

Bi-weekly Random Bits from the Internet

2015-09-07

(NAH, SANTA BARBARA IS OKAY.)

Eileen Myles, The Art of Poetry No. 99

P2, Ben Lerner, The Paris Review, August 31, 2015

Call me Maybe: MariaDB Galera Cluster

P27, Aphyr, September 1, 2015

A Neural Algorithm of Artistic Style

P39, Leon A. Gatys, Alexander S. Ecker, and Matthias Bethge, ArXiv, September 2, 2015

Eileen Myles, The Art of Poetry No. 99

Ben Lerner, *The Paris Review*, August 31, 2015

Eileen Myles has lived in her East Village apartment, where this interview took place, since 1977, and yet, entering her studio, I got the feeling she'd just moved in or was ready to move out — both make sense, because she is itinerant (since our conversations last summer, she's been in Dublin and Lisbon and Oakland and Paris and Provincetown, Massachusetts, and several other places) and because she and her work are unsettled in the best sense: restless, disturbing, changeable. She has no imitable manner, no manners. She has an interesting art collection, but no clutter. Her only built-in furniture — a sort of combination bed-desk-bookshelf — was constructed by her painter friend Philip Shinnick while he was dog sitting in 1995 (Myles was in Russia). Otherwise, she has her brilliance and her stabilized rent.

The author of nineteen books of poetry and prose, Myles is often referred to as an “institution” — the way one speaks of a terrific restaurant that's endured the waves of gentrification as a “New York institution.” But the word bounces off her: there is nothing official about her, nothing staid or still. She is exemplary for more and more young writers precisely because she has gone her own way.

During our first talk, we ate scrambled eggs in front of an ineffectual fan. On subsequent visits, I brought iced coffee and stone fruit. There was no small talk and no dead air. — Ben Lerner

INTERVIEWER

The cliché is to end an interview with a question like, What's next for Eileen Myles? I thought we could start with that. You're working on a book about dogs, aren't you?

MYLES

A particular dog. It's about the first dog of my life — a pit bull named Rosie I got in 1990. She was my longest relationship. She lived until 2006. When she was dying, I was living in San Diego. It was so boring and I spent the whole five years I lived there anticipating her death, really aware of time. At one point, I got a video camera and I would just take it with me on these long walks with Rosie. I thought I was making personal cinema, chatting while I shot, though in fact what I said was completely boring. But the actual walk and my bad camerawork were beautiful, so I transcribed it exactly. The book's a lot of things but one is our walks, which are inter-

cut with everything else. his is really what I do — on some level my writing's just a really medieval account of what's there. A loose and meticulous copy. What's there is often fantastic. Like when I first got Rosie, I looked into her eyes and thought, his is my father. I was eleven when he died and I was always obsessed with him. It was a joke I had for the sixteen years of Rosie's life — that my father came back as my dog just to hang out a little more.

Also, I find myself thinking that maybe this is the last AIDS book, which is not to say that people aren't still dying of AIDS, but not like they were in the eighties and nineties. I lived through so much dying that it almost became commonplace. And I had Rosie through all of that time, too. I've had her through the time of so many relationships that bloomed and failed. She's the metonym for so much stuff.

INTERVIEWER

How old was your father when he died?

MYLES

He was forty-four. It was an alcohol-related accident. He was successfully drinking himself to death. He'd reached the point where he had convulsions whenever he stopped. There was a piece of furniture in the house that was there when my parents bought it, and they wanted to get rid of it but they couldn't angle it down the stairs, so they decided to push it off the roof. They got it out on the roof outside my brother's bedroom, and we were all told to go downstairs and watch. It was Saturday and it was my sister's birthday. We're all waiting downstairs, in the gap between the houses, and my father yells, Here it comes, and he comes lying down and lands at our feet. The story was that he had a convulsion. Because he was exerting himself, he had a seizure and he went over the roof instead of the dresser. We lived near Arlington Center, which is northwest of Boston, and fire trucks and everybody in the town were suddenly on our street. But he came to — he woke almost instantly after he hit the ground — and went to the hospital. He had pounding headaches for the next two weeks, and he went for an EEG, but they had crappy instruments in 1961, and they sent him home. I was in school that day, in seventh grade, and there was going to be a party that night, the first girl-boy party I was allowed to go to. Junior-high sex. There was total excitement that day in school. I got in trouble for laughing as we were all going down the stairs, so the nun said, Eileen Myles, write "I will not talk in the corridors" five hundred times as punishment. It was like a prompt. Does Sister Ednata know how many poems she unleashed with that command? When I got home, my dad was taking a nap on the couch. My mother said, Why don't you just set up a table in the den and do your punishment there so you can keep an eye on

him. I'm going to go out and hang clothes. She's in the yard and I'm writing "I will not talk in the corridors," and he starts to make weird sounds and he died there right in front of me, his face changing colors and the death rattle and the whole thing and I just kept writing "I will not talk," "I will not talk." And then he's dead.

INTERVIEWER

Jesus.

MYLES

Yeah, that's my trauma. And I don't know if this is my family or the working class in general, but nobody ever talked to me about it. It was never discussed, not one time. Except once about twenty years later. I went home to visit my mother in Arlington and out of the blue she goes, I know you were alone in the house when your father died — like we'd been just talking about it.

INTERVIEWER

How do you understand that moment?

MYLES

It's the insane scale of my family. Is it humor? Our brand of silence? I think she was just verifying the facts to me, that I wasn't invisible, that she did understand what had happened. I think. Part of her knew that was necessary. My mother was an orphan, and she wore her own suffering like a badge. We were like children of concentration-camp survivors. She was always telling us we had it so easy — she was four standing on a chair doing the dishes, they made her work hard in all those houses she was passed around in. I'm sure she was beat up and sexually abused, but she's always in control. She's also sort of a miracle in a way — very strong. I mean, I really love and admire my mother.

INTERVIEWER

Did she and your brother and sister notice that you had a way with language? Do they say, We always knew Eileen would be an artist?

MYLES

Oh sure. I was like the family clown. the middle child entertaining. I was a lousy student, but interestingly the nuns always let me write plays or do drawings, endless special projects. I made art when I was a kid so I wouldn't fail in school.

INTERVIEWER

I remember reading somewhere in your work that your brother was “the smart one,” the one on whom the family hopes were pinned.

MYLES

I was the creative one and Terry was the genius. He was supposed to go to Harvard. It put incredible pressure on him, to his detriment I think. My younger sister was kind of on her own. My brother and sister never liked each other. They are still fighting. I always had the freedom to come out from behind — which is kind of a female position. I was reading on my own for years — novels, a lot of sci-fi — and doing my special projects in school, which even meant organizing a band and writing songs. It was the sixties, so there was a little bit of a peace-love, Sister Corita tone to my education, and I made things as a way to survive it. If I had been a good student and an achiever, I might have been excited by a more systematic approach to writing than what I do. People loved to throw around the word rigorous in the eighties. I'd go bleh. When I started to pull something out of the pool of incoherence, it was exciting in itself. Later, I found theory next to the bed. I had girlfriends who went to college after I did, and they'd be reading Fredric Jameson or the Situationists or Deleuze. My girlfriends introduced me to those books, not the poets of my generation. If it didn't come through the bathroom or the bedroom, I didn't find it.

INTERVIEWER

Your fiction uses all kinds of material from your life. What does the frame of fiction or nonfiction do for you? How and when do you choose to present your work as one or the other?

MYLES

What's fictional is arrangement — what follows what. If somebody is lying to you, part of what they're doing is hiding things, omitting stuff, changing the order of things. And that's fiction.

INTERVIEWER

You address the question of working from life explicitly in your fiction and poetry. In *Chelsea Girls*, for instance, you as narrator anticipate that people will be furious seeing their lives used as material.

MYLES

Right. And I answered by saying, his is my life, not yours. It's my material. he tone's maybe a little bit hurt. It's Schuyler-esque. You tell it cause you're lonely – you're the only person inside that life.

INTERVIEWER

You took care of James Schuyler at the end of his life. How did you come to know him?

MYLES

I was at a reading at St. Mark's and I needed a job. I was telling Charles North how broke I was, and he said, I know a job you could have. I don't know if you'd want it ...

INTERVIEWER

What was the job description?

MYLES

Really, they needed someone to give Jimmy his pills and spend time with him. He had burned through all the friends who had been willing to let him stay in their homes and live with their families. That was over. He'd had a lot of breakdowns, and there was nobody left who was going to let him into their intimate space. His friends were already in late middle age. Jimmy was in his fifties. They got a lawyer and created a trust fund for him. They got him a room at the Chelsea. They got him me. And I had to move his papers and his clothes out of the Allerton, which was this dive hotel around the corner where he'd moved after burning down his room someplace else, and I moved his crap from there over to the Chelsea. his is all exactly in *Chelsea Girls*. My writing has made me a redundant human being.

INTERVIEWER

I don't remember if the burning down of the room was considered deliberate.

MYLES

He did what anybody on a lot of pharmaceuticals does at some point – he took them all. Whether it's suicidal or absentminded. He was definitely smoking a cigarette, and the room went up. The day after I spoke with Charles, I went to the Chelsea, where they had already installed him. He was lying on his bed with this long hair.

INTERVIEWER

He was a big man at that point, wasn't he?

