

SpamStream

¹ I know the style and tone of these shows were (are?) totally different, but I don't have a clear memory of whether it was needlessly sensational, like *Inside Edition*, or sort of a folksy-local, avuncular telecast, as I remember *Chronicle* being.

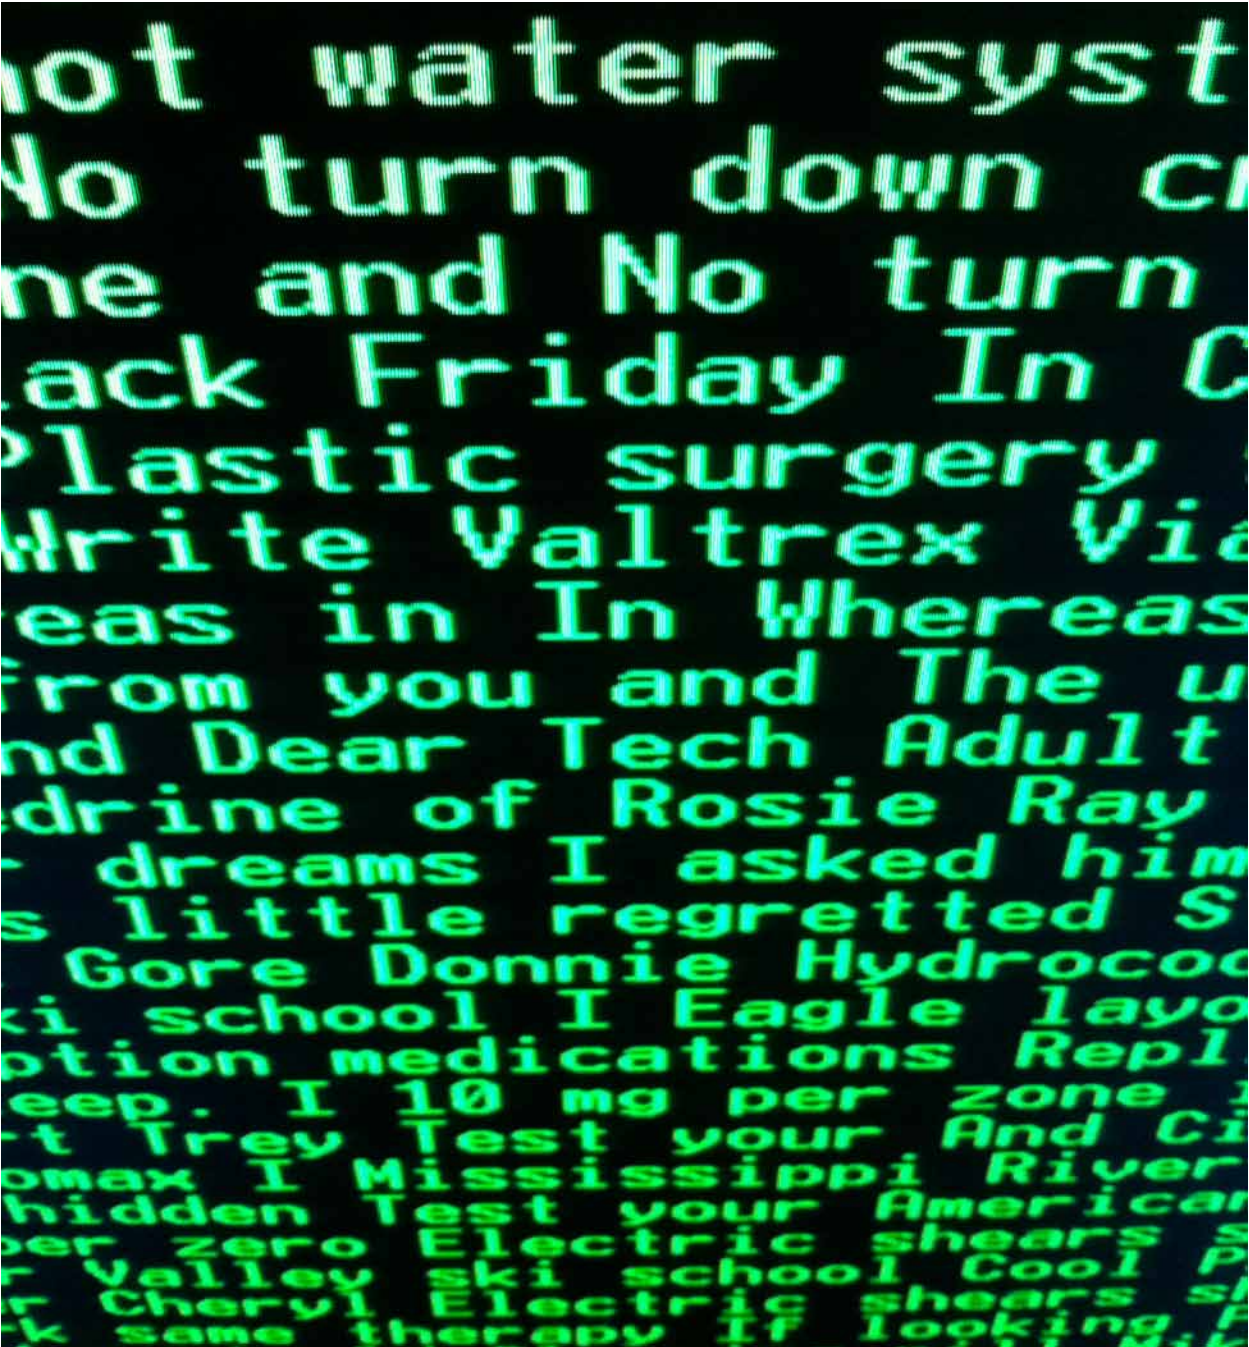
² *Bottle Rocket*-through-*The Royal Tenenbaums* Wes Anderson, before he grew his hair out and started wearing contact lenses.



