease to treat any cases as supporting factive seeing-in. If we did, then, while some digital photographs would indeed offer us the distinctive epistemic benefit I described, we would not be in a position to exploit it. Nor would we take these photos to put us in the distinctively intimate and direct relation to their objects described. Our experience of these photos would lack that distinctive

[10] Walton also toys with weakening the idea of transparency. Photographs do not allow us to *see* their objects, perhaps, but they do offer us a form of contact with them that is genuinely perceptual (ref). Though weaker, this claim remains much stronger than mine. Some of the objections to Walton's original position that fail to touch my account, continue to apply to his weakened view.

phenomenology. And if we took manipulated digital images to be indistinguishable, not just from unmanipulated ones, but from traditional photography, it is indeed possible that our attitudes to the latter would shift in the same ways.

Still, whether any of this occurs is largely an empirical matter. It depends on how often we encounter manipulated images that we can't tell from unmanipulated ones. More precisely, it depends on how often we *take ourselves* to be in that situation. And it depends on our psychology, on how we react to whatever frequency of false positives we take ourselves to encounter. All this remains to be seen, as digital photography becomes ever yet more common. Not a great deal can be said about it *a priori*.

It is more interesting, from a philosophical point of view at least, to consider whether anything in the constitution of digital photography prevents it sustaining factive seeing-in. This does promise to bear on whether such photography can offer the benefits yielded by photography of the traditional kind. Any threat here lies in locating differences between the normative contexts in which digital photographs and traditional ones are produced. After all, it was only in relation to norms of proper functioning that we were able to make sense of the idea that photography supports factive pictorial experience. Change the norms, and that idea may be rendered inapplicable. There are two ways in which the norms governing digital and traditional photography might differ.

The first returns us to the theme of manipulation. Above we worried that frequency with which digital images are manipulated (or thought to be) might affect our attitude to them. Here the thought is rather that, whether or not possibilities of manipulation are ever exercised, or believed to be, if the system has been *designed* to make manipulation possible, that may affect the epistemic benefits it can offer. For, if the system has been designed in part to allow what is seen in the image to be changed, it is no longer clear that there *is* a norm of proper working that will, when followed, secure that what we see in the picture is how things really were. Of course, the issue here is again in part empirical. Have digital photographic systems been designed to enable manipulation, or is the greater ease of doing so a mere side effect of those ends they have been designed to achieve? What is not an empirical matter is what the consequences might be of answers to that question.

The second possible source of normative difference lies closer, perhaps, to the very heart of digital photography. Consider interpolation, and in particular its role in the way some digital cameras process colour. Digital cameras record the information in light coming from the photographed scene and falling on a photosensitive screen. To record colour information, discrete sensors need to be deployed: one for each of the three additive (?) primary (?) colours. This creates the engineering problem of how to locate three such sensors at each site on the screen. One solution commonly adopted is to place a single sensor at each site, alternating sensors across adjacent sites. This is cheap, in engineering terms, but leads to loss of information. If a site maps a particular spot on the photographed object, the colour of that precise spot cannot be recorded at the site: that would require sensors of all three types, and only one is available. Interpolation provides the means to handle this difficulty. Information from different sensors at adjacent sites is compiled, and the camera's software calculates the likely colour of larger areas on the photographed object by comparing the limited colour information available from points within those areas. In effect, the camera makes an educated

guess as to the colour of a given area, based on the colour that would typically generate that pattern of stimulation across the sensors at adjacent sites.

There is no human intervention in this process, beyond that involved in designing the device. Thus interpolating colour is not a matter of representing it to be a certain way because someone took it to be so. The camera does the 'guessing', not some person; and it does so on the basis of what it statistically likely, not how it 'saw' things to be (whatever that would mean). Nonetheless, interpolation is not clearly an information-preserving process in the sense required to sustain factive seeing-in. Once the software has done its work, the system represents a certain area on the photographed object as of a given colour. However, its representing it as so does not entail that the area is that colour, not even when everything has worked as it should. The limited information received about colour is in fact consistent with the area being coloured differently. Perhaps, for instance, it is not a uniform colour at all. Thus interpolation threatens factivity, yet interpolating is precisely what these cameras have been designed to do.

The threat to factivity from interpolation is limited. True, interpolation is not only deployed in capturing colour—digital zooms, for instance, also make use of the process. Nonetheless, it is, at least as far I know, and at least for the moment, a process of restricted application. Not all cameras solve the problem of recording colour by interpolating. And no camera deploys interpolation to represent every property of the photographed object. Still, processes for handling information in digital form, whatever kind of information that may be, are always under pressure to capture more by recording less. So we may see more interpolation in future. Whether we do or not, we can conclude the following. To the extent that colour is interpolated in a system, experience of the pictures it produces will not be factive, with respect to colour properties of the objects photographed. More generally, to the extent that digital photography deploys interpolation capturing a certain feature, to that extent it cannot support pictorial experience that is factive, with respect to that feature. How many features are handled by interpolation, and to what degree they are, is thus one determinant of how far digital photography sustains factive seeing-in.

**Robert Hopkins**
**University of Sheffield**

## References

Cohen, Jonathan & Meskin, Aaron. 2004. 'On the Epistemic Value of Photographs' *Journal of Aesthetics and Art Criticism* 62, 197-210.
Currie (crit. of KW)
Analogy with perc/t.y?
Judson?
Savedoff, Barbara. 2000. *Transforming Images* Ithaca: Cornell University Press.
Snyder and Walsh Allen. 1974.
Walton, Kendall. 1984. 'Transparent Pictures: On the Nature of Photographic Realism' *Critical Inquiry* 11, 246-276.
Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford: Clarendon Press.
Wollheim, Richard. 1987. *Painting as an Art*. London: Thames & Hudson.