MYLES

He was skinny. He got fat after I got my hands on him. I would make him cheeseburgers and bring him whatever he wanted, and he just blew up. I was like twenty-nine and pretty androgynous-looking, and I think he thought I was a cute young boy when we were introduced, so he said, Great! Hi, babe. It was that kind of interview – where you walk into the room and you're instantly hired. Most of the jobs of my life are like that, even UCSD. A lot of wanting, and then blam. I spent something like five hours a day, seven days a week with him for about half a year.

INTERVIEWER

Did he talk?

MYLES

He would lie on his bed and read, and I would sit in my chair and read. Every day I arrived with the papers and his pills, and I would leave him his night dose before I let. At the beginning, I think he was on Thorazine, and they changed the drugs a number of times while I was with him. Most days I would come in with a hangover, and he'd be like, What'd you do last night, babe? And I would tell him about my adventures of being a young queer.

INTERVIEWER


Would he give advice?

MYLES

A little bit every now and then. Oh, get rid of her. Dump her. Stuff like that.

MYLES

We talked about who we liked. he floor was covered with books, people just sent him books. He read some of them, not most. His room really introduced me to what it looks like when you send your book to somebody. I'd pick things up and say, Do you like this? Do you like that? Often, he'd say, I think he's hot, he's a nice guy, she's lovely.



INTERVIEWER

Did he ever read your poetry?

MYLES

He did. He actually said a certain poem I showed him was a masterpiece. I don't think he was kidding. You don't get that every day.

INTERVIEWER

Do you think of him as an influence on your poetry?

MYLES

Huge. The first poem of his I read was when I was going to Queens College. It was "Poem," the one that begins, "How about an oak leaf/if you had to be a leaf?" That poem still feels like a miracle. In all of them, it's like his sensibility is bursting through the description all the time, so there was an air of magic or surreality, like things were about to change. He is an adept at this surprise surreality, like that poem where suddenly it's "What the Dentist Saw" and you're inexplicably looking down a red throat. His work is dangerous the way fairy tales are.

INTERVIEWER

I want to return to the question of fiction and nonfiction. Have you had to deal with fallout from using real people and experiences in your fiction?

MYLES

Absolutely. It's not so different from the way people first responded to cameras, like they were afraid it was stealing their souls. And it is. Art hurts your animal eyes. Your family feels like you are stealing their souls if you write about them. Most artists like to be written about in a certain way. When somebody's in love with you, they think it's amazing you've written them a poem, and when they don't love you anymore, they hate those poems. They wish those poems would go away.

INTERVIEWER

Does your family read your work? Are they proud? Are they scandalized?

MYLES

It's changing right now. People are retiring and they have more time. My sister's partner has been really generous. I think people's partners in general often express the emotional reality of the situation first. My family kept my work at arm's length. Working-class people are a little uncomfortable with you thinking you're special. Just the fact of publishing, having a career, means you must think you're special. It's a violation of the code of being a cog in the wheel. Not that my family truly thinks of themselves that way. We were brought up to feel that we were geniuses, but the world didn't see us that way and everyone's a little grumpy about it. I'm one of them. My writing even deliberately hugs anonymity and grossness so I won't be punished for thinking I'm special. It's humble, conceited work. I have mostly dated middle-class or upper-middle-class women, and I would watch their parents excitedly

talking to my girlfriends about their accomplishments. Then they'd ask about what I was doing, and I always became nervous because I thought I was being tested. But that's really what they do — examine the goods lovingly. But I felt like a criminal for a long time. Part of an artist's social skills is the ability to comfortably talk about yourself. In a more working-class social situation, that is verboten. You'd be mocked for talking about yourself. You may talk about your body, your weight, the weather, work, but to talk about what you're doing is to pass a sort of judgment on everybody else. So I think it's the code of my family to be secretly proud of me but to never talk to me about my work. I have a little box in my storage unit called "The Mother Lode" because when my mother went from the house I grew up in into a senior center, she said, Well, you probably want these — and she meant all my books.

INTERVIEWER

You frequently use your own proper name in your work. In *Cool for You*, for example — "I gained the attention of the room. Eileen Myles. Everyone watching my body, I felt." Why does your name appear within your writing and not just on the title page?

MYLES

It's like a boundary against annihilation. When I arrived in New York, I was just glad to be in the phone book. But to use your own name in your writing always felt like a dare, with a lot of feeling underneath it, a lot of need and bravado. In *Chelsea Girls*, there's a story where she is — and I find it very easy to talk about the Eileen Myles character in the third person — where she is raped. She's sitting on a boat at Cape Cod, and she writes her name in the sand with her foot. And then erasing it, it's like the precarity — such a stylish word — but the precarity of the name and the self is so real.

Remember the movie *Monster* and how amazing it was that Charlize Theron, such a beautiful woman, would make herself so ugly, playing a sex worker, playing a violent woman? But a huge number of women in the world live that way. She's a monster? I don't pretend to live there, but I've tasted enough to know that the primary fact of one's name is quite precious. It's the writer's body. I suppose I do feel like a monster in some way. All my friends are monsters, Dodie and Kevin and Bob. Dennis Cooper? *Monster*. The New Narrative writers all use their own names. And I have turned my own name into a fake name that inhabits my fictions.

INTERVIEWER

You mention the New Narrative moment. The Language poets led an avant-garde attack on narrative forms, arguing that the politics of a work of literature reside in its form and not in its content — that narrative, no matter what it was narrating, was always in some sense bourgeois. And they insisted that the speaker in a poem or work of literature was a fiction that should be undone. What is your relation to Language writing or avant-garde poetry?

MYLES

I think the world I landed in poetically in the seventies was the avant-garde. Everything coming out of St. Mark's regarded itself under that banner and opposed itself to more academic stuff, or what stood as the mainstream of American poetry. Nobody at St. Mark's would have read, say, Louise Gluck or Mark Strand. They were regarded as the other team. I think there was a feeling that we were inherently more political, that the work felt more political. Language writing came out of this same St. Mark's scene. Charles Bernstein and Peter Seaton were both in Bernadette Mayer's workshop. Everyone read Clark Coolidge, John Wieners, Lorine Nie-decker. And everyone was in the same magazines for a while. But when the little L=A=N=G=U=A=G=E pamphlets appeared in New York, that was when it seemed like the critical discourse had become almost more important than the work, and definitely the Language poets on the West Coast were clocking you over the head with what they believe mattered. You got the sense that suddenly there were rules. It's tricky though, because we all came up together. We make many of the same moves in our work. It's all types of appropriation and what people now call formal constraints and erasures. Maybe the innovation for Language poetry was that often it was simply that, and I probably should admit that I just like a friendlier avant-garde. Writing that's more of a mess and has more of a relationship to bodies — that's what we were living and dying in, and many of us weren't bored with that yet.

INTERVIEWER

Do you think there's a tension between a working-class identity and the avant-garde's suspicion of personal experience? If you feel the tension, how do you respond to it?

MYLES

It just makes me go further. *Not Me* is a book that is personal and dissolves the personal at the same time. I don't think there's any conflict, certainly not a class conflict. Rae Armantrout and Ron Silliman are both working class. And we're all sort of in the same caboose now. And friends, of course. When they were younger,

they were sort of cops about it – it being their poetics. Rae is brusque, Ron was ... dogmatic. But Language poetry is itself a class. And recently there's been a way of talking that makes Language poetry or conceptual poetry be like this pancake that covers all of the avant-garde, and that's not true.

INTERVIEWER

Do you have a sense of why you're drawn to a short line in your poetry?

MYLES

Speed. And often I carry little notebooks, so that's just what it is. They're also comic-strip balloons but going vertically, not horizontally.

INTERVIEWER

I think Robert Creeley said somewhere that he was devastated when he heard William Carlos Williams read because he had developed a poetics based on the idea that those line breaks were moments of hesitation – that you are supposed to read the breaks – and Williams just ran right through them. Do you think of the line, of the text on the page, as a score for performance?

MYLES

No, they're a score for attention. They tell you how to read when you're reading on the page, though they also guide me when I read out loud.

INTERVIEWER

A score for attention?

MYLES

Just take this much in and keep moving. I don't want to shovel poetry into people like it's information. I want it to be perceptual, like a dot. I feel like if it's only that much, who wouldn't read it?

INTERVIEWER

Do you tend to read your poems the same way each time? Do you feel like there's a right way for you to read them?

MYLES

No. That seems obscene to me.

INTERVIEWER

Obscene how?

MYLES

I hear them a certain way. The most exciting thing is to read a poem out loud for the first time. There's a whole kind of inside thing bursting out, and I'm always dying to hear it. I do hear it in my head, but I never read it out loud to myself until I'm in front of people.

INTERVIEWER

You never read a poem out loud when you're writing it?

MYLES

I don't want to hear the sound of my own voice. It's the sound of something in me, but it's not my voice. It isn't a literal voice — at all. But there is a murmuring. I have some very... I don't know if sentimental is the word, but I have some thoughts about what poetry is and how old it is and what it means and what it could be. I feel like it's this old thing mumbling inside of me. When I first started to write, in my twenties, I did associate poetry with being fucked up, and poetry definitely managed something for me — the dissonance between the world just as we're being invited to enter it and that whole world inside of you and what to do with that gap. I had a mentally ill grandmother who I had, and have, a deep attachment to. She kind of mumbled and had a weird little West Ireland accent, and when I started to write I'd swear it was her. I felt this grrrrr, like something inside of me that was not me. It's very weird — I feel like I'm working for Eileen or something, like I have this job being her performer, learning this argot of hers, finding it. But I think it's older than me and it's older than my grandmother.

