

The World We Learn to See

On Photography, Artificial Images, and the Future of Perception

When I walk through a city with a camera, almost everything that will later appear in my photographs is already there.

The street was built before I arrived. Cables were laid, barriers installed, timetables written. People go to work, wait for a train, buy things, move goods. Light falls across façades. Signs regulate access. Traffic determines movement.

None of this needs me in order to exist.

And yet I later claim to have created something.

That claim belongs so naturally to photography that its strangeness is easy to overlook. A writer can invent a city. A painter can place a figure where none existed before. The photographer, at least within a certain tradition of the medium, works with a world that precedes him.

What, then, does he create?

I decide where to stand, what remains inside the frame and what stays outside it, when to wait and when to move on. Later I decide which image remains, which disappears, what is placed beside what, where something repeats, and where a deviation begins to matter.

The world provides the elements. The photographer brings them into new relationships.

A barrier remains a barrier. But placed beside an image of people waiting, a delivery rider, a construction site, or a system of access, it can become part of a larger perception. It is no longer only an object. It begins to speak about access, movement, labour, or order.

That relationship was not simply contained within the object. It emerges through the work.

For me, this is part of what creation in photography means.

But the artist's control ends earlier than we sometimes assume. He does not simply produce a meaning and deposit it in an image. He may intend something, recognise a relationship, work with great precision. Once the image encounters another person, however, he is no longer the only one involved in what it becomes.

The viewer does not arrive empty.

We bring experience, memory, knowledge, fear, and desire. A barrier can mean protection or exclusion. A uniform can suggest safety or threat. An empty public space can signify freedom, abandonment, or control.

Meaning therefore belongs neither entirely to the work nor entirely to the viewer. It arises in the encounter between them.

That does not mean that everything is arbitrary. An image also sets limits. Some readings are more readily supported than others. A work is not an empty mirror. But neither is it a container from which the viewer simply retrieves the message the artist placed there.

Perhaps the artist creates less a final meaning than the conditions under which meaning can emerge.

Photography, in this sense, does not only produce images. It produces situations of seeing.

And none of those situations begins at zero.

We have already learned how to see.

We learn which images stand for success, danger, beauty, poverty, power, or intimacy. Films show us what loneliness can look like. Advertising gives visible form to prosperity. News media shape images of conflict. Museums contribute to what appears culturally significant. On digital platforms, certain bodies, rooms, faces, and gestures return again and again.

What we experience as our own way of seeing is therefore never entirely our own.

We see through experience. But those experiences themselves took place in a world already filled with images.

Then we give something back.

We photograph, build houses, design spaces, write texts, buy things, make decisions. What we have perceived becomes part of what we produce in turn. Other people encounter those things and are changed by them as well.

We create cultural forms. Those forms shape how we see. Out of that altered seeing, new forms emerge.

Artificial intelligence did not invent this cycle.

For that reason, I find it unhelpful to treat AI as though an alien force had entered what was previously an entirely human culture. These systems are also made by people. People develop and finance them, operate them, produce much of the material from which they learn, write prompts, select results, reject others, modify them, and distribute them.

Human origin, however, does not imply human control over the whole.

No individual created a city either. Yet cities are human-made. They emerge from countless decisions accumulated over decades: planning, ownership, labour, politics, technology, migration, consumption, conflict, and accident.

Once those decisions have taken material form, they confront the individual as reality.

A road was built by people and can still determine my path. A barrier is human-made and can nevertheless stop me.

What we create becomes a condition of our own actions.

Technical systems belong to the same movement. We build them, use them, and respond to their results. What follows changes the culture in which further decisions are made and new systems developed.

Images occupy a particular position within this cycle. They are outcomes, but once seen they can become starting points themselves.

An image can influence how someone perceives their body, how an architect thinks about a space, or how a future city becomes imaginable. A photographer may photograph differently after encountering an image. A child also learns through images what is considered beautiful, normal, desirable, or threatening.

This is where generative AI becomes culturally interesting to me.

The familiar question of whether AI can be creative, or whether an image made with it can be art, seems too narrow. An image can move a person even though the system that produced it experiences nothing. An artist can bring personal experience into form through a generative system, selecting, rejecting, changing, and combining. The absence of experience on the part of the tool does not decide the meaning of what is produced with it.