I remember seeing a news story on some sort of *Inside Edition*-type show (or maybe it was *Chronicle*¹), years ago, like pre-millennium days, where the anchorperson said something like, “Up next: Do you know what a ‘blog’ is? Well, find out, after the break.” It was a story about how people were starting these things called, yes, “blogs” and how it was short for “web log” (kind of). They interviewed some nerdy dude who looked like Wes Anderson², who rattled on about the wave of the future and personal datastreams and cyberspace, and when they cut back to the anchorpeople, they seemed (rightly) confused, and sort of shrugged their shoulders before moving onto the next story about a celebrity tribute or local dairy farm, again, depending on which show actually broadcast this segment. I remember being baffled, too, even though I had used a computer and early, dial-up-driven forms of the internet at that point. One thing was clear: based on the nerdy guy, I didn’t want to have anything to do with blogs.

Over the next fifteen years or so, as blogs became something that people actually knew about and used, I maintained my distance. The diaristic nature of the content and purpose of most of them seemed self-indulgent and boring. And that image of the nerdy guy stuck with me. So, when we were asked to start a blog to document our thesis development process in *Design Seminar II*, I was skeptical and did probably a little bit of the anchorperson -shoulder-shrugging: Ok. I’ll do it. I guess.

The basic communication concept of the blog is that it provides a space for people to write and post whatever they want, and for other people, if



3

Bathroom tiling, driving, and *Breaking Bad* were some topics covered, to give you an idea of how much I fell into the “blogosphere” mentality of “everything I have to say is interesting.” Has there ever been a lamer word than blogosphere?

4

Any self-respecting blogger will have a filter set up to prevent any derogatory comments from being automatically posted to their site.

interested, to access this information. One distinction between a “normal” web site and a blog is that the latter usually includes a built-in feature that allows visitors to comment on posts/entries. For example, someone might have a portfolio site containing their work that people could view without having the option to comment and keep their opinions to themselves since very few of these sites would have a section called “Tell Me What You Think of My Work”. A blog has more of a “I’m interested in what you have to say about what I have to say” vibe to it. Thus, when I started begrudgingly posting entries on my blog about my thesis progress, or whatever else I was writing about³, part of me was curious if anyone would comment on what I had to say. When I started posting songs that I had recorded as part of my thesis prep work, I instinctually started checking the “Comments” section of the *Wordpress* backend interface to see if anyone had written to me, praising my totally awesome take on an Afghan Whigs song. The very structure of the blog elicits a digital sort of needy vulnerability; if people have the ability to tell me that I’m great, why aren’t they telling me I’m great?

As my daily and increasingly obsessive comment-checking continued, I began to encounter some, ahem, weird shit. I found comments attempting to be posted to my page⁴, not by admirers overflowing with praise of my musical acumen or my pithy remarks on dynamic media, but by computers acting as blog readers. I had comments like “Great post. Really interested in what you had to say.” And then the name of the post would be <http://viagrulzforcheap.ev/3423423442/534544>. Or some would say “You make a good point, but I think you should read this:”, which would be followed by a link to a body-building supplement or a torrent site, or whatever. This kind of spam directly taps into the narcissistic blogger mentality by pretending to be interested in what the blogger has to say, and by using this ruse to get his or her attention long enough to click on a link in order to generate web traffic and revenue. It’s using the system somewhat ingeniously vs. regular email spam, which is a direct and unsolicited attack on one’s inbox, and the majority of which most likely gets deleted without being opened. And it’s all based on co-opting the idea that people are putting their writing or pictures or opinions out there for the world to see, and would probably really





like it if someone told them how awesome their ideas were⁵.

Amongst the pseudo-praises and mild dismissals of the spam-bots, there existed another kind of comment, also spam, but fundamentally different in its form. Rather than pretending to be a person commenting on a post and masking URLs in usernames or as hidden links, these comments were simply collections of keywords with embedded links, presumably to online stores or shady sites; I'm not one to click on a "hot moms" link. You know, usually. As I waded through these keywords, I noticed combinations of words and phrases that seemed strange and incongruous on their own and sort of pretty when juxtaposed with each other. *Chevrolet super sport nova. Karma lyric. Chameleon club culture. Coruna killing. Civil war battle front royale. The indicated nymphs.* And so forth.

Whereas the other comments attempted to simulate human interaction in order to get people to let their guard down enough to read the message and possibly click on a link, these keyword-only messages seemed to side-step that tactic entirely, opting for a pure language approach: dangle these words in front of the blogger and let them try not to click on them. The poetic (as I saw it) juxtaposition of the words seemed to be unintentional.

