

The Devil in the Digital Details

Digital abstraction and concrete reality

Brian Cantwell Smith*

1 Introduction

It cannot be denied that computing and information technology have had—and are continuing to have—a monumental impact on the creative and documentary arts. The only plausible precedents are such epochal transformations as the invention of writing, or the press.

Less clear, however, is *what it is* about computing that is responsible for this upheaval. What makes information technology *special*, so that rendering art and music in computational terms wrecks such havoc with our understanding of *identity, materiality, ownership, originality, performance*, and perhaps even *value*?

One obvious place to look is to the notion of *information*. But information a curious notion. From one perspective, the concept of information can seem so vapidly general as to verge on the banal. Yet at the same time, specific technical notions of information are being intensely researched in as many as a dozen fields. I do not deny that the role of information in the arts a critical topic—one that would warrant its own book or conference. Still, information alone cannot explain our current predicament; no one could be so arrogant as to claim that information was invented in our own era. In one form or other, information has been around for millennia—perhaps since the dawn of time.

Another idea about what is fueling the computational revolution, and its impact on our lives, is the notion of *digitality*. Intuitively, it makes sense to ask whether digitality might be key to what makes information technology special, since, at least in automatic machinery,

*© 2002 Brian Cantwell Smith
Kimberly J Jenkins University Professor
Philosophy and Computer Science
Duke University

Franklin Center Room 223
2204 Erwin Road
Durham, NC, 27708-0402, USA
bcsmith@duke.edu

digitality seems genuinely novel. The notion of digitality has also captured the public imagination. Whereas talk of the “Information Age” seems almost passé, scholarly texts and popular accounts still trumpet the coming of the “Digital Age.” Best-sellers such as Negroponte’s *Being Digital* are just the tip of the iceberg. As of this writing, Barnes & Noble reports more than 8,000 books with the terms ‘digital’ or ‘digitality’ in their title.

In the creative arts, digitality’s impact has been immense. But scale of impact has not been matched by depth of understanding. It is widely agreed that the *rendering into digital form* of images, sounds, records, and ideas has unleashed considerable conceptual confusion. Somehow or other, in ways we need to understand, digital images, recordings, and texts break the bonds of time-honored norms of *identity, production, ownership, reproduction*, etc. If we can understand what it is to be digital, therefore, we should thereby get a leg up on disentangling some of the most vexed issues about the nature of art in the 21st century.

These, then, are the goals of this paper:

1. Constitutively, to understand *what it is to be digital* or discrete—as opposed, say, to being *continuous*.
2. Pragmatically, to understand *what digital systems are good for*—and what they are not good for; and
3. Consequentially, to understand *the impact of digitality* on our understanding of: creation, ownership, identity, materiality, reproducibility, and the like.

2 Properties

Three properties of digital systems are immediately identifiable: their *perfection*, their *abstractness*, and their *dynamics*. Describing these three does not constitute a theory of digitality. It merely spells out what a theory of digitality must explain.

2.a Perfection

Digital systems are, in a remarkable sense, perfect. When encoded digitally, a system can be *flawlessly copied*, without error, an infinite number

of times. No loss, no corruption, no friction, no accumulating impact of dirt or rust. “Perfect sound forever,” said Sony, in when it introduced the compact disc. Even if we know better, now, there was something right in their proclamation. No scratches, no noise, no irritating static. And no decay. In Bangladesh, religious manuscripts perpetually disintegrate, inexorable victims of insects and humidity. If only we could record them in digital form, we are told, they would be immune to rot. *Scanning as transubstantiation!* This is truly heaven on earth: the abstract purity of Plato’s realm rendered incarnate, in an endless string of os and is.

Something special is required, for this digital perfection to be achieved. There must be a determinate set of judgments, or properties, or types, in terms of which the system can be completely characterized. Constitutively, that is, in order to be digital, a phenomenon must succumb to a finite series of informationally-complete black-and-white judgment calls. This holds of such ordinary “digital” notions as: scoring a basket in basketball; moving the pawn to $\kappa 4$; writing down the letter ‘A’; making a copy of a text, a poem, or a musical score; cutting a board between six feet and six feet one inch long. All these things can be determinately accomplished—without error, ambiguity, or matter of degree.

This is where “information” comes in—a digital system is a system about which *complete information can be given in such (finite, black-and-white) terms*. If we know the answers to all those “yes/no” questions, we have “captured” all that matters about the system. Thus a chess game can be restarted, even if the board is dropped, if we know exactly which pieces were on what squares—a finite, compact list. By contrast, such systems as the state of a billiards table, a haunting smile, a painting or a musical *performance*, or cutting a board *exactly six feet long*, are non-digital because there is no finite, absolute, discrete set of facts of the matter that fully “capture” what is going on.