More interesting is the part such systems may come to play in determining which images we encounter before we begin to make meaning from them.

Culture has always prestructured our seeing. Parents do it. Schools, religions, media, museums, advertising, and artists do it. There is no such thing as perception without conditions.

With generative systems, however, the scale may be changing.

People in very different countries can ask the same small number of systems to show them a successful person, a happy family, a modern city, a beautiful apartment, a romantic encounter, or a desirable future. The system has to give those concepts visible form.

A viewer may interpret that form in very different ways. But before interpretation begins, something has already been selected: bodies, spaces, gestures, light, relationships, forms of proximity and distance.

The system does not finally decide what success, beauty, or the future mean. But it can influence which visible forms become most readily associated with them.

A road does not tell me where I must go. It makes some routes easier and others more difficult.

A technical image system does not tell me what to think. But it can influence which images are most readily available to thought.

Generative systems could, in this sense, increasingly become part of an infrastructure of seeing.

Perhaps this concerns me because it resembles the city so closely. A city consists of decisions that have become environment. It opens possibilities and closes others. It directs without determining every movement. The individual remains capable of acting within it.

Cultural systems also rarely operate through commands alone. More often, they create probabilities.

The obvious concern is that if people everywhere use the same systems, eventually all images will look the same.

I think that is too simple.

Cultures do not disappear because people use the same software. They bring language, history, religion, local experience, memories, and conflicts with them. The subjects can remain different while something else begins to converge.

Different faces. Different clothing. Different buildings. Different rituals.

And yet similar ideas of how success should become visible. Similar forms for intimacy. Similar light for melancholy. Similar compositions for power. Similar images of what a modern city or an attractive future should look like.

The cultures would still be different. But they might begin to make their differences visible in increasingly similar ways.

The result would be less a uniformity of content than a convergence of visual grammar.

That, to me, is one possible danger of an automated image world. Not that it produces too little that is new. Generative systems can produce almost unlimited numbers of new images. The more important question is whether those images also allow us to see something differently.

A billion new images do not amount to a billion new ideas.

We could see constantly changing faces while retaining the same idea of beauty. New houses and the same idea of prosperity. New cities of the future and the same idea of progress.

The surface changes while much beneath it may remain remarkably stable.

A culture could then produce enormous quantities of visual novelty without expanding its idea of the world to the same degree.

This is where art becomes interesting to me.

Art arises inside a culture and carries its assumptions with it. But it can do something with those assumptions. It can establish relationships that were not previously self-evident. It can make the familiar strange, or arrange something we encounter every day in such a way that an existing category is no longer enough.

To do this, art does not necessarily have to invent something materially new.

A photographer can show the same city thousands of people move through every day and still make visible a relationship that was not previously part of their perception.

Photography, of course, is not innocent either. Photographers repeat clichés. The history of the medium is full of images that reinforced stereotypes and structures of power rather than breaking them open. The photographer also enters the world with images already in mind.

But those images must pass once more through a world that was not built entirely according to them.

I have repeatedly experienced this in my own work.

I began by seeing cities largely through their visible orders: routes, thresholds, infrastructure, rules.

Then I stood before an entrance that formally existed and watched a person with a walker encounter that same entrance very differently from the way I would. My own experience had offered me an apparent self-evidence that was not self-evident for someone else.

At another place, a sign said: Please do not lean bicycles here. The bicycles were leaning there anyway. Not one, but many. The rule was visible, and so was its everyday circumvention.

People wait within systems, use them, bypass them, or resist them. Eventually the order that had first interested me could no longer be separated from the people who had created it and now lived within its conditions.

At Potsdamer Platz in Berlin, that shift became even clearer.

Beside surviving remnants of the former Berlin Wall, a man stood directly on the line in the ground that still marks its course. In front of him, present-day traffic moved through the city. He remained there for several minutes, looking into the flow. The sun was in his eyes; again and again he raised a hand above them so that he could continue to see.

I did not know what he was doing there. Perhaps he was simply watching the traffic. Perhaps he was waiting for someone. Perhaps the line meant nothing to him like what it had begun to mean to me.

Precisely because I did not know, the situation remained open.