Up until this point in the DMI program, I had avoided data-visualization projects. Part of the reason was that I hadn't ever found any kind of data source or collection that was meaningful enough to me to elicit the considerable amount of enthusiasm that necessitates the evolution from flight of fancy to something becoming a tangible project idea. I think it's also hard to make a good data visualization project, especially one that does something other than give you an initial, "Oh, that looks interesting" followed by the subsequent realization that no, in fact, the world does not stop because someone made a poster of the political tide of Tunisia that features different colored icons—with different scales!—or a web site that allows users to see numbers turn into rainbow lasers and to sort those lasers by category⁶. The blog spam content finally gave me a datastream⁷ that I found compelling enough to visualize, or at least compelling enough to investigate its potential visualization.

But so how to visualize the spam? I basically had a large collection of words. How should the words be treated? Should each word be its own

⁵ This might seem like a stretch, but it reminds me of the "depth approach" that advertising agencies began deploying in the 1940s and 50s to target consumers. Basically, rather than trying to sell people a product, advertisers began to look at the reasons that people actually bought a product, which often weren't obvious or rational and rarely had to do with the product itself. For example, toothpaste companies originally thought that their customers bought their product and brushed their teeth to promote oral health. That's fine; a good amount of people do. Eventually, they realized, according to their research, that the majority of people brush their teeth specifically because they want their mouths to feel clean. So they started marketing toothpaste as having a "clean mouth taste"^a. The point being that if you want to reach people and possibly get them to buy something, there's an initial level of their consciousness that you have to get past in order to tap into what they really desire. I think the blog spam does this by piquing the curiosity of the blogger just enough to get them to open the comment and make sure that it's not someone telling them how bitchin' they are.

^a Vance Packard, *The Hidden Persuaders*. New York: Ig Publishing, 1957. pg. 40.

⁶ Though, shit, that might be cool.

⁷ Goddamn you, *Inside Edition* or *Chronicle*-featured nerd guy.

8
For images and project descriptions, check out <http://sosolimited.com>.

9
<http://davidsmall.com>

10
I'm not the world's biggest poetry fan, but when I do read it occasionally, I prefer to read it in the privacy of my own home, devoid of the kinds of people that hang around at poetry readings, if you get my drift. (Yes, this was meant to be spoken in the voice of Dos Equis' *Most Interesting Man in the World*.)

11
I've known him since freshman year of high school, have played in a band with him for ten years, have written screenplays with him, and we were each other's Best Men at our respective weddings. The point is: I trust the guy.

12
When you say it this way, doesn't this sound like it should be some kind of quasi-news/ interview show on like MSNBC? :
"And now it's time for *Design Studio II with Brian Lucid*. Here's your host...
BRIAN LUCID!!!"
(applause)
Brian: "Thank you, Kent. Today we're going to talk about information hierarchies..."

13
Ok, let's just say no.



discrete element, or should they maintain some of the juxtapositions that made them compelling to me in the first place? Some quick research into data visualizations of word systems and text chunks, just even locally, led me to stuff like Sosolimited's *The Long Conversation*, *Reconstitution*, and *Evening News Remix*⁸, as well as any of interactive media saint David Small's⁹ work that uses text and typography as a primary visual element. These projects are all done well, and, well, have been done. So I quickly abandoned the idea of a typography-based project. Also, I wasn't interested in the frequency of use of the words in the comments, or their parts of speech, or their etymological origins. Mainly what struck me was the *sound* of the words in my brain-voice as I read and sifted through them. The word "poetic" kept lingering in my brain. Poets want their work to be read aloud; that's where it lives and exists¹⁰. I decided to give a selection of blog spam to my friend and frequent collaborator, Brian "Hoss" Coughlin, and have him record himself reading the words and phrases. I didn't give him instructions on tone of voice or articulation; I wanted to just leave it up to him¹¹. And so he did. I cut the audio into small pieces containing a single word or phrase. I now had the spam both as text and as audio. Ok, so what now?

In *Design Studio II* with Brian Lucid¹², we had the assignment of designing a holistic encyclopedia of the arts. We had to sort of curate whatever collection we wanted in the encyclopedia and design an interface to access the collection we chose. Trying to narrow the scope of my particular collection, I focused my encyclopedia specifically on contemporary art and artists that I know about and like. I attempted to design an interface using abstract forms to signify different moderate to high-brow art concepts and themes, like, say, the occult, transfiguration, immortality, deconstruction. You know, easy things to represent in abstract color forms. The viewer/user would be able to manipulate these shapes by combining them into new forms in order to see overlapping patterns in the work on a purely conceptual and formal level. Contemporary art is mainly about concept, so in some ways an interface that stressed conceptual characteristics seemed to be the most appropriate for the subject matter. Whether or not I succeeded in creating a successful interface is arguable¹³, but the project was

significant in that it raised the personal specter of representing complex ideas concepts using a visual "language" *per se*.