Sure enough, we can approximate the state or character of a non-digital system (a billiards game, a painting, an image), by using ever finer samples, to any degree that we choose; more on that in a moment. But the phenomenon itself, *at the level at which it is the phenomenon that it is*, is not discretely constituted. Unlike chess, that is, such non-

digital phenomena as billiards and paintings are not defined—cannot be wholly and completely accounted for—in terms of a finite set of well-defined “yes/no” questions.

Music is an interesting case, in this regard—because of the difference between a score and a performance. Musical *scores*—at least traditional scores, composed of the familiar suite of notes, staves, markings, etc.—*are* digital, plus or minus a bit. They can be perfectly copied. If one score gets wet, or starts to decay, or has coffee spilled on it, a new one can be made without (so we have chosen to ascribe value) desecrating Beethoven’s composition. But the performance itself—which, in virtue of its concreteness, invariably adds an unutterable wealth of detail to the sparse information of the score—is not discrete in the same way. In fact this is one way to understand what performers do: they fill in the infinitely rich detail between the skeletal sparseness of a digital score, and the ultimate thickness of a concrete, continuous (i.e., non-digital), musical utterance.

Issues of ownership and value in music can be made intelligible in terms of this divide. When we credit a (classical) composer with being the “creator” of a work, we do not view the work in question as a concrete in-the-world musical utterance, but rather as a work *under description*—a work “abstracted” according to the conventions that dictate what aspects of a work are captured in a traditionally-notated (digital) score. This “restriction of credit” to an abstracted version makes room for subsequent performances to be viewed as loci of genuine artistry, creativity, originality, etc., in their own right. When the Guarneri Quartet performs late Beethoven chamber music, Beethoven is given credit for the “digital” content of the score; the players, for the non-digital aspects of the ensuing performance. Similarly, one way to understand Factum Arte’s project is as giving us a “score” of the tombs in the Valley of the Kings, to be “interpreted,” in the future, in different performances, by different Egyptologists and archeologists.

2.b Abstraction

A second manifest property of digital systems is that they are at least apparently *abstract*. Programs, bit maps, digital data—none of these things weigh a certain amount, or have energy or momentum, nor can

they be eaten by moths, or otherwise decay. Digital data often *relates* to physical stuff, in the way that an arrangement relates to what is arranged, or a configuration relates to what is configured. But qua arrangement or configuration, digital entities are more like things which are truly abstract, such as numbers and pure ideas, than they are like concrete paintings or hand-hewn log cabins.

One way you can tell when something is (at least relatively) abstract is when it can be realized in a wide variety of materials. Famously, chess games do not have to be played with wooden or ivory pieces; salt shakers would do, or people, or a spate of suitably hovering helicopters. Similarly, letters can be formed of ink, or pencil, or jet contrails, or by arrangements of sports teams band-members during intermission. Similarly, it is because they are “abstract” that chess, unlike billiards or fencing, can be played by mail—or over the internet. Digital systems, to use a technical term of art, are *medium-independent*, in a way that non-digital systems are not.

Issues of medium-independence, it should be noted, lie at the heart of raging debates about the possibility of Artificial Intelligence. Are your thoughts digital, like moves in a chess game—implying that your identity could be uploaded onto a digital computer? Or is your mind more like billiards: inexorably tied, at the level at which you are you, to specific irreproducible facts about your material embodiment? The fate of our children depends on the answer.

2.c Dynamics

The third obvious property of digital systems, along with their perfection and abstraction, has to do with their *dynamics*. On the face of it, the most obvious dynamic property of digital systems is their *stability*. “Perfect sound *forever*,” was Sony’s claim. And think of those manuscripts in Bangladesh; what digitality promised was protection against the ravages of time. Indeed, purity, perfection, and *stability*—a kind of eerie immutability or invulnerability to the ferment of life, to the eruptive activity of concrete existence, seems almost defining of the digital realm.

Yet if perfection and epochal stability are the marks of the digital,

so too, curiously, is change. On the web, you can create, duplicate, modify with unprecedented abandon. Switching between a zero and a one takes so little energy it is essentially free. Want to adjust that memo you posted last night? No problem! Click, click, click; just one more email. And it's not just we people who change things; convergent networks and routers are in the business of moving things around, not keeping them fixed. Computing itself, in fact, once one thinks about it, is the epitome of change. "Mathematics *plus time*," it has been called. Whitehead redux: it is the *processing* of symbols, not the symbols themselves, that ultimately matters. In fact the symbols and media are increasingly dynamic: streaming video, QuickTime movies, virtual reality enactments, all pouring by at megabytes per millisecond. And what is true of the technology is equally true on the human side: eruptive start-ups, multi-mega-mergers, dot-com demise—a dizzying pace of change. For a revolution based on stability, the digital world sure moves fast.