INTERVIEWER

I just assumed that all poets read aloud in the process of composition. I'm intrigued by this idea that there would be something sacrilegious about reading it to yourself

before you read it to a community.

MYLES

What is so great — I'll even say holy — about reading a poem for the first time in front of people is that you're sharing what you felt in the moment of composition, when you were allowing something. When I'm writing the poem, I feel like I have to close my eyes. I don't mean literally, but you invite a kind of blindness and that's the birth of the poem. Writing is all performance. Something's passing through. When people talk about formal constraints that's just technology, that's fashion. I like fashion, but you keep adjusting those things to let the other thing happen. Performance is us writing what's using us, remarking upon it.

INTERVIEWER

Do you think of yourself as a performer in a way that's distinct from other poets?

MYLES

I'm intrigued by all kinds of vocal performance. I worship singers. I always wanted to be one. When I first came around, I could not figure out how to get work published, and what was available to me was open mics. To me that's the first level of publication. You write the poem. You edit the poem. You type it up or you print it out. You show up publicly. You read the poem out loud. And then after that, it goes into some magazine, and then after that, it goes into a book. But I think reading is the most live, the most accurate, the truest version of the poem.

INTERVIEWER

So if you had it your way, you would always want your live performance of the text to accompany the book. Maybe have the book come with a recording?

MYLES

That would be great. I went to a studio in London and recorded *Inferno* for three days. It was the most fun I ever had. The recording engineer was a really good guy, and smart. I was reading to him. And he had a sense of humor, and he'd laugh at certain things. It was very intimate. Like he would say, Could you do that line again? Or I'd say, Let me read that one again, finding a place where we would agree that this was the right way to do it.

I've been experimenting with performance since the eighties. I memorized for a while, but I was stiff. After that I started to improvise, because in fact I did speak with my hands, I actually did move my body in a purposeful way, so I did a whole series of improvised performances – talking, basically. When somebody would ask me to come do a reading, I would decide for this group to talk about this thing, that story. And it was very satisfying because I would look at people, and I would feel the room in this way. That's what led me to run for president in 1992, actually – my sense of the energies of a room as political energies, campaigning as performance. It was an amazing political education to feel the sentiment of the room. Having done all that, when I came back to only reading, I realized that there was no reason not to move my body when I read a text. I mean, I wrote it. It's me. And if you've ever watched people play jazz there's a lot of twitching and lips pursing and strange movement. It's about getting the sound out. I became more aware of how music is composed of silence and pauses, and I know what the poem sounds like in my head, so I started trying to be more loyal to that. Like if I hear a three-count, I stop there and take that long piece of silence. That's really important to me. Now I have all that to think about if someone describes me as a performer rather than fearing that there's a little dis in there, like, Hey, you're a spoken-word poet.

INTERVIEWER

It implies you're not as good on the page. That you require a supplement?

MYLES

Or that you need to wow them with your message.

INTERVIEWER

You have a Boston accent ...

MYLES

In my writing more than my life. here's just no true working-class vernacular in my life. My parents were the children of immigrants. My dad had two different kinds of Irish accent. My mother wanted us to speak good English because her first language was Polish. But the kids next door, who were lower class than us, spoke like Huck Finn. I wanted that. Part of it was my longing then to be real, like in books or in comic books. And when I briefly went to graduate school, people were talking about black English. I thought, Isn't there some white equivalent? here is, but there isn't. White people are too afraid to be trash. I was told the other day that I have the

most hated accent. She meant I have the accent considered most racist to black listeners. That's working class.

INTERVIEWER

How has a conception of the vernacular — the speech, rhythm, and accent of the spoken — mattered to you as a writer?



MYLES

The vernacular is the place where everything meets. It's a gathering of people. Think

of Sons and Lovers when Paul Morel goes to the pay window at the coal-mining office and talks in his local vernacular even though he's an educated guy. I lived in Provincetown with a girlfriend, and we owned a house and workers would come over. And then I'd step into my Massachusetts accent to get the guys to not fuck us over. I think that's avant-garde — the meeting of need and language. Take the Happenings in the sixties, which were the beginning of postmodernism — what they were mainly interested in was the interface. How do we put film and bodies and poetry together? Now we're in a world where everything's sampled. Is there any place where the recording doesn't meet the live event? What used to be new art is now life. We're reeling around in this giant performance. When I wrote theater, I would think of each play as a Christmas tree, and the moments, the scenes in the play, were like ornaments. You can't see it, but the tree's there, and we're putting on the ornaments, we're putting the tinsel on, layers and layers, and moving with a confidence that there's a place. And you compose fearlessly because there is a tree, there is a place. I feel like the vernacular is that.

Which reminds me, the poet I have been most excited about lately is Fred Moten. He is piling vernaculars on and singing. Even the academic vernacular. He is so smart. I liked all five of the books we had as finalists for the National Book Awards. Even Louise Gluck's book. I had never read a book of hers. Beautiful trembling lines, and I loved the prose. But she doesn't know how to end a poem. That's what really marks the difference between the mainstream and the avant-garde. It's a different sense of the whole. If you're in danger you have to know how to get out.

INTERVIEWER

You've often referred — both in your writing and in interviews — to the career of the poet, to your career. I'm wondering how this fits into the discussion of class.

MYLES

There's a fundamental problem in working-class families. It's like you revere art, you believe in reading, you believe in books, but you don't understand their production. That's the disconnect. Those are the keys you can't have. And that's the non-lineage that cuts people from other classes out of the art life. Art looks like a lottery from out there.

INTERVIEWER

So what's at stake for you in insisting on writing as a career?

MYLES

I have made myself homeless. I have cut myself off from anything I knew prior to living in New York. I did this to myself, so I know exactly how it happened. Yet in the poetry world, people need to act like they don't know how this happened. Like a *je ne sais quoi*, but it's them. Here's a faux vernacular, as though the ambition must be hidden at all times, to be more, I don't know, attractive? It's the loafer posture, the veneer of I don't really need this. People loved to talk about how Frank O'Hara didn't really care about getting published. That doesn't jibe with my experience.

INTERVIEWER

Many poets — O'Hara, Ashbery — claim not to work very hard. Like O'Hara's quip about knowing a poem is done when the phone rings.

MYLES

It's like we're not doing business, we're golfing. And there might be a little gender in there, too. When I was in Ireland I met a man who was Beckett's favorite director. He talked about Beckett and how he wasn't ambitious at all and how he had no idea how to get his manuscripts to publishers. But he had these ladies who he would have sex with who worshipped him, and they would type up his manuscripts and bring them around for him.

INTERVIEWER

We could go back to Milton's daughters, right? Taking dictation.

MYLES

There's a whole female industry engaged in materially supporting the illusion that the artist doesn't work directly on his legacy, his immediate success. He's just a beautiful stoner boy or an intellectual. All thought. No wife? I like turning that illusion inside out. And making the work be literally about the field and the failures and even the practice. I wrote about these things in *Inferno* because Dante did. We should let the writing world and its ways of distributing awards be part of fiction. We should expose the very cultural apparatus that is affecting the reception of the book you're reading. What's dirty is that we're not supposed to talk about how it has sex and reproduces.

INTERVIEWER

You're both a maverick and very social. You've been involved in various "scenes" — St. Mark's, parts of the art world, performance groups like Sister Spit — without being assimilated into any of them.

MYLES

People render categories like modes of dismissal. If someone considers themselves a Language poet or a conceptual poet, they generally will introduce me as a New York School poet. It's a way of othering, but it's also a way of building something, monumentalizing their own efforts. Voila! his edifice. Where's Marjorie Perloff to break the bottle on it? It's the opposite of vernacular. In fact I have just discovered that vernacular comes from verna, meaning "slave," a native — a born slave.

INTERVIEWER

Are there any categories — whether pertaining to sexuality or class or style — that don't rub you the wrong way?

MYLES

I've been called a folk poet. I think that's kind of cool.

INTERVIEWER

I don't know what that means.

MYLES

I don't either. Well, if there are folk singers, what does that mean? Does it mean folks hear you? Or does it mean that you're covering old songs? I don't mind being called a New York School writer, it just isn't necessarily true.

INTERVIEWER

Well, the New York School has always been a fiction, right? It's not like Barbara Guest's writing and John Ashbery's writing are obviously part of a single school.

MYLES

Same with Language poetry. Rae Armantrout and Charles Bernstein — not so much

resemblance. Is Harryette Mullen a Language writer? But these poets agree to see themselves in conversation together. They form a cloud.

INTERVIEWER

On the topic of insider-outsider, here I am interviewing you for the Paris Review, which has an experimental history but is also in many ways a mainstream publication. Are you at all afraid of being canonized?

MYLES

I'm not, actually. Everything will ruin you, why not this. But if you do your work, somebody else will celebrate you, too. And we are having an interesting conversation. I'm talking my head off. This is a good experience. It could only bring good. When I took the job at UC somebody came up to me at a reading and said, How does it feel to have sold out? I'm like, What did I sign up for? A life of poverty?

INTERVIEWER

Do you reject work because you feel like you've done it before? Do you kill a poem or piece of prose because you feel it's too familiar?

MYLES

I never kill a style. I like the idea of writing a poem I could have written thirty years ago. I'm the factory. My writing fears manifest more on the order of my inability to stop being Eileen Myles. I guess I don't worry about my poems so much. I worry about me. It's really creepy to be addicted to yourself or the performance of yourself. Like looking at your phone too much.