So, my first idea for the visual component of what I was now calling SpamStream was to create an algorithmically-generated abstract visual language using Processing, that would essentially take Hoss' audio files as an input, do some dynamic-media-*Yahtzee*-style shaking up of the dice, and roll out a visual form. This had also already been done with spam, to an extent, with Alex Dragulescu's *Spam Architecture* series. He created a system that uses spam as an input. Then, "various patterns, keywords and rhythms found in the text are translated into three dimensional modeling gestures."¹⁴ The result of these gestures are abstract forms that (to me) have basically no decipherable connection to the spam content itself. The spam content is first translated into numbers, which then makes the subsequent translation into visual form somewhat arbitrary. He could have been using traffic pattern data from Mumbai; it didn't matter. The output was indecipherable as derived from its source. So, my goal was to maintain some kind of conceptual connection between the source material and the final visual output.

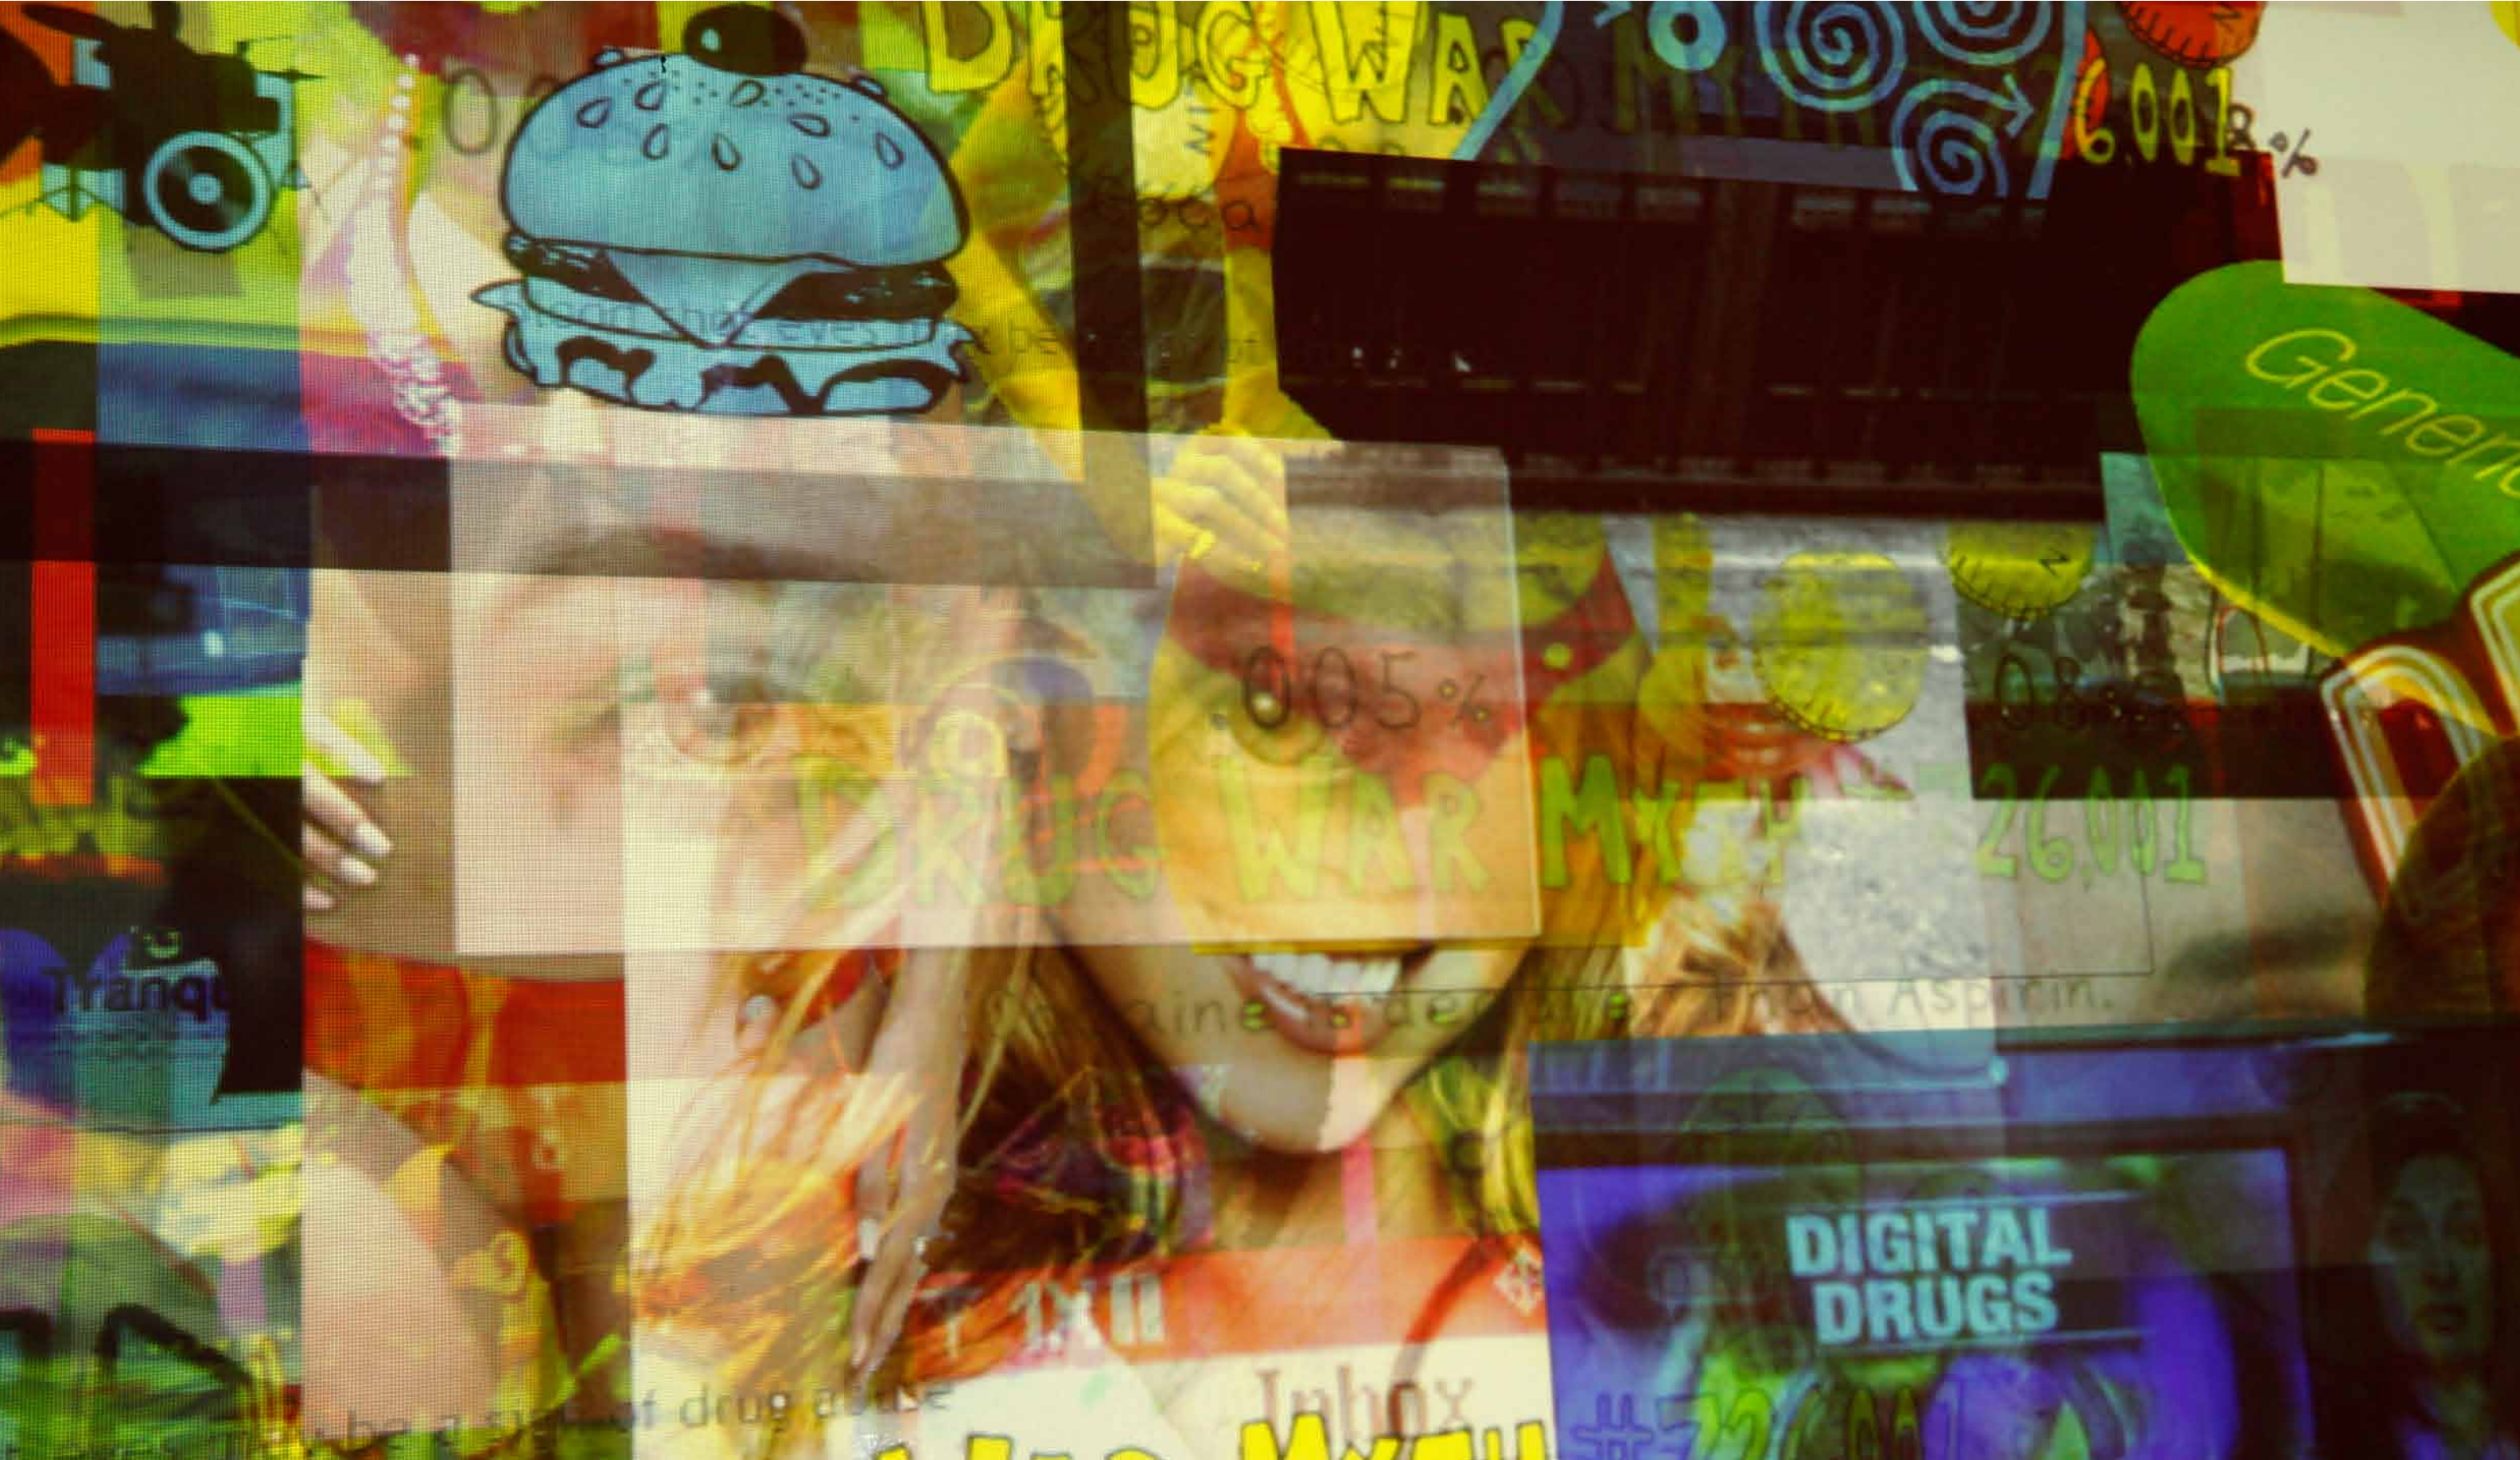
Having also eschewed learning *Processing* up until this point in my DMI journey, I realized that this algorithmic direction would be a dead end without someone to lead me down the Java-laden path. I asked my sister's boyfriend, Matt Quesada, if he'd be interested in collaborating on the project. Matt is a software engineer for a company called *EScription*. He had no prior experience using *Processing*, but had worked with other code languages before. He suggested that we use *Python*¹⁵. Also, because his company does a lot of work involving speech recognition software¹⁶, he thought it might be cool to incorporate that technology into the project. During our first working session, he loaded a single image into a window. I asked him what sort of parameters or filters *Python* allowed for image manipulation. He said "Um. Well, I loaded the image." There was silence.