Figure 1 — Pulse in an electronic circuit

Fixity *and* fluidity, in other words—digital dynamics crucially involves both. And both in ideal form. If you want stability, it will stay. If you prefer change, it will change—in exactly the ways you specify. *Perfect* dynamics—that is what powers the digital miracle.

3 Physical realisation

Of these three properties—perfection, abstraction, and dynamics—it is the perfection that is ultimately the most important.

It is odd, moreover—absolutely astounding, in fact—that such perfection is pragmatically achievable, in this our messy world. It is astounding because (this is our first crucial insight) nothing, in the end,

is *really* digital. Attach an oscilloscope to a digital circuit (figure 1), and all you see are splattered variations of bewildering complexity. Seriously: how long would it take a Martian to figure out that these intricate whiplashes of electronic alternation are, in fact, *digital*: “naught but os and 1s”? The discovery would merit a Nobel prize. Why? Because electrical signals, all those signals running around inside your personal computer, are not, *in fact*, digital. Rather: the parcels and patches of concrete reality that we call “digital,” like all patches of concrete reality, are really continuous.¹ And not just continuous, but, like everything that exists, perfused with an unutterable richness and texture and complexity of fine-structure that stupefyingly

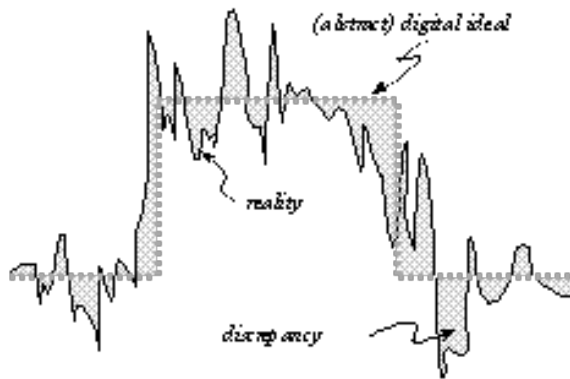


Figure 2 — The discrepancy from digitality

defies finite description. What makes them digital—or rather, more accurately, what *allows us to call them digital*—is that they are continuous patches that we can treat *as if* they were digital, without getting into trouble.

Or so the story goes.

3.a Discrepancy

Pure digitality is a myth—an abstraction in terms of which, with Orwellian abandon, we (re)interpret reality. As indicated by the dashed line in figure 2, the austere digital ideal is never achieved. Rather, real-

¹Ignoring quantum mechanics—which does not bear on current computers.

ity differs from the ideal by an unavoidable *discrepancy* (indicated in grey). Sometimes, as we will see—far more often than people realize—the discrepancy cannot avoid making an impact. To illustrate, I will presently argue that it is conceptually impossible for two pressings of the “same” audio CD to sound identical. But it is not the discrepancy that is mysterious. What is magic is that sometimes—in so-called “digital circuits”—*the discrepancy doesn't seem to matter*.

We are all familiar with computers that have gigabytes of memory, run at billions of instructions per second, and are linked to an untold myriad of other computers on the network. If every one of these systems—every bit, step, move, fragment—is “fallen,” failing to meet the inaccessible standard of digital perfection, why does the discrepancy not bring the whole thing down? Or to put it positively: how do we build perfection, on top of such inexorably messy foundations? How on earth does the digital idealisation work so well—how can it work at all—if reality is so unerringly defiled? And work well it does; those gigahertz processors and terabytes of memory really do achieve their digital goals.

The answer, or anyway the beginning of an answer, has to do with *containment*. Rather than eliminate discrepancy (a hopeless task), digital circuits *control* it. The genius of digital engineering involves figuring out ways to ensure that the *discrepancy does not propagate*. Whenever a signal gets (dare we say it?) *noisy*, we reshape it, clean it up, put it back on the strait and narrow—with Stalinesque efficiency. Memory on your laptop computer is “refreshed” 50 or 60 times a second, in order to stay stable. If it were not, then, like those Bengali manuscripts, it too would rot away, decay, collapse in frangible chaos. It takes work (and battery power) to prop up a digital myth—even to maintain the digital illusion of doing nothing at all.

What is stunning—and after thirty-five years in the field I am still amazed—is that we have figured out how to build devices to maintain the illusion—for a while. In the end, they, too, will fail (figure 3). Not even digitality can forever escape damnation by those deuced moths and rust. But this side of heaven, digitality comes as close to perfection as we can get. And we can get stunningly close—as close, in fact, as we want. Just tell the engineers what error rates you can accept: 1 in 10^7 ? 1 in 10^{20} ? 1 in 10^{25} ? Whatever you want; no problem.