INTERVIEWER

I think of how the topic of fame comes up in your early work. Sometimes it seems like a joke but sometimes it seems quite serious — I always knew I'd be famous, I'm fated to do this work. And of course you are famous, but it's not as though you can't go shopping without being recognized. Didn't Ashbery say something about famous poets not being like famous human beings?

MYLES

There's nothing more ambitious than a young poet. You feel omnipotent. You're on

the upswing of bipolar. And that enrages older poets — which, to a certain sensibility, only makes you want to be more vapid and fame obsessed and glib. No one can tell you what the limits are. As it should be. No one's got the keys to the kingdom.

A piece that's missing in our talk about the avant-garde is gayness, campiness, queerness. Somebody like Arthur Russell was where avant-garde and discos and Buddhism met. Nightclubs. That was in the air when I was young. What do you do with someone like Ariana Reines right now? Claudia Rankine is famous. She's political and she's glamorous. I think it's more interesting to think about how poets could have ever been so drab, why they would have made the choice to be poor, to be obscure, to not want fame. I guess it was a form of resistance, but lots of us were feeling the other thing. We weren't leaving America, we were making another one. I hope we still are.

INTERVIEWER

To what degree does a literary posterity ever enter your consciousness? I mean that old kind of poetic fame, that dream of immortality?

MYLES

What do I care? I'll be dead. I really care much more about the present, having this life that I want and being able to write the things I want to write and see them well published and not feeling obstructed, ambition-wise. I have a poem that contains the line "Fame is merely advanced sentiment." It's sort of like this extra feeling you're putting forward. And what do you call that extra feeling? We don't know where this is going to go. Like when you're in a vehicle or on a train and you see fragments of vistas — a sign, a little town — and for some reason you remember that little town all your life. I feel like the repetition of fame is like that. It doesn't have a place, but it has an amount of space that you want to know more about and you'd like it to be there. It's not quite posterity — it's like building a periscope, needing to see more, wanting to write already from that vista.

INTERVIEWER

On the topic of seeing more, when did art criticism — if you think of your prose about art as criticism — start for you? How does it relate to your other kinds of writing?

MYLES

Growing up, I was supposed to go to art school, but I went to college to prove that I was, I guess, intelligent. I stopped drawing at that point, so when I first came to New York, I would look at art, but I was a little sad and resentful. I mainly looked at photography in the seventies. I was friends with Peter Schjeldahl, though, who took me to galleries and pushed me toward being an art writer. But it was Rene Ricard who actually talked me into it. He did a performance at the Guggenheim in 1981 and insisted I review it for Art in America. I was deeply insecure — I didn't know if I could write art reviews. I wasn't the greatest prose writer at that time. Art writing, working with editors, is how I learned to write prose. That's the fact.

Then once I got sober, I looked at art and just got high. The pleasure overrode any resentment I had that I wasn't a visual artist. And part of that pleasure was to write. For ten years all I did was go to the Poetry Project, and after being the director, I turned toward the art world and performance. I wrote plays and all these other things. Not Me is about that, about being alive in this completely new space.

INTERVIEWER

Do you think that change — getting sober — shows up in the way you write?

MYLES

Getting sober rescued my writing, because at the end of my drinking my writing had gone to shit. It had not been a problem for a while, but then suddenly it was really a problem. My mind sucked. I was becoming foggy and full of shit.

INTERVIEWER

Were you killing yourself?

MYLES


Absolutely. I would drink until I fell down. Or I would quit smoking — as if that were the problem. The bars would close and I would go running up First Avenue, jogging, full of energy, and then I would stumble and hit my head. I hit my head a lot. I would wake up and have black eyes. I was big for head injuries, which is how my dad died.

INTERVIEWER

And was it just like you decided one day ...

MYLES

My best friend, Tom, who was also a drug addict and a junkie – he had very little patience for my always going on the wagon. Having seen a parent die of alcoholism, I knew what was waiting for me. Tom had no respect for that. But then he came back after visiting his family one summer, and he had stopped drinking. We always talked about our alcoholic dads, but it turned out his had actually gotten sober, and mine had died. He urged me to join him in this new thing. I preferred to die, thank you. We really talked that way then. I remember David Rattray and I were invited to read at the Ear Inn in April, but we both independently told the person we thought we'd probably be dead by then.



INTERVIEWER

Were you afraid that if you stopped drinking, you wouldn't make work?

MYLES

I decided that even if I couldn't be a poet, I was going to stay sober. Which was kind of huge — that was the only identity I had. I was a dyke too, but . . . Immediately upon making that deal, I began to write better poems. It was like reorganizing my mind in some new way. I found this new condition to write in. It was both more and less. Everything became an incitement to do other kinds of writing. Can I write for the Village Voice? Can I be a performance artist? For me, the eighties, when I was in my thirties, was just tooling around and doing all these different things.

INTERVIEWER

What did you take from punk and other kinds of music?

MYLES

Those people were clearly part of my aesthetic family, and I felt like if they didn't have to apologize, I didn't have to. Here's an energy you get from music and popular culture, and it always felt closer to me than "fine art" did. I grew up in the era of low art, really great television — the fifties and sixties were amazing.

INTERVIEWER

What kind of TV did you find amazing?

MYLES

Dobie Gillis, Ernie Kovacs, Ed Sullivan. Everybody watched TV together, everybody was watching the same shows in the same way that everybody got Rubber Soul at Christmas, the way everybody got Revolver. Performance art in the eighties was just vaudeville by people who had grown up watching TV. TV was great as long as nobody knew what it was, because it was people coming right out of vaudeville, straight out of radio, and onto television. People would always come out to talk at the beginning of a TV show. There was always a monologue before these shows. Dobie Gillis sat on a rock and talked about girls. One stepped out of TV and performed or wrote a poem.

INTERVIEWER

So you've never conceived of poetry as antagonistic to mass culture in some way?

MYLES

No. Since it was what I was going to do, I had to make it out of things that were plausible for me. The plan was not to fail. To succeed meant translating the culture I already had or liked into what I was now doing.

INTERVIEWER

You've mentioned that Frank O'Hara's manifesto "Personism" is very influential for you. Can you say why?

MYLES

Because of the pause! Because of that pause in the piece where he's describing writing a poem and suddenly — I mean, I describe it as "the pause" but I've looked at it, and he doesn't actually pause, he just says, "I realized" that instead of writing a poem he could just call the person. What a transitional moment in the history of poetry, to be writing and to realize that now we have a technology that's in the world and doesn't have anything to do with poems and writing — and then suddenly it does. Suddenly the poem is the phone call. I think poetry was never the same. The poem is not in letter time, it's in telephone time. I think every kind of technology changes the entire structure of the culture, and that's why literary is such a cloying term. We're making books, but I think the metaphor for the book is not the book, it's all these other things. You may not use social media, but it's using you. You're writing in tweets, like it or not. To me, the notion of literary is what's academic now — to be literary is to write archaically. It's archaic to insist the book be a book and even to behave as if there is a realm that is "literary" that refers only to letters.

But O'Hara and the phone was the beginning. He's putting it out there in "Personism" so you'll read his poems differently.

INTERVIEWER

There are a lot of cell phones in your recent poems. And you're active on social media.

MYLES

I never wanted to be a poet. I wanted to be a spaceman. So this is it. My wildest fantasy of being a poet exists in social media because I feel like you are walking down

the street in your connected notebook. It's sort of like when you realize the poem is no longer a whole, it's just stabs at trying to notate a vivid, fleeting experience, so maybe a line comes to you – and it's really hard to figure out what's poetry and what's a tweet at this time – but the line comes and you literally can show the thing that you saw, or not. I mean, they can't be in the sensorium with you, but you make the joke, Here's the thing I saw and here's the line that came to me and I'm sending this out to seven thousand people right now and I'm gleeful that I'm not alone in this particular way.

INTERVIEWER

Every few days there's some important New York artist who says the city's dead and that everybody should go to Detroit or wherever. Do you feel that way, too?

MYLES

I figured you and I would talk about New York somehow. here's a way in which, if you lived here for a lot of your adult life, you're sort of screwed for living anywhere else, because it's like having a computer – you're used to it and to all your information being in one place. I may not go see any art, but all the art I want to see is here. I may not want to be with anybody, but anybody I want to be with is right out there, there are so many people who are here, and if not, well, if I wait a month, they'll come.

People who came to New York from Middle America in the thirties and forties were making a big move from a rural culture to an urban culture – to here, where they could have sex with a lot of people and drink themselves to death and make art. here was a chemistry they had to figure out in order to become something or somebody no one had ever heard of. I feel like it's a different equation now, but it's the same relative truth.

Call me Maybe: MariaDB Galera Cluster

Aphyr, September 1, 2015

Previously, on Jepsen, we saw Chronos fail to run jobs after a network partition. In this post, we'll see MariaDB Galera Cluster 7.4.7 allow transactions to read partially committed state.

Galera Cluster extends MySQL (and MySQL's fork, MariaDB) to clusters of machines, all of which support reads and writes. It uses a group communication system to broadcast writesets and certify each for use. Unlike most Postgres replication systems, it handles the failure and recovery of all nodes automatically, and unlike MySQL Cluster, it has only one (as opposed to three) types of node. The MariaDB Galera packages are particularly easy to install and configure.

Galera Cluster uses the normal InnoDB isolation levels locally—but we're interested in cluster-wide consistency guarantees. Between nodes, Galera claims to implement Snapshot Isolation—a reasonably strong consistency model.