By the end of the session, we had started dabbling in *Processing*. Based on the limitations of our collective knowledge, I couldn't imagine how we were going to make any kind of interesting abstract/quasi figurative representations of the audio files. I had envisioned Jeremy Blake-esque floating

14
Mr. Dragulescu, or one of his minions (he's now at the MIT Media Lab), quoted in Geoff Manaugh's *The Bldg Blog Book*, San Francisco: Chronicle Books, 2009. pg. 14.

15
A Java-based programming language that is more of a nuts-and-bolts language than something like *Actionscript 3* or *Processing*, both of which are languages created specifically for designers. For example, in *Python*, you'd have to write a function manually to, say, load an image, whereas in *AS3* or *Processing*, you could simply call the "image" function, which has already been created and embedded into the program, so that you, Mr. or Ms. Graphic Designer Dabbling in Code don't have to waste your time getting your hands dirty with "real" code. For *SpamStream*, Matt came into the project having worked only with the nuts-and-bolts-level stuff, and I'd had an only mild level of experience with the "designer" languages. His lack of awareness of data visualization projects using *Processing* and *AS3* was one of the major reasons I wanted to work with him.

16
Their company was just purchased by *Nuance*, the company responsible for *Dragon* and all of its shockingly annoying radio commercials.



0.05%
DRUG WAR MYTH 26,001

...aine is der... than Aspirin.

DIGITAL
DRUGS

Tranqu

urban... of drug abuse

Inbox



MCLE
CREDIT
AVAILABLE*

50%
discount

day!



WHO NEEDS DRUGS?

NO SERIOUSLY, I HAVE DRUGS

BY C. J. RAY



17
SpamStream Fun Fact: Approximately 35% of the images that come up when one searches for “Sex” on *Google* are *Sex in the City*-related.

18
In Spring 2010’s *Design for Motion and Sound*, in the context of a group project^a, we were given the words “desire”, “interrupt” and “collide” and told to make a story, or a “narrative”^b. This story had to be told in three different ways, or, more accurately, via three different media forms: text, sound, and image. At some point, as part of the assignment, my group’s words became “desire”, “love”, and “collide”. These words, for me, conjured up classic film noir posters of the 30s and 40s. Long story short, we ended up making a translation system that used text as an input and had a dynamic collage of film noir imagery as the output. This, however, was a *linear* system, in that you put something in at one end, then saw the end result on the other. You sort of had to remember what you put in. With *SpamStream*, you would be experiencing the translation simultaneously and continuously. The juxtaposition of the sound and image and text on the speech-recognition screen would allow for multiple levels of real-time affinities to be made between the discrete components of the project.

a
I always thought that teachers assigned group projects to encourage collaboration. Now, after years of disliking them (the group projects), and in the midst of what is most likely my last semester of education—ever—I think they assign them so they’ll have to sit through less critiques. They tell me I’m an optimist.

b
Just, yikes.

digital paintings, with each brush stroke or canvas wash being algorithmically driven by the amplitude and tonality of Hoss’ voice. Instead, we were able to load one image at a time. No audio. And, Christ, it had taken us a while to figure out how to do *that*.