3.b Discrepancy and noise

What about noise? Is discrepancy noise? Often—but not always. A small company in California listens to random cell phone calls, throwing away the signal. It turns out, for every handheld unit, that its discrepancy serves as a kind of analog “signature.” Given inevitable contingencies of manufacture and materials, each device is slightly different. Those differences are reflected in the particular shape of the error or discrepancy signal. (Tolstoy should be happy: we diverge in our own peculiar ways.) The company’s job is to monitor the character of the discrepancy, and sound an alarm when it inexplicably changes. That is how phone companies detect



Figure 3 — Demise of a CD

when a phone number has been stolen and implanted on a different unit. There is money in being a discrepancy sleuth.

Admittedly, discrepancy is sometimes painful. It degrades the music, distracts the image, crashes the machine. But it can be valuable. Cell-phone discrepancy is useful, because it *correlates*—with particular, concrete handsets. The same is true of pirated software: discrepancies on CD-ROMs contain tell-tale traces of illegal duplicating factories.

Does that mean noise is uncorrelation? No, that can’t be right, either. Sometimes lacking correlates is a priceless advantage: a childhood hideout, the Kohinoor Diamond, that night in Kathmandu.

3.c Abstraction

These remarks about discrepancy and error, singularity and correlation, tie directly into what I said earlier about the distinction between digital phenomena, such as chess and musical scores, and non-digital phenomena, such as billiards and paintings and musical performances.

The point is simple. *Digitality is not a property of entities per se.* Nothing either is, or is not, digital, intrinsically. Rather, whether something is digital or non-digital is *relative to a level of abstraction*—relative to a level of description at which it is characterized.

It follows that the perfection of digital systems, though *real*, is not *absolute*. Rather, *the perfection, too, is level-specific*—relative to a level of description or abstraction. The digital miracle, therefore, is not that perfection is achievable *at the physical level of abstraction at which the world is, in fact, messy*

(that genuinely *would* be impossible). Rather, the miracle is that messy physical stuff can be arranged so that, while staying physically messy, it can nevertheless *implement perfection at a higher, digital level.* (In terms of physics, this involves non-linear phenomena, attractors, and a host of other technical notions.)




Figure 4 — Implementing Digitality

Perfection, we might say, cannot be achieved *in* the physical realm; that would contravene friction, thermodynamics, and those moths and inexorable rust. Rather, as indicated in figure 4, digitality can be achieved *on top of* the physical world—by building it up, at a higher level of abstraction, on top of the underlying messiness.

This “level-specificity” of digitality’s perfection is going to matter a very great deal.

3.d Digital implementation

We finally have enough equipment to understand compact discs—to say nothing of Adobe Photoshop, digital cameras, and scanned paint-

ings. (Note: I will mostly talk here about music and CDs, because their traditional medium—sound—is a single-dimensional variable progressing through time, which makes for easier pictures. But the points I will make apply equally to two-dimensional static phenomena, such as pictures and paintings as traditionally conceived, as well as to two and three dimensional dynamic media and representations, such as video and virtual reality.)

The picture we have reached is one of a messy, continuous, underlying physical substrate, on top of which we can implement digital perfection. So far so

good. If the phenomenon we are interested in—chess, say, or written (at least printed) language, or musical scores—is *itself* digital (i.e., con-




Figure 5 — Interposing Digitality

stituted in terms of a finite set of black-and-white, “yes/no” distinctions), then we are essentially done. But what about phenomena that are *not* intrinsically discrete—such as musical performances, or paintings? How can we achieve perfection in their case?

We cannot. That is an intrinsic truth. But we can *fake* it.

A moment ago, I talked about implementing digital perfection on top of a messy, continuous physical substrate. For music, paintings, and other continuous phenomena, we can pull the inverse trick: recursively implement messy physical performances on top of a lower, *digital* level. That is, we can construct the three-level structure shown in figure 5.

What makes this all work—or at least what makes it work as well as it does or can work, which is something I am going to want to pursue in a moment—is our ability to make the implementing digital distinctions be *sufficiently fine-grained*, with respect to the upper-level continuous phenomena, that they are not noticeable, or anyway not unduly noticeable. This is called *sampling*, and is the fundamental strategy

behind CDs, digital cameras, bitmapped images, and the like. The details are familiar: you make a digital approximation to a continuous signal so that the *discrepancy*, as defined above, stays small. You can make it as small as you like, by using higher-and-higher sampling rates, more and more megabytes of storage.

Digitality is *continuously implemented*, in other words (i.e., is implemented on top of a continuous substrate), as we saw before; that is the relation between the lower two levels of the figure. Continuous artwork, in turn, is *digitally implemented*, in the sense of being implemented, in turn, on top of a digital substrate; that is the relation between the upper-two levels. When properly executed, the digitality of the middle level will be largely, or at least relatively, unnoticeable, at the top level. This is how the top-level phenomenon (the music, painting, image, whatever), unlike the chess position or poem, in spite of being implemented *on top of* digitality, need not *itself* be understood *as* digital. Imagine a Thelonus Monk CD: the growl, the bending of a note, a sigh. These phenomena are not *themselves* discrete. Their continuity is preserved, more or less, in spite of the digital *implementation*.