Transaction isolation also occurs at the cluster level. Between transactions processing on separate nodes, Galera Cluster implements a transaction level called SNAPSHOT-ISOLATION. The SNAPSHOT-ISOLATION level occurs between REPEATABLE-READ and SERIALIZABLE.

This is not strictly speaking correct—but to understand why, we'll need to know a little more about SQL isolation levels.

SQL Consistency Models

The standard ANSI SQL isolation models define four isolation levels for transactions: read uncommitted, read committed, repeatable read, and serializable. Each level prevents an additional kind of unwanted phenomena: dirty reads, fuzzy reads, and phantoms. The standard's definition of these phenomena is somewhat ambiguous, and has been extended to a more rigorous spec by Adya and others.

For instance, Read Uncommitted (RU) must prohibit Dirty Write, which Adya terms P_0 . In P_0 , one transaction modifies a record written by another before that transaction commits.

Dirty Write (P_0): $w_1(x) \dots w_2(x)$

In this notation, $r_1(x)$ means “transaction T1 read x”. A w denotes a write, a means abort, and c indicates transaction commit.

Read Committed (RC) prohibits P0 and also P1: Dirty Read. Repeatable Read (RR) goes one step further and prohibits P2: “Fuzzy Read”. Finally, Serializability (1SR) adds one final exclusion: P3, or “Phantom”. Serializability is a very strong constraint: all transactions must appear to execute atomically—results should be equivalent to a system in which the transactions ran strictly one after the next, rather than concurrently.

Dirty Write (P0):	$w_1(x) \dots w_2(x)$	Prohibited by RU, RC, RR, 1SR
Dirty Read (P1):	$w_1(x) \dots r_2(x)$	Prohibited by RC, RR, 1SR
Fuzzy Read (P2):	$r_1(x) \dots w_2(x)$	Prohibited by RR, 1SR
Phantom (P3):	$r_1(P) \dots w_2(y \text{ in } P)$	Prohibited by 1SR

There’s a neat kind of symmetry here: P1 and P2 are duals of each other, preventing a read from seeing an uncommitted write, and preventing a write from clobbering an uncommitted read, respectively. P0 prevents two writes from stepping on each other, and we could imagine its dual $r_1(x) \dots r_2(x)$ —but since reads don’t change the value of x they commute, and we don’t need to prevent them from interleaving. Finally, preventing P3 ensures the stability of a predicate P , like a where clause—if you read all people named “Maoonga”, no other transaction can sneak in and add someone with the same name until your transaction is done.

If you’re having trouble figuring out what these isolation levels actually allow, you’re not alone. The anomalies prevented (and allowed!) by Read Uncommitted, Read Committed, etc are derived from specific implementation strategies. If you use locks for concurrency control, and lock records which are written until the transaction commits (a “long” lock), you prevent P0. If you add a short lock on reads (just for the duration of the read, not until commit time), you prevent P1. If you acquire long locks on both writes and reads you prevent P2, and locking predicates prevents P3. The standard doesn’t really guarantee understandable behavior—it just codifies the behavior given by existing, lock-oriented databases.

Locks aren’t the only ways we can protect transactional integrity. Some MVCC databases, like Oracle and Postgres, offer a different kind of model: Snapshot Isolation.

Snapshot Isolation

Snapshot Isolation arises from Berenson et al’s critique of the ANSI SQL isolation levels. The paper is subtle and dense, but the gist of it is that there are some anom-

aliases allowed by the original SQL definitions that are undesirable, and we might define a new isolation level that gives us most, but not all, of the benefits of Serializability, while still being relatively efficient.

In Snapshot Isolation, every transaction operates on an isolated snapshot of the committed data, taken at any time prior to the transaction's first read. The transaction can freely modify its snapshot, and those values are visible within the transaction, but are not visible to other transactions unless the transaction commits.

A transaction T1 may commit if the data it plans to write has not been written by any other transaction which committed after T1's initial snapshot was taken. If a committed transaction conflicts with T1's writes, T1 must abort. This is called First-committer-wins, and prevents transactions from stepping on each other's updates without seeing them. This is the "lost update" anomaly: P4.

P4: $r_1(x) \dots w_2(x) \dots w_1(x) \dots c_1$

Like Repeatable Read, Snapshot Isolation prohibits P0, P1, and P2, but allows some types of Phantom anomalies (P3). It is therefore weaker than Serializability, and strictly stronger than Read Committed. However, it differs from Repeatable Read in two respects:

- Repeatable Read allows an anomaly called A3—a type of Phantom, which SI prevents.
- Snapshot isolation allows an anomaly called A5B, called Write Skew, which RR prevents.

A3 occurs when T1 reads some predicate P, T2 modifies that predicate set and commits, and T1 subsequently reads that predicate set again and commits.

A3: $r_1(P) \dots w_2(y \text{ in } P) \dots c_2 \dots r_1(P) \dots c_1.$

Note that A3 is a specific case of P3: it only applies if both transactions commit, whereas P3 includes all cases where a predicate set is read and modified by another transaction. Snapshot Isolation precludes A3 (but still allows some other anomalies in P3) because T1 always reads the data from its snapshot. So with respect to this particular anomaly, Snapshot Isolation is stronger than Repeatable Read.

On the other hand, SI allows A5B, or Write Skew. Write Skew occurs when two transactions read two distinct values, then write on top of the other transaction's read. Formally:

A5B: $r_1(x) \dots r_2(y) \dots w_1(y) \dots w_2(x) \dots (c_1 \text{ and } c_2)$

A5B is prevented by Repeatable Read, but allowed by Snapshot Isolation since the write sets of both transactions do not overlap. We also know of one additional anomaly allowed by Snapshot Isolation, in which a read-only transaction can sneak in between two transactions which legally commit out of order because their write sets are disjoint.

A6: $r_2(x) \dots w_1(y) \dots c_1 \dots r_3(x) \dots r_3(y) \dots c_3 \dots w_2(x) \dots c_2$

In light of this, we should revisit Galera's claim that "SNAPSHOT-ISOLATION level occurs between REPEATABLE-READ and SERIALIZABLE." This is not quite correct: Berenson and Adya are clear that SNAPSHOT ISOLATION lies between Read Committed and Serializable, but is neither a superset nor subset of Repeatable Read.

Given Snapshot Isolation allows these anomalies, a natural question to ask is whether we can prevent the anomalies, obtaining full serializability, by restricting the classes of transaction that can run on the system. It turns out the answer is yes: we can force the transactions to conflict by promoting some reads to writes (forcing the write set to intersect and preventing commit), or by analyzing dependency cycles in transactions. These techniques allow us to turn a Snapshot Isolation system into a Serializable one. This is how PostgreSQL's Serializable isolation level is implemented—on top of an underlying Snapshot Isolation system.

So: we've learned about the ANSI SQL isolation levels (RU, RC, RR, and 1SR), and seen that Snapshot Isolation fits in alongside Repeatable Read, and below Serializability. It prevents Dirty Writes, Dirty Reads, and Lost Updates, and some types of Phantoms. However, it still allows an anomaly called Write Skew, so not all transactions in an SI system are guaranteed to serialize correctly. In order to verify Galera's claims of Snapshot Isolation, we'll have to design a test that fits within those guarantees.

Designing a test

Jepsen has a linearizability checker in Knossos, but that won't work here—Snapshot Isolation, and even Serializability, don't require that operations take place now. They simply have to take place atomically at some point in the history—maybe in the past or future. We need a different kind of checker!

Imagine a system of two bank accounts, each with a balance of \$10.

```

create table if not exists accounts
  (id      int      not null primary key,
   balance bigint not null);
INSERT INTO accounts ( id, balance ) VALUES ( 0, 10 );
INSERT INTO accounts ( id, balance ) VALUES ( 1, 10 );

```

Then we'll generate and apply transactions that transfer random amounts of money from one account to the other, so long as no account goes negative:

```

SET SESSION TRANSACTION ISOLATION LEVEL SERIALIZABLE
set autocommit=0
select * from accounts where id = 0
select * from accounts where id = 1
UPDATE accounts SET balance = 8 WHERE id = 0
UPDATE accounts SET balance = 12 WHERE id = 1
COMMIT

```

Because these transactions write every record they read, they must be serializable under Snapshot Isolation. This is not the same as Galera offering serializability for all transactions—we're just asserting that these particular transactions appear to take place atomically. By serializable, remember, we mean any history which is equivalent to the transactions executing in some sequential order. It's OK for operations from two transactions to interleave, so long as their outcomes are the same as if T1 had executed by itself, then T2 had executed after.

Claim: Given a set of transactions, each beginning with a read and writing every value they read, a Snapshot Isolation system must always produce serializable histories.

Proof by contradiction: assume these transactions may not always serialize. Then there must exist some possible history in which one transaction T1 appears to execute interleaved with another T2.

Lemma 1: Since the start time for a transaction must precede its first read, and the first operation in every transaction is a read, every snapshot's start time must occur before any of its operations. Similarly, each transaction's commit time must occur after its last operation. Therefore, the [start-time, commit-time] interval for a transaction covers all its operations.

Without loss of generality, assume T1 starts before T2 starts.