Knowing now our technical limitations, I started thinking about images. Specifically web images. What if the audio files could generate web imagery based on content? I created a spreadsheet and classified all of the words and phrases into seven categories: Advertising, Sex, Violence, Money, Entertainment, Religion, and Drugs. These categories seemed to comprise the majority of the content that I had. I’ve since somewhat-cheesily dubbed them the Seven Deadly Sins of Spam. But, it’s true; the majority of stuff out there fits into these categories. And, I loved how some of the words and phrases could fit into multiple categories. The ambiguity/double meaning of a phrase like “online caribbean stud”, tagged to “money”, “entertainment” and “sex”¹⁷, seemed like it could be a project in and of itself. And the idea that the computer would be hearing Hoss’ voice saying these things, interpreting that and then displaying its translation seemed fascinating.¹⁸

So, I thought it made sense to have the audio files tagged as belonging to one or more of the categories. The system would identify the file and pick one of the categories at a time, then search *Google* or *Flickr* (or whatever), for that category, grab an image, and then display it. Each category would also have its own method of display (fast or instant appearance for Violence, long slow fade for Religion, etc.). As Hoss’ audio files played in a random sequence, a layered collage of internet imagery would be created and projected onto a wall—a combination of spam-bot poetry and web pixels, never the same composition twice—and a monitor would display what it heard Hoss’ voice saying via speech-recognition software. The effect would be a sort of multi-layered feedback loop, a constant back and forth between human and computer interpretations.

We started coding the project with this general blueprint in mind. I went down to Roslindale every Sunday afternoon and we sat in front of Matt’s 50-inch flatscreen television with his PC hooked up to it. During these sessions, I tried to take all of the thoughts and ideas and whims I was having about blog spam and imagery and translation of media¹⁹ and

whittle them down into useful, palatable chunks for Matt to process into code. I could rant and rave about this stuff for extended periods which was probably somewhere between insufferable and not helpful. I also tried to be patient and take one section of the project at a time. It’s hard not to blurt out stuff like “What about the audio?!? What about the speech recognition?!?” when you’re trying to simply get one image to stay on the “stage” while another loads in, while weird shit keeps happening for reasons that are frustratingly unknowable, and mainly you just want to have a beer because it’s getting on five in the afternoon on a Sunday already. But that’s how we worked; we slowly went through the possibilities of loading the images into the system, exploring different filters that *Processing* offers and the various ways of faking and making a time-based collage²⁰. We focused on creating captivating visual treatments of the images that would later make the experience somewhat interesting, if not meaningful, when added to the context of the audio files.

We encountered a fair amount of technical problems (as is the case with pretty much any foray into the actual coding of a project versus keeping it encased in conceptual lucite), some as a result of the basic limitations of *Processing* and some probably as a result of our incompetence/naiveté in coding. Some problems were frustrating and limiting and just kind of sad²¹, but became obstructions that had a sort of Leth/von Trierian effect on the project. A major concern from the outset was whether the viewer/perceiver would be able to make any sort of connection between the random words and phrases that he or she would be hearing and the random images being projected on the wall. A one-to-one correspondence would be too concrete and dumb; too much randomness would be too abstract. I mean, you have to give the viewer some context from which to derive meaning.

The initial idea, as detailed above, was that a phrase would be pulled from the database, the system would pull a category from that phrase, and an image would be displayed based on the attributes assigned to that category. The first image would be a “background” image that would serve as the canvas on which the subsequent images would be overlaid. Then the next phrase would be pulled, a category would be pulled, an image displayed. And so on. However, when we tried to code it this way, we couldn’t get the

19
When an idea takes hold and you’re really in it, when you try to tell someone about it, it can be like you’re either basically insane, or someone’s grey sheep uncle who goes heavy on the nachos at Thanksgiving and has little bits of tortilla chips around his mouth while telling you about something called “future shares in vertigo theory” that he wants you not to invest in, like he has, but just to “chew on it for a while.”

20
Whereas *Flash* has the idea of a “movie” and a timeline inherent in the program, *Processing* requires the use of timers and other jerry-rigs to create something that unfolds over time.

21
Early on, when we couldn’t get two images to appear at the same time, Matt and I discussed having me ask Brian Lucid (who was advising me on the project in a class called either *Thesis Exploration* or *Thesis Explorations* depending on which document or web page you find yourself perusing) about it, before we both decided that this was so rudimentary a question that we would both be too spiritually embarrassed to actually ask it, both me, in person, and Matt, safely faceless and tangentially involved in the project.

22

This is code/nerd speak for “list”.