Why go to all this work? Why implement continuity on top of digitality, and then implement the digitality on top of more continuity? Because, by interposing perfection, between the bottom-level messy *realisation* and the top-level messy *phenomenon*, you can largely *insulate* the continuous richness of the upper level *phenomenon* from the continuous richness—which is to say, the moth and rust and disintegration—of the bottom level. The result is that the particularity of the “performance,” as it were, is insulated from the particularities of the recording—which in turn gives you extraordinary portability, stability, and immunity to decay.

Or anyway that is the theory. As we will see in a moment, reality is more complex. Still, this analysis answers our second question, by showing us the “why” of digitality: it is an engineering strategy, pure and simple, for insulating the continuity of one phenomenon from the continuity of another, in terms of which the former is carried. As John Haugeland has put it:

*“Digital, like accurate, economical, or heavy-duty, is a mundane engineering notion, root and branch. ... It only makes sense as a practical means to cope with the vagaries and vicissitudes, the noise and drift, of earthly existence.”*²

In an ultimate sense, as I have already suggested, it is the lower half of figure 5—the implementation of digital (i.e., perfect) systems on continuous substrates—that is the miracle. Why it is that the world is such that, as far as we can tell, digitality is the only way to achieve perfection, and why it is, correspondingly, that the world, at the messy physical level, is such that digital perfection can be achieved on top of it—these are the sorts of metaphysical question that keep me awake at night. They are questions that no one, I believe, has yet satisfactorily answered. But in terms of concrete, pragmatic impact on our lives—our third opening question—what matters is not so much digitality per se, and its possibility, as what we have revealed here: the consequences of *digital implementation*. It is the full three-layer structure of figure 5, not just the two-layer structure of figure 4, that is transforming the world of art.

Negroponte’s book was mistitled. It should not have been called *Being Digital*. It should have been called *Being Digitally Implemented*.

4 Mediation

Conceptually, most of the official story is in place. But it is instructive to pursue an example, to understand its profound limitations. Because the theory, as I have presented it so far, is not quite right. It is not *bad*; no one would turn down owning digitality’s patent. But as usual, the devil is in the details.

To see what’s wrong, I want to show how this analysis, if followed out with relentless logic, contravenes what is almost universally assumed: that all instances of the same digital “signal” are *absolutely*, not

²Haugeland, John, “Analog and Analog,” *Philosophical Topics* (Spring 1981); reprinted in J. I. Biro & Robert W. Shahan, (eds), *Mind, Brain, and Function: Essays in the Philosophy of Mind*, Norman, Oklahoma: University of Oklahoma Press (1982), pp. 213–225. Quote is from p. 217; emphases added.

just *relatively* identical. The example is taken from on-going debate in high-end audio circles about whether it is possible for two different pressings of the “same CD”—i.e., two different token polycarbide discs, each of which contains “exactly the same sequence of 0s and 1s,” to sound different.

According to the official story, they must sound the same. According to me—according to reality, that is, I will claim—they must not. And as usual, there is nothing special or peculiar about this result. The conclusion will hold of any digital implementation whatsoever.

4.a Different bit streams

As I’ve said, the debate concerns pressings of “exactly the same” CD, where the two tokens have “identical” bit streams. It should be noted, however, that this is a difficult case; there are huge issues, in the art world, about much easier cases, when the CDs—or digital representations in general—encode *different* bit streams. Here, what is important to realize is that *identity*, like digitality (and like just about everything else we are talking about), is level-specific. In cases of music and art, unlike texts and scores, where the “original” is continuous, identity at the top level need not correspond to identity at the middle level. Or so it is argued. And so people interpose digital “watermarks,” or lossy compression (e.g., to MP3), or digital stamping, and so on—claiming that they can do this without altering the upper level. I don’t want to consider these examples, here, except to say that such changes *have* to make *some* difference. If the upper level is *continuously identified*, at the top level (as in a performance, or painting), then any difference at the middle level is a real difference. What’s at stake is not whether there will be an upper-level difference (there will be), but whether that upper-level difference matters.

4.b The debate

But turn to the example at hand: of whether there are or even can be aural differences between pressings of the *same* CD—i.e., between two pressings that encode the very same, identical bit stream.

The debate takes predictable form. So-called “golden-eared audio-

philes” claim to hear differences between and among such different pressings. With great vehemence, self-styled “rationalists” deny the objectivity and validity of these golden-ears’ subjectivist claims. “It is *impossible* for the two pressing to sound different,” they cry. “They cannot sound different, because they are digital, and, as digital recordings, they are *identical*.”