- Case 1: T1 commits before T2's start time. Operations from T1 and T2 cannot interleave, by Lemma 1, because their intervals do not overlap.
- Case 2: T1 and T2 operate on disjoint sets of accounts. They serialize trivially.
- Case 3: T1 and T2 operate on intersecting sets of accounts, and T1 commits before T2 commits. Then T1 wrote data that T2 also wrote, and committed in T2's interval, which violates First-committer-wins. T2 must abort.
- Case 4: T1 and T2 operate on intersecting sets of accounts, and T1 commits after T2 commits. Then T2 wrote data that T1 also wrote, and committed in T1's interval, which violates First-committer-wins. T1 must abort.

Cases 1 and 2 contradict our assumption that T1 and T2 do not serialize. Cases 3 and 4 contradict the Snapshot Isolation invariants. Therefore, if the system provide Snapshot Isolation, T1 and T2 must serialize. Since every pair of transactions must serialize, the total history of transactions must also comprise a serializable history.

Note that we've restricted ourselves to transactions that write every value they read in order to avoid phenomena A5B and A6. If we had transactions with intersecting read sets and disjoint writes sets, A5B and A6 might occur, and nonserializable histories could result!

Now introduce read-only transactions, listing all balances:

```
SET SESSION TRANSACTION ISOLATION LEVEL SERIALIZABLE
set autocommit=0
select * from accounts
COMMIT
```

Read-only transactions trivially serialize with one another. Do they serialize with respect to transfer transactions? The answer is yes: since every read-only transaction sees only committed data in a Snapshot Isolation system, and commits no data itself, it must appear to take place atomically at some time between other transactions.

So the set of all transfer and read transactions must serialize. This gives us two important invariants:

1. Since each transfer conserves money, the total amount of money in the system remains constant.
2. Since read transactions serialize and read every balance, every read should see the same total.

With invariant 2, we can write a function to verify that every read sees the correct total balance. Then we generate a mix of randomized transfer and read operations, and apply them to our Galera cluster.

Results

At low levels of concurrency—say five clients, performing about one op per second, things are fine.

```
{:valid? true,  
 :perf ...  
 :bank {:valid? true, :bad-reads []}}
```

But increase the probability of conflicting transactions by running, say, 20 clients pushing an aggregate ~150 transactions/sec for about a minute, and things go terribly, terribly wrong.

```
INFO jepsen.core - Analysis invalid!
```

```
{:valid? false,  
 :perf ...  
 :bank  
 {:valid? false,  
  :bad-reads  
  [{:type :wrong-total,  
    :expected 20,  
    :found 18,  
    :op {:value [6 12], :time 1717930325, :process 15, :type :ok, :f  
:read}}  
    {:type :wrong-total,  
     :expected 20,  
     :found 16,  
     :op {:value [2 14], :time 3253699251, :process 13, :type :ok, :f  
:read}}  
    {:type :wrong-total,  
     :expected 20,  
     :found 17,  
     :op {:value [8 9], :time 5110345929, :process 17, :type :ok, :f  
:read}}  
    ...
```

Every transaction should have seen a total balance of 20—but instead we read different values, like 18, 16, or 17. This means our read transactions don't see a consistent snapshot of the database. They actually see intermediate results from transfer transactions.

These inconsistent reads aren't just limited to read-only transactions—transfers can see them too. Here's an excerpt from the query log, showing a transfer transaction which saw an inconsistent snapshot of the world.

```
66 Connect    jepsen@192.168.122.1 as anonymous on jepsen
66 Query show variables like 'max_allowed_packet'
66 Query SELECT @@tx_isolation
66 Query SET SESSION TRANSACTION ISOLATION LEVEL SERIALIZABLE
66 Query set autocommit=0
66 Query select * from accounts where id = 1
66 Query select * from accounts where id = 0
66 Query UPDATE accounts SET balance = 8 WHERE id = 1
66 Query UPDATE accounts SET balance = 9 WHERE id = 0
66 Query COMMIT
66 Query ROLLBACK
66 Query set autocommit=1
66 Query SET SESSION TRANSACTION ISOLATION LEVEL REPEATABLE READ
66 Quit
```

8 + 9 is only 17: the client read an inconsistent snapshot of the two accounts, tried to move some money from one to the other, resulting in new balances 8 and 9, and tried to write those values back. Luckily this transaction failed to commit: it rolled back—as did every inconsistent transfer transaction in this test case. In the ticket I filed for this issue, I initially believed the problem might be limited to read-only transactions—not a great result, but not as catastrophic as permanent data corruption.

Unfortunately, not all inconsistent transfer transactions correctly abort. Some can commit, permanently creating or destroying money. In this run, for instance, an inconsistent read in a transfer transaction fabricates \$2 out of thin air:

```
...
{:type :wrong-total,
 :expected 20,
 :found 22,
 :op {:value [10 12], :time 200102098751, :process 12, :type :ok, :f
```

```

:read}}
  {:type :wrong-total,
   :expected 20,
   :found 22,
   :op {:value [6 16], :time 200109803013, :process 7, :type :ok, :f
:read}}
  {:type :wrong-total,
   :expected 20,
   :found 22,
   :op {:value [10 12], :time 200113103237, :process 6, :type :ok, :f
:read}}
  {:type :wrong-total,
   :expected 20,
   :found 22,
   :op {:value [6 16], :time 200128852818, :process 3, :type :ok, :f
:read}}]}}

```

The transfer transactions should have kept the total amount of money at \$20, but by the end of the test the totals all sum to \$22. And in this run, 25% of the funds in the system mysteriously vanish. These results remain stable after all other transactions have ended—they are not a concurrency anomaly.

```

...
{:type :wrong-total,
 :expected 20,
 :found 15,
 :op {:value [15 0], :time 130519175659, :process 14, :type :ok, :f
:read}}]}}

```

In summary, Galera does not provide Snapshot Isolation. A transaction does not operate on an isolated snapshot of the world; other transactions may modify the data it's reading.

Dirty reads?

So reads can see data that's being modified by a transaction. Is this P1: Dirty Read? We saw this behavior in MongoDB, where a read could see invalid data from a transaction that never committed.

To measure this, we'll design a similar test. This time, each operation writes a unique value to every row in a single transaction, allowing us to identify precisely

which transaction was responsible for the values a read sees.

If Snapshot Isolation held, every read would see the same number in every row. We know that Galera allows inconsistent reads, so we expect to see a mixture of different numbers in each row. But we can go one step further, and distinguish between values that were written by successful transactions, and those written by failed transactions. If we see data from transactions that never committed, that would be P1: Dirty Reads.

Luckily, the test results suggest that Galera does not allow Dirty Reads.

```
INFO jepsen.core - Everything looks good!
```

```
{:valid? true,  
 :perf ...  
 :dirty-reads  
{:valid? true,  
 :inconsistent-reads  
 [[21462 21466 21466 21466]  
  [21462 21466 21466 21466]  
  ...  
  [34449 34449 34460 34460]  
  [34460 34460 34463 34463]],  
 :dirty-reads []}}
```

We have plenty of inconsistent reads here—writes in this test always set every row to the same value, but we see different values in reads. However, none of the values we read came from a transaction which did not (eventually) commit. So: there are no dirty reads—at least, not in this test. There might be other conditions that allow Galera to expose uncommitted data, but I haven't found them yet. This suggests that Galera could support Read Committed.

Given these results, I suspect the inconsistent reads we're seeing could be A5A: Read Skew. In a Read Skew anomaly, a transaction reads x , but before it can read y , a second transaction sneaks in and updates x and y together. Because both transactions commit this isn't a Dirty Read, but it has similar effects: two records which should only change together can be changing independently.

```
A5A: r1(x) ... w2(x) ... w2(y) ... c2 ... r1(y) ... (c1 or a1).
```

Snapshot Isolation prohibits A5A by taking isolated snapshots for reads. This anom-

ally is supposed to be impossible in Galera—but, as we’ve seen, snapshots aren’t correctly isolated.

The Galera team responded by explaining that Galera does not honor first-committer-wins for performance reasons. No first-committer-wins, no snapshot isolation. No snapshot isolation, well... I’m not sure exactly what Galera does guarantee, but it’s not what it says on the tin.

Recommendations

Galera is easy to install—I spent weeks trying to set up MySQL Cluster to no avail, and got Galera Cluster running in a matter of hours. It offers support contracts, reasonable documentation, and doesn’t require you to navigate the choose-your-own-adventure of Postgres replication strategies. With homogenous nodes, simple configuration, and the wealth of MySQL tooling available, it seems like a solid choice from an operational perspective.

Unfortunately, even in totally healthy clusters, with no node failures or network failures, Galera Cluster does not satisfy its claims of Snapshot Isolation. At moderate concurrency, you should expect to read inconsistent state from other transactions, and to be able to write that state back to the database. Designing your applications to avoid write skew is insufficient: even if writes completely cover the read set, SI fails to hold.

Since Galera appears to forbid reading uncommitted data, I suspect (but cannot show) Galera supports Read Committed isolation.

The probability of data corruption scales with client concurrency, with the duration of transactions, and with the increased probability of intersecting working sets. Placing nodes further apart (for example, on the other side of the planet) will dramatically increase commit times, raising the probability of consistency errors.

I’m not aware of a workaround for these issues. I assumed materializing transaction conflicts or promoting reads to writes would be sufficient, but those techniques failed. If you’re a Galera user with a support contract, you can try asking them to address the issue. You might adopt a different database—though since Galera is the first distributed SQL system I’ve analyzed, and FoundationDB disappeared, I’m not sure what to recommend yet. Or you may decide that your transactions don’t overlap often enough, or aren’t high-value enough, to be a serious problem. Not all consistency errors are critical!