23

After initially trying to use a *Google* or *Flickr* API to pull the images dynamically into the project, I decided to simply download 100 images per category by searching *Google Image Search* for the category names (“Sex” [more tame than you’d think, as long as your *SafeSearch* is on “moderate”], “Religion”, “Money”, etc.). This added another level of computer/human interchange, as I was the one picking the images from search results that had been created both by system algorithm and tags that people had affixed to the images.

program to apply a different category on an image-to-image basis; there was some problem with the array²² that prevented us from doing it. We could, however, keep the same category for a series of words/images. Thus, if the first word/phrase had the “Drugs” category, then the next phrase/images would also be from the “Drugs” category. This limitation, after initially leading to a near-meltdown by yours truly who had to be assuaged by an earlier-than-normal Sunday afternoon IPA dosage, potentially led to a more meaningful experience for the viewer; there was already enough randomness in the project, from the origins of the words and phrases and images²³ to the order in which they were displayed, to the position where they were displayed (random x and y coordinates on the “stage”), to the characteristic way in which they were displayed via filters (a randomly chosen filter from a set of ten). A sequence could now live within one category, and the category could give the sequence the framework for context and meaning; random material isn’t random when viewed through a common filter. In fact, the categorical framework provides a context from which meaning can be applied by the viewer e.g., why is there a woman wearing a bikini followed by a syringe followed by a smiling silver fox-type character? Oh, we’re in the Drugs category. Hmmm.

Having been interested in the nature of repeated viewing and perception intrinsic to looped media, and the idea of the loop allowing repeated viewing and perception, and the possibility of unlocking layers of meaning based on these repeated viewings, I envisioned *SpamStream* as functioning on its own internal clock, with the viewer being simply that: the viewer; not an interactor or user, but simply one who beholds the collages and audio and who derives or doesn’t derive meaning from the experience. This seemed maybe a bit lazy and kind of lame, I’ll admit, but I liked the idea of this system functioning autonomously, with a person being someone who sort of just pops in to take a quick look-see, while the system continues to churn on and on at all hours of the night on its own digital cogs and gears. During our programming sessions, Matt and I would use a “mouse click” command to forward the collage, rather than waiting for seven seconds (or whatever time period we would decide worked) for each image to appear. We could step through the images and the audio files quicker this way to



24
Think Stay-Puft Marshmallow Man in
Ghost Busters.



see if the filters that we created were functioning. When the programming phase for the first functional version of the project was close to complete, we ended up deciding to keep the mouse click as the way of forwarding the collage. By offering the user a button, basically the most simple of interfaces, he or she could control the rate at which the audio files are played, and, subsequently, the rate at which the collage is created. We could then add parameters limiting the amount of audio files/images the system would run before the collage restarted by wiping the stage clean. But the speed of the experience would literally be in the hand (for lack of a sexier word) of the user. Pressing the button would trigger the playback of an audio file and the appearance of an image, enacting a direct physical correspondence between the user and the installation.

We started thinking about the simplest button possible. A mouse would be too straightforward and familiar. In a Ray Stantz-esque²⁴ epiphany, we thought of the Staples' *Easy Button*. A big red button that served one purpose. No right clicks or scrolling. One function. Push me.

I went to Staples and bought an *Easy Button* for \$4.99. I took the thing apart with a 1.4mm screwdriver. Matt, having a background as an electrical engineer, modified it to function as a wireless controller by disassembling a wireless mouse and installing the hardware within the casing of the button. Voilà. *SpamStream* was now interactive.

We eventually got the project working and stable enough to include in the student show *MediaLuscious*, which was held in the Pozen Center at MassArt in March 2011. We projected the collage onto a large wall. To the left of the projection was a 50" plasma screen that displayed the speech recognition software output of the audio files in real time. A podium about fifteen feet away from the wall held the modified "Easy" button (now the "Push Me" button). Two speakers below the collage played the audio.

Over the course of the night, people stopped by and checked out the project. Some would give the button a single press, look for a connection on the projection, and then move on. Some people stayed for longer periods of time, watching the collages build and change as layer upon layer was added to the composition. Some people glanced at the wall and walked away. One woman, probably in her late fifties, pulled me aside, congratulated me, and

told me that she "got what I was doing." Matt and I had an extended conversation with Dana Mozer²⁵ about the technical details of the project. He thought that Hoss' voice was computer-generated, which was interesting, and raised the idea of the back-and-forth between human and computer from a different perspective. Fred Wolfink watched the project comfortably for an extended period of time.

It's hard to know whether people valued, or how much they valued, the interaction component of the project. It was kept as a simple button push, as detailed above, to shift the focus onto the collages, which, with their visual beauty and depth, were the aesthetic strength of the project. If the end result is captivating and draws people in and makes them consider the connections between the media, maybe that's enough.

Ultimately, *SpamStream* is about the connection between words and phrases originating somewhere in the cracks of the internet, filtered somewhat through human consciousness via their being spoken aloud and then linked to images that are tied to those words and phrases via concepts and keywords. It raises the question of what images can mean in different contexts, and how language can influence those contexts as both hierarchical information and specific description. The interaction and idea exchange between the human and the computer is becoming increasingly inescapable and decidedly commonplace for the majority of people in cultures both Western and beyond. The meaning that we apply to this interaction doesn't exist in URLs or passwords or virtual shopping carts or blogs, but in the fringes of our experience and the repetition of simple actions. We search for things—for concepts, for keywords, for answers—and the images we encounter become part of our screen memories. Embedded. Meta-data. Forever.

25
From my understanding, one of the head
honchos of the SIM department.