The rationalists are wrong. To think that two pressings of the same identical bits *must* sound identical is simply a conceptual—perhaps ideological—mistake.³

To see why, we need to understand the impact, on the continuous (audio) signal at the top level, of what I called the “discrepancy,” at the middle level of our three-level diagram (figure 5), between concrete reality and digital abstraction. In particular, consider again Figure 2



Figure 6 — Sampling of a continuous original (i.e., top-to-middle layer discrepancy)

(page 7), showing the inevitable discrepancy between the real in-the-world voltage and the digital abstraction superimposed upon it. This image depicts a “digital” *signal*, such as a track on a CD, a bitmap downloaded from a digital camera, the output of a Photoshop session

³Similar disputes, I might note, arise throughout the community: about the rationality of colouring the edge of your CDs with green felt-tipped markers, of using isolation transformers on the CD player’s power cord, etc. My general reaction, in such debates, is not only to feel that the golden-eared audiophiles are phenomenologically correct, but also that their conclusions, far from involving anything mystical, are straightforwardly scientifically explicable, if only one’s conceptual analysis is sufficiently powerful. Colouring the edge of your CDs is perfectly rational, it turns out, if one is scientifically awake.

tweaking a digital scan. That is: it is a picture of (a piece of) reality *within the digital realm*. But the image is reminiscent of something more familiar: pictures, such as the one in figure 6, of how we *sample* or *scan* or *convert* continuous phenomena (paintings, music, 3D-scenes etc.) into digital form. In these input cases, too, as always, the digital abstraction (the “information that can be encoded”) diverges from the infinitely-rich concrete reality.

Every digital craftsman is taught this: that digital encoding—the so-called “analog to digital” (A-D) conversion process that takes place in every scanner, digital camera, DAT recorder, etc.—will invariably miss all sorts of small or minuscule variations, subtleties, and nuances in the original or “source” phenomenon. This is the realm of bit-depths, sampling rates, compression, etc. Suppose figure 6 depicts the (continuous) acoustic intensity of a live music performance. The areas marked with hash-marks are those aspects of the original that the digitisation process will fail to capture. By the same token, digital cameras and scanners analogously “abstract away” from: (i) any fine-grained structure of the original image or scene that is too small to be “caught” within the temporal or spatial sampling rate; and (ii) all variations in intensity that are less than one bit’s worth of gradation in the system’s dynamic range.

And what is true of input is equally true of output—though the point is less familiar. Suppose we print or render a bit map or other digital encoding: on a television screen or monitor, cheap ink-jet printer, or expensive imagesetter. Printers, monitors, etc., produce *real images*: concrete, continuous, full-blooded denizens of the world. These output images, too, being actual, will, like everything else, have an infinitely rich and detailed fine-structure. They, too, will look like figure 6. Or consider listening to a CD. Once again, the digital-to-analog (D-A) converter will take as input a digital signal, and produce as output something that is continuous, analog, and (as usual) infinitely-detailed. Being actual, these outputs, like everything else in the universe, will in fact have an infinitely rich fine-structure.

It follows, from all this, that if two images (sounds, whatever) are produced from one digital source, that they can potentially differ in some or even all of their fine-structure, in all their discrepancy. In fact they can—and will—differ in an infinity of ways, in spite of having

been produced from the same bit stream, because the “digital” bit-stream doesn’t determine that in one sense superfluous but in another sense absolutely necessary fine-structure. This is the point that underwrites Lowe’s work on *Digital Prints*. Starting with multiple copies of an “identical” bitmap, he printed eighteen high-quality prints, using eighteen different printing/rendering processes. They look different—radically different, even, when examined closely. And from what we have said we can easily see why. They look different because they differ *in their fine-structure*—in their *discrepancy from the* (common) *digital abstraction*.

4.c Fine-scale interactions

So prints, outputs, sounds, all differ. What does this have to do with different pressings? Because of this punch line:

The discrepancy intrinsic to the (continuous) physical realisation of a digital signal (i.e., the discrepancy endemic to the relation between the middle and bottom layers of figure 5) invariably influences the variation at the top layer of the resulting performance (i.e., the discrepancy endemic to the relation between the top and middle layers of the figure).

Not only *can* it have an influence; it *must* have an influence. It is a theorem of physics.




Figure 7a — How Digital-to-Analog Converters *Don't* Work

Why? Because, as we said at the outset, the “digital” signal is *not really real*. All that “really” exists is the underlying, physically messy

carrier. The so-called “digital signal” is only an idealizing abstraction.

The point is that D-A converters, the devices that *produce* a performance, given a digital signal, do not—and can not—work as indicated in figure 7a. This figure illustrates how people *think* things go—but it is a fantasy, based on the idea that the digital abstraction is real. Rather, the way they *really* work is indicated in figure 7b. They work this way because this is all that really exists.