Present and history occupied the same surface. A person stood between them. I had not brought that relationship with me. I had encountered it.

I can go out with one idea and return with another.

That does not happen automatically. A photographer can just as easily find only what he was already looking for. He can use the world to confirm an idea that was complete before he arrived.

But the possibility of contradiction remains.

Perhaps this is one future significance of photography. Not because photography is automatically truthful, or because a photographic image is inherently more valuable than a generated one. A generative system can also surprise us, produce unexpected relationships, or lead an artist towards something they had not seen before.

In photography, however, the contradiction can come from something that does not exist in order to become an image.

The street, the body, the weather, the room, and the social situation have an existence outside the photograph. They do not have to conform to my idea of them.

The world can contradict the image before the image exists.

The openness of a cultural cycle may therefore depend less on how many new images it produces than on whether something can enter it that has not already emerged entirely from its own existing ideas.

Photography does not guarantee such an interruption. But it contains that possibility.

Between an idea and an image, there is still a world.

A world that can refuse us.

At this point, control begins to mean something different as well.

Culture has never developed entirely according to plan. People meet, adopt forms, misunderstand them, alter them, and connect things that were previously separate. New generations use what they inherit differently from the way the previous generation intended. Art would be difficult to imagine if its results were fully determined in advance.

Development therefore always contains some degree of non-control.

With artificial intelligence, however, this is not an innocent thought. Powerful technical systems can influence decisions, concentrate power, and multiply errors on a scale that extends far beyond art and images. Where the possible consequences become large, loss of control is no longer an aesthetic experiment.

Yet a contradiction remains.

If we permit only outcomes that stay within our existing ideas, we do not merely limit risk. We also limit the possibility that something might emerge which changes those ideas.

This does not make loss of control valuable in itself. Unpredictability is not the same as development. A technical system could surprise us endlessly while still varying within the same cultural assumptions. It could become difficult to control without taking us culturally beyond what we already know.

What matters, then, is not the loss of control itself, but what the opening makes possible.

A living culture needs repeatable forms. Without them, we could barely understand one another. But it also needs situations in which what already exists is not simply confirmed. That can happen through other people, other cultures, history, the material world, or through something that initially has no designated place within an existing order.

Technical systems may also be capable of playing such a role.

If so, two demands begin to conflict. We need control because the consequences of our actions are real. At the same time, we need a domain in which not every result is fixed in advance, because complete predictability would also mean allowing only what our existing imagination can already accommodate.

The task can therefore hardly be to choose between control and loss of control.

It is to create conditions under which deviation remains possible without confusing openness with irresponsibility.

This becomes more urgent as the boundaries between tools continue to shift. Photographers will use generative systems. Generative systems will be connected to cameras, sensors, archives, and other sources of data. Human decisions, technical processes, and encounters with the world will increasingly overlap.

The distinction between a photographic and an artificial image may therefore become less clear. The conditions under which images are produced and circulate will become more important.

Whoever determines those conditions does not determine every individual meaning. But they influence the space in which some meanings become more probable than others.

Human beings have always created structures that later act back upon them.

We build cities and then move through the routes we have laid down. We create institutions and learn to think within their categories. We produce images and begin to see through them.

Artificial intelligence does not invent this principle.

But it may alter its speed, scale, and reach.

What we create returns to us.

And when it returns, it is no longer only the result of our ideas. It becomes part of the environment in which our next ideas emerge.

Neither complete control nor maximum unpredictability therefore seems convincing as an answer. A culture cannot renew itself if it excludes every deviation. But nor does it become more open simply by abandoning control.

The more difficult task may be to preserve something in between: a space in which contradiction remains possible without giving up everything that allows such a space to exist.

For photography, that space is comparatively easy to describe.

Between my idea and the image stands a world that does not have to obey me.

For the technical systems we are now creating, the answer is less simple.

We will continue to produce images. Those images will shape perception, and out of that altered perception further images, spaces, and realities will emerge.

We cannot leave this cycle.

The decisive question is therefore not whether we can control it completely.

It is whether we can create conditions under which it remains open without openness itself becoming a danger.

A culture does not renew itself by producing ever more new images.

It renews itself where something still has the capacity to change the images through which it has learned to see the world.