Galera has indicated that they may provide actual Snapshot Isolation, and possibly full Serializability, in future releases. So... stay tuned!

A Neural Algorithm of Artistic Style

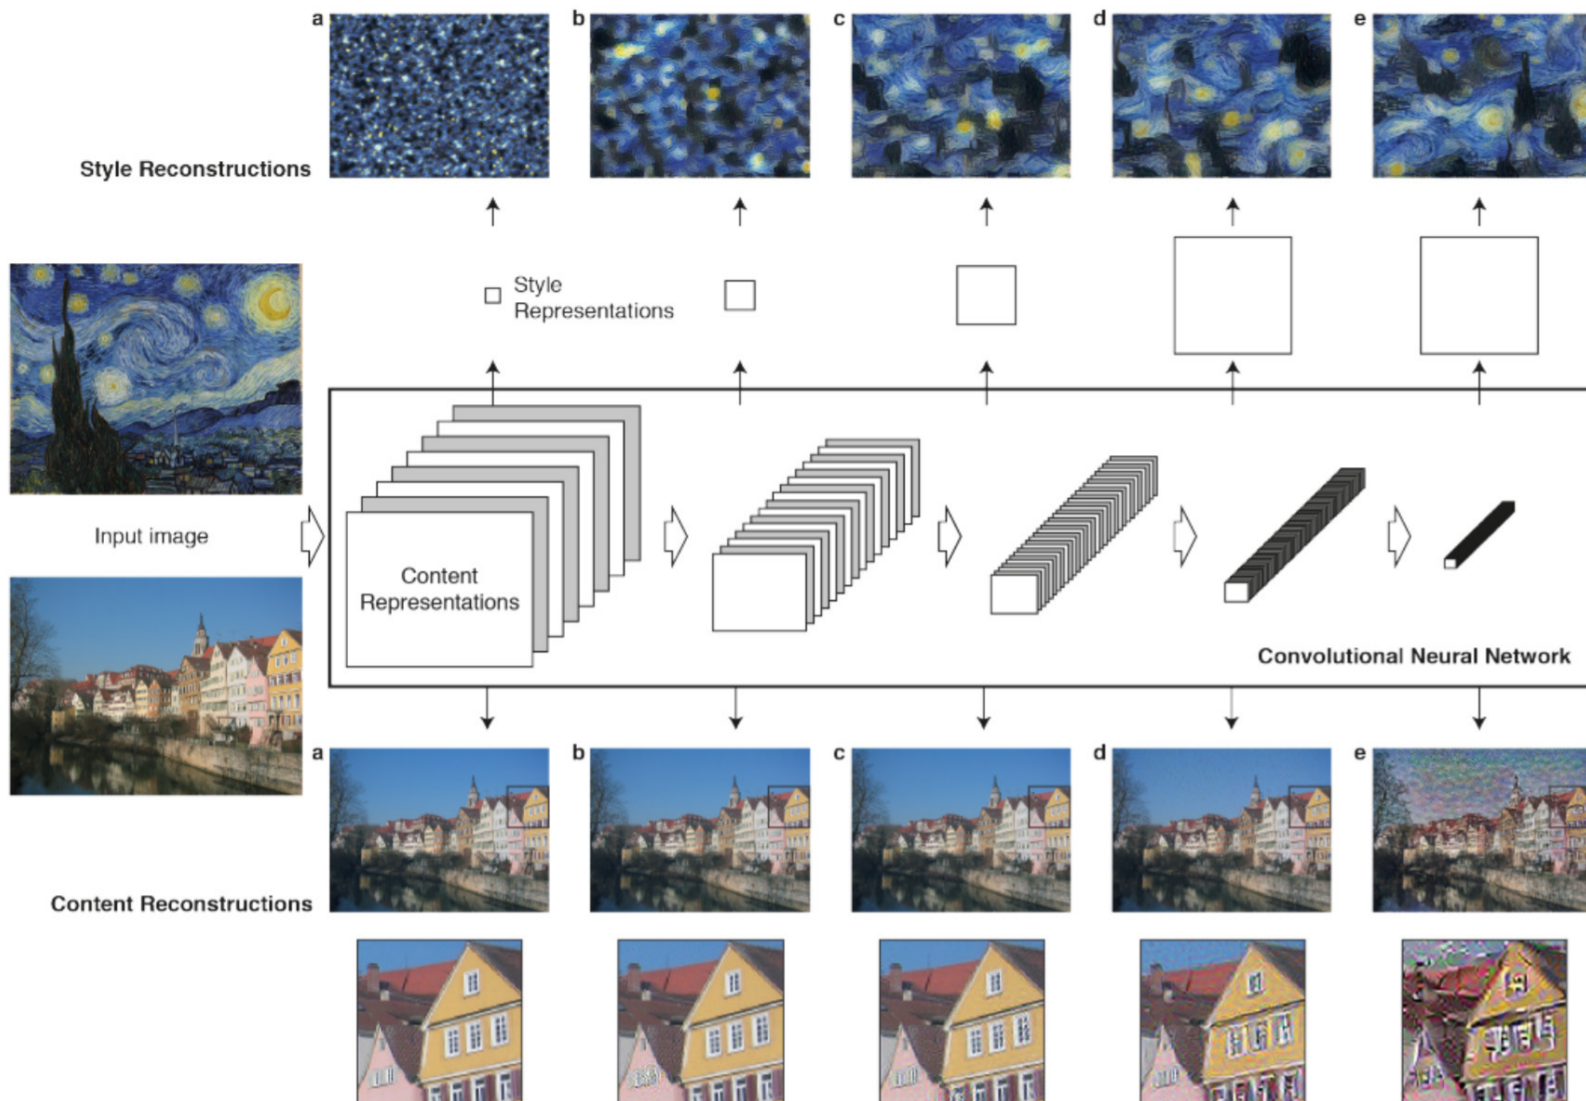
Leon A. Gatys, Alexander S. Ecker, and Matthias Bethge, ArXiv, September 2, 2015

In fine art, especially painting, humans have mastered the skill to create unique visual experiences through composing a complex interplay between the content and style of an image. Thus far the algorithmic basis of this process is unknown and there exists no artificial system with similar capabilities. However, in other key areas of visual perception such as object and face recognition near-human performance was recently demonstrated by a class of biologically inspired vision models called Deep Neural Networks. Here we introduce an artificial system based on a Deep Neural Network that creates artistic images of high perceptual quality. The system uses neural representations to separate and recombine content and style of arbitrary images, providing a neural algorithm for the creation of artistic images. Moreover, in light of the striking similarities between performance-optimised artificial neural networks and biological vision, our work offers a path forward to an algorithmic understanding of how humans create and perceive artistic imagery.

The class of Deep Neural Networks that are most powerful in image processing tasks are called Convolutional Neural Networks. Convolutional Neural Networks consist of layers of small computational units that process visual information hierarchically in a feed-forward manner (Fig 1). Each layer of units can be understood as a collection of image filters, each of which extracts a certain feature from the input image. Thus, the output of a given layer consists of so-called feature maps: differently filtered versions of the input image.

When Convolutional Neural Networks are trained on object recognition, they develop a representation of the image that makes object information increasingly explicit along the processing hierarchy. Therefore, along the processing hierarchy of the network, the input image is transformed into representations that increasingly care about the actual content of the image compared to its detailed pixel values. We can directly visualise the information each layer contains about the input image by reconstructing the image only from the feature maps in that layer 9 (Fig 1, content reconstructions, see Methods for details on how to reconstruct the image). Higher layers in the network capture the high-level content in terms of objects and their arrangement in the input image but do not constrain the exact pixel values of the reconstruction. (Fig 1, content reconstructions d,e). In contrast, reconstructions from the lower layers simply reproduce the exact pixel values of the original image (Fig 1, content reconstructions a,b,c). We therefore refer to the feature responses in higher

layers of the network as the content representation.



To obtain a representation of the style of an input image, we use a feature space originally designed to capture texture information. This feature space is built on top of the filter responses in each layer of the network. It consists of the correlations between the different filter responses over the spatial extent of the feature maps (see Methods for details). By including the feature correlations of multiple layers, we obtain a stationary, multi-scale representation of the input image, which captures its texture information but not the global arrangement.

Again, we can visualise the information captured by these style feature spaces built on different layers of the network by constructing an image that matches the style representation of a given input image (Fig 1, style reconstructions). Indeed reconstructions from the style features produce texturised versions of the input image that capture its general appearance in terms of colour and localised structures. Moreover, the size and complexity of local image structures from the input image increases along the hierarchy, a result that can be explained by the increasing receptive field sizes and feature complexity. We refer to this multi-scale representation as style representation.

The key finding of this paper is that the representations of content and style in the

Convolutional Neural Network are separable. That is, we can manipulate both representations independently to produce new, perceptually meaningful images. To demonstrate this finding, we generate images that mix the content and style representation from two different source images. In particular, we match the content representation of a photograph depicting the “Neckarfront” in Tübingen, Germany and the style representations of several well-known artworks taken from different periods of art (Fig 2).



As outlined above, the style representation is a multi-scale representation that includes multiple layers of the neural network. In the images we have shown in Fig

2, the style representation included layers from the whole network hierarchy. Style can also be defined more locally by including only a smaller number of lower layers, leading to different visual experiences (Fig 3, along the rows). When matching the style representations up to higher layers in the network, local image structures are matched on an increasingly large scale, leading to a smoother and more continuous visual experience. Thus, the visually most appealing images are usually created by matching the style representation up to the highest layers in the network (Fig 3, last row).