It is perfectly obvious, in fact, that the fantasy could not be real. Just think of what it would require! It would mean that an engineer would have to build a concrete, physical device that (i) responded to




Figure 7b — How Digital-to-Analog Converters *Do* Work

the non-existent digital signal that the actual analog signal was ideally meant to encode, but (ii) that ignored the actual variation or “discrepancy” in the actual, real physical signal, which is the underlying physical realisation of that digital abstraction.

And that, needless to say, cannot be done.

We have talked of three kinds of fine-structure, each more detailed than is captured in any governing digital abstraction: (i) fine-structure in the original input, if there is one, before it is entered (converted, scanned, etc.) into a digital realm; (ii) discrepancies within the digital realm, in the fine structure of the signals that “carry” or “encode” the digital abstraction; and (iii) fine-structure in the output (prints, sounds, images) produced from those digital encodings. The striking fact is that it is a fundamental theorem of physics that these fine-structures not only will, but must, influence each other.

To see why, think about the encoding process. Suppose we start

with a routine continuous signal—an acoustic wave, painting or image, or 3D-scene. We've all been told, thousands of times, that analog-to-digital encoding processes take continuous signals as input, and produce digital signals as output. But do they *really* produce digital outputs? *No—of course not!* Digitality, as we've said, is an *abstraction*. Analog-to-digital converters, in contrast, are concrete: made out of physical stuff—the same stuff that we are made of, the same *continuous* stuff of which field-theoretic physics holds true. As a result, it would be contrary to the laws of physics for them to produce something abstract. Rather, like all physical processes, take in, *and produce*, concrete, continuous, signals or waves. More specifically, what analog-to-digital converter really does is: (i) take as input a continuous, concrete, real-world signal, and (ii) produce as output *another* continuous, concrete, real-world signal, where (iii) the output signal, *if interpreted under a digital abstraction*, can be seen to “encode” the digitised version of the input.

Analog-to-digital conversion processes, in other words, don't mediate between what is concrete and what is abstract. No real-world process could do that; it would be magic. Rather, A-D conversion is a *concrete-to-concrete transformation*, both ends of which are as a result genuinely continuous.

Exactly the same moral holds true, of course, at the output end: in processes of digital-to-analog conversion. Just as with A-Ds, D-A converters don't *really* take a digital signal as input, and produce a continuous one as output. Rather, they (i) take as input a continuous signal that supports a particular digital abstraction (i.e., lies within its acceptable constraints), but that, like all “digital” signals, is complete with discrepancies and fine-structure; and (ii) produce as output another continuous signal, the continuous signal which the digital abstraction of the input encodes. Just as in the input case, that is, output D-A conversion is a process of *concrete-to-concrete* mediation.

Once we have recognised the inalienable concreteness of the signals at both ends of A-Ds and D-As, it immediately becomes clear that it is an absolute necessity—a veritable theorem of physics—for the fine-structure (or discrepancy) in the inputs of A-Ds and D-As to have an affect on the fine-grained structure of the outputs. In particular: the fine-structure of the “digital” input to a D-A will effect on the fine-

structure of the continuous output. It *must* have such an impact, because (as usual) physical devices are continuous. Sure enough, engineers can strive mightily to minimize the effect. But there is no way, in this world we inhabit, for an engineer to build a concrete physical device that (i) responds to *the (non-existent) perfect digital idealisation* that a signal is “intended” to implement, but nevertheless (ii) to ignore the fine-structure of the incoming signal *as it actually is*.

This is why, ultimately, no two pressings of the same CD will (or even could) sound exactly the same. Or rather, to put the point more exactly: this is why no two pressings of *the very same (digital) bit stream* will ever lead to *exactly the same (continuous) acoustic wave*. They will sound different even if we assume, for simplicity, that they are played on the same stereo system, in identical states. They will sound different because, although each CD will carry the same digital idealization, each will do so complete with its own unique fine-structure—i.e., with its own distinctive way of diverging from the putative digital ideal. After all, the fact that they are the “same” CD means no more than this: that if we were to abstract away from the two infinitely-rich continuous patterns, in the way mandated by the digital idealisation, the two pressings would be discovered to “carry” the same stream of digital bits. Any aspect or fine-grainedness of structure that is not relevant to this digital abstraction is free to differ between the two CDs. And as we have already seen, the D-A converter is mandated by the laws of physics to respond differentially, in the two cases, *to those different fine-structures*. Perhaps not *very* differently; but nevertheless *some* differently. So when the continuous signal is extracted from the D-A, sent to the power amplifiers, and propagated to the speakers, it will carry its own distinctive characteristic signature. There is no way in which it could be any other way.