Of course, image content and style cannot be completely disentangled. When synthesising an image that combines the content of one image with the style of another, there usually does not exist an image that perfectly matches both constraints at the same time. However, the loss function we minimise during image synthesis contains two terms for content and style respectively, that are well separated (see Methods). We can therefore smoothly regulate the emphasis on either reconstructing the content or the style (Fig 3, along the columns). A strong emphasis on style will result in images that match the appearance of the artwork, effectively giving a texturised version of it, but hardly show any of the photograph's content (Fig 3, first column). When placing strong emphasis on content, one can clearly identify the photograph, but the style of the painting is not as well-matched (Fig 3, last column). For a specific pair of source images one can adjust the trade-off between content and style to create visually appealing images.

Here we present an artificial neural system that achieves a separation of image content from style, thus allowing to recast the content of one image in the style of any other image. We demonstrate this by creating new, artistic images that combine the style of several well-known paintings with the content of an arbitrarily chosen photograph. In particular, we derive the neural representations for the content and style of an image from the feature responses of high-performing Deep Neural Networks trained on object recognition. To our knowledge this is the first demonstration of image features separating content from style in whole natural images.

Previous work on separating content from style was evaluated on sensory inputs of much lesser complexity, such as characters in different handwriting or images of faces or small figures in different poses.

In our demonstration, we render a given photograph in the style of a range of well-known artworks. This problem is usually approached in a branch of computer vision called non-photorealistic rendering. Conceptually most closely related are methods using texture transfer to achieve artistic style transfer. However, these previous approaches mainly rely on non-parametric techniques to directly manipulate the pixel representation of an image. In contrast, by using Deep Neural Networks

trained on object recognition, we carry out manipulations in feature spaces that explicitly represent the high level content of an image.



Features from Deep Neural Networks trained on object recognition have been previously used for style recognition in order to classify artworks according to the period in which they were created. There, classifiers are trained on top of the raw network activations, which we call content representations. We conjecture that a transformation into a stationary feature space such as our style representation might achieve even better performance in style classification.

In general, our method of synthesising images that mix content and style from different sources, provides a new, fascinating tool to study the perception and neural representation of art, style and content-independent image appearance in general.

We can design novel stimuli that introduce two independent, perceptually meaningful sources of variation: the appearance and the content of an image. We envision that this will be useful for a wide range of experimental studies concerning visual perception ranging from psychophysics over functional imaging to even electrophysiological neural recordings. In fact, our work offers an algorithmic understanding of how neural representations can independently capture the content of an image and the style in which it is presented. Importantly, the mathematical form of our style representations generates a clear, testable hypothesis about the representation of image appearance down to the single neuron level. The style representations simply compute the correlations between different types of neurons in the network. Extracting correlations between neurons is a biologically plausible computation that is, for example, implemented by so-called complex cells in the primary visual system (V1). Our results suggest that performing a complex-cell like computation at different processing stages along the ventral stream would be a possible way to obtain a content-independent representation of the appearance of a visual input.

All in all it is truly fascinating that a neural system, which is trained to perform one of the core computational tasks of biological vision, automatically learns image representations that allow the separation of image content from style. The explanation could be that when learning object recognition, the network has to become invariant to all image variation that preserves object identity. Representations that factorise the variation in the content of an image and the variation in its appearance would be extremely practical for this task. Thus, our ability to abstract content from style and therefore our ability to create and enjoy art might be primarily a preeminent signature of the powerful inference capabilities of our visual system.

Methods

The results presented in the main text were generated on the basis of the VGG-Network, a Convolutional Neural Network that rivals human performance on a common visual object recognition benchmark task and was introduced and extensively described in. We used the feature space provided by the 16 convolutional and 5 pooling layers of the 19 layer VGG- Network. We do not use any of the fully connected layers. The model is publicly available and can be explored in the caffe-framework. For image synthesis we found that replacing the max-pooling operation by average pooling improves the gradient flow and one obtains slightly more appealing results, which is why the images shown were generated with average pooling.

Generally each layer in the network defines a non-linear filter bank whose complexity increases with the position of the layer in the network. Hence a given input image x is encoded in each layer of the CNN by the filter responses to that image.

A layer with N_l distinct filters has N_l feature maps each of size M_l , where M_l is the height times the width of the feature map.

So the responses in a layer l can be stored in a matrix F^l of $R^{N_l \times M_l}$ where F^l_{ij} is the activation ij of the i th filter at position j in layer l . To visualise the image information that is encoded at different layers of the hierarchy (Fig 1, content reconstructions) we perform gradient descent on a white noise image to find another image that matches the feature responses of the original image. So let p and x be the original image and the image that is generated and P^l and F^l their respective feature representation in layer l . We then define the squared-error loss between the two feature representations

$$\mathcal{L}_{content}(\vec{p}, \vec{x}, l) = \frac{1}{2} \sum_{i,j} (F^l_{ij} - P^l_{ij})^2 .$$

The derivative of this loss with respect to the activations in layer l equals

$$\frac{\partial \mathcal{L}_{content}}{\partial F^l_{ij}} = \begin{cases} (F^l - P^l)_{ij} & \text{if } F^l_{ij} > 0 \\ 0 & \text{if } F^l_{ij} < 0 . \end{cases}$$

from which the gradient with respect to the image x can be computed using standard error back-propagation. Thus we can change the initially random image x until it generates the same response in a certain layer of the CNN as the original image p . The five content reconstructions in Fig 1 are from layers ‘conv1 1’ (a), ‘conv2 1’ (b), ‘conv3 1’ (c), ‘conv4 1’ (d) and ‘conv5 1’ (e) of the original VGG-Network.

On top of the CNN responses in each layer of the network we built a style representation that computes the correlations between the different filter responses, where the expectation is taken over the spatial extent of the input image. These feature correlations are given by the Gram matrix G^l of $R^{N_l \times N_l}$, where G^l_{ij} is the inner product between the vectorised feature map i and j in layer l :

$$G^l_{ij} = \sum_k F^l_{ik} F^l_{jk} .$$

To generate a texture that matches the style of a given image (Fig 1, style reconstructions), we use gradient descent from a white noise image to find another image that matches the style representation of the original image. This is done by minimising the mean-squared distance between the entries of the Gram matrix from the original image and the Gram matrix of the image to be generated. So let a and x be the original image and the image that is generated and A^l and G^l their respective style

representations in layer l . The contribution of that layer to the total loss is then

$$E_l = \frac{1}{4N_l^2 M_l^2} \sum_{i,j} (G_{ij}^l - A_{ij}^l)^2$$

and the total loss is

$$\mathcal{L}_{style}(\vec{a}, \vec{x}) = \sum_{l=0}^L w_l E_l$$

where w_l are weighting factors of the contribution of each layer to the total loss (see below for specific values of w_l in our results). The derivative of E_l with respect to the activations in layer l can be computed analytically:

$$\frac{\partial E_l}{\partial F_{ij}^l} = \begin{cases} \frac{1}{N_l^2 M_l^2} ((F^l)^T (G^l - A^l))_{ji} & \text{if } F_{ij}^l > 0 \\ 0 & \text{if } F_{ij}^l < 0. \end{cases}$$

The gradients of E_l with respect to the activations in lower layers of the network can be readily computed using standard error back-propagation. The five style reconstructions in Fig 1 were generated by matching the style representations on layer ‘conv1 1’ (a), ‘conv1 1’ and ‘conv2 1’ (b), ‘conv1 1’, ‘conv2 1’ and ‘conv3 1’ (c), ‘conv1 1’, ‘conv2 1’, ‘conv3 1’ and ‘conv4 1’ (d), ‘conv1 1’, ‘conv2 1’, ‘conv3 1’, ‘conv4 1’ and ‘conv5 1’ (e).

To generate the images that mix the content of a photograph with the style of a painting (Fig 2) we jointly minimise the distance of a white noise image from the content representation of the photograph in one layer of the network and the style representation of the painting in a number of layers of the CNN. So let p be the photograph and a be the artwork. The loss function we minimise is

$$\mathcal{L}_{total}(\vec{p}, \vec{a}, \vec{x}) = \alpha \mathcal{L}_{content}(\vec{p}, \vec{x}) + \beta \mathcal{L}_{style}(\vec{a}, \vec{x})$$

where α and β are the weighting factors for content and style reconstruction respectively. For the images shown in Fig 2 we matched the content representation on layer ‘conv4 2’ and the style representations on layers ‘conv1 1’, ‘conv2 1’, ‘conv3 1’, ‘conv4 1’ and ‘conv5 1’ ($w_l = 1/5$ in those layers, $w_l = 0$ in all other layers). The ratio α/β was either 1×10^{-3} (Fig 2 B,C,D) or 1×10^{-4} (Fig 2 E,F). Fig 3 shows results for different relative weightings of the content and style reconstruction loss (along the columns) and for matching the style representations only on layer ‘conv1 1’ (A), ‘conv1 1’ and ‘conv2 1’ (B), ‘conv1 1’, ‘conv2 1’ and ‘conv3 1’ (C), ‘conv1 1’, ‘conv2 1’, ‘conv3 1’ and ‘conv4

1' (D), 'conv1 1', 'conv2 1', 'conv3 1', 'conv4 1' and 'conv5 1' (E). The factor w_l was always equal to one divided by the number of active layers with a non-zero loss-weight w_l .