And finally, to bring this back to images, the same holds true of printing. It is not just that two printings of the same bit-stream (even: of the very same CD encoding that bit stream) can produce different concrete images, when printed on different printing devices, as Lowe showed so compellingly. It is also that two different CDs of that “same” bit stream, when printed on the *same* printer, will *also* produce different prints. This will be true independent of how the digital bit stream was

produced: entirely within the digital realm (Photoshop or painting programs), or scanned or sampled from a continuous original.

5 Conclusion

What have we learned?

Six things, already. And a seventh lies just below the surface, with which I will conclude.

First, digitality is not an intrinsic property of anything. Whether or not something is digital is a higher-order characterization of it: a characterization of a characterization. It is *characterizations* of objects, “takes” or cuts on objects, that are, or are not, digital—not objects per se. As a poem or score, a text may be digital, even if as an arrangement of ink, it is not. As a CD, a recording may be digital, even if, as a reflector of laser light, it is not.

Second, some objects—such as musical scores—are digital at the level at which we identify them as the sorts of thing that they are. That is why we say that a score can be “perfectly” copied, or think that we know *exactly* what sonnet Shakespeare wrote. It is not that we (or anyway most of us) *have* the fully-concrete sonnet that issued from his pen. Rather, what society or culture or history has settled on, about sonnets and scores, is that what constitutes their identity, as the kind of object that they are, is their characterisation under a given set of descriptors or types, which can be exhaustively specified in terms of a finite set of “yes/no” decisions.

Third, in spite of this identification of some things (such as scores) as constitutively digital, nothing *actually* is—or anyway, nothing concrete, nothing actual. The physical world is messy, and so any material thing, as a material thing, is, far from being perfect, in fact a messy, decaying, piece of stuff.

Fourth, phenomena that are not only (of course) not digital per se (nothing is), and that are also not digital as physical entities (as we have just seen that nothing is that, either), and that are not digital at the level at which we take them to be constituted—such as paintings and musical performances—can be *digitally implemented*, at some loss, but with the benefit that one thereby largely insulates their high-level con-

tinuity from the low-level continuity of the substrate—paving the way for extraordinary longevity, transportability, reconfiguration, modification, etc. This is the realm of the digital CD and the digital image; a digital implementation of a continuously-constituted phenomenon.

Fifth—in spite of the undeniable success of this three-level strategy (of digitally implementing continuous phenomena), the underlying discrepancy is never avoided entirely. Because, as we have seen, physics is continuous, the discrepancies from the ideal in the lower level of implementation (digital on top of a messy substrate) can never be wholly isolated from the discrepancies at the upper level (the loss or violence to the continuous upper-level phenomenon that comes from digitally sampling or representing or encoding it).

Furthermore, it is a theorem of physics that this unavoidable underlying discrepancy will always influence the output.

Sixth, because nothing physical is in fact digital, and because, as we have seen, the discrepancies can never be entirely removed, it follows that *digitality itself is an abstraction*. This is the reason why, even though we *say* that a digital implementation insulates the continuity of the constituted phenomenon from the messy continuity of the implementing substrate—to say nothing of the moth and rust—it is never actually so. Sure enough, as I have just said, we can go to a lot of work to minimize the impact of the inexorable discrepancy (different pressings of the same CD can be arranged to sound *pretty much* alike). But the metaphysical truth remains: digitality is not only a property of abstraction; it itself is an abstraction. When we say, of an abstraction, that it is an abstraction—for example, when we say of it that it is a digital abstraction, as for example in the case of musical scores—we are (recursively) engaging in a higher-level abstraction of our own.

Seventh and finally, what goes around comes around—one more final time. Even the idea that we are *abstracting* is an abstraction. The whole edifice of “levels of description” is a *way of describing* what we do. It is a cut, a take, on our epistemic practices.

Nothing that is actual, actually abstracts. Rather, for us to say that something abstracts—a recorder or scanner that performs an analog to digital abstraction, say, or a printer or amplifier that performs a digital

to analog abstraction—is an abstraction of ours, which, *as* an abstraction, like all abstractions, under-describes what is going on. To say that something abstracts is to do an injustice to it. To make a claim about an *actual* process is to commit oneself to an abstract characterisation of a (concrete) process—a process that, like all physical processes, mediates between *one thing that is concrete and something else that is also concrete* (or perhaps we should say, more carefully: between one thing in its full concreteness and something else in *its* full concreteness).

Put it this way: it is not just that *digitality* is an abstraction. Or even, though this is also true, that the *perfection* of digitality is also an abstraction. The bottom line is that *abstraction* is an abstraction, of which digitality is an instance. As I have said, nothing that is actual, actually abstracts. We might as well get used to it. The world is utterly and inexorably concrete.

— end of file — 🍏🍏