

# The AI Disruption is Here

June 8, 2024

DOUGLAS MCLENNAN: First of all, are any of you in here recording this session on your phone? Okay. And if you were, what you would do is take the recording, and then afterwards you would tell the AI, summarize everything that Doug said. If he said anything interesting, just put it into a form that I can use, right? So in the future, this is going to be the way that you're going to be able to interact with people and be able to remember what it is that's important and how you want to use it.

So as Simon said, I'm a musician and journalist and strategist in communications in digital technology. I see my niche as at the intersection of arts, journalism, and technology. I'm also, as he said, the — hold on here. Advance, okay. I'm also the editor of Arts Journal. And every day we scan thousands of stories from all over the world and collect those that speak to issues of the day in the world of arts and culture. I started Arts Journal 25 years ago, and since then, we've day by day chronicled a change in the culture during a truly transformational time. I've never really thought of Arts Journal as a news source, but more a lens through which to work — to look at the world. There are many lenses you can use. Politics, business sports. I prefer culture. And Arts Journal aims to give you a sense of the issues and context of how culture is evolving.

So artificial intelligence. What I'd like to try and do is give you an overview of the ways that it might impact culture broadly, more specifically orchestras. And then finally, even more specifically, how you personally might want to use it in your work. You're probably sick of the letters A and I by now. The hype is intense. Every product you encounter these days seems like it has to have AI in it.

[0:02:09.0]

Yet, for many of you who no doubt have tried it and tried to figure out how you might use it in your work, AI is still more promise than practicality. Given the formidable pressures of trying to just keep the doors open, who has time or resources to explore? You're looking for things that you can use now. And the good news is that those things already exist. There are many tools you can use to become more efficient, extend your reach, and make it possible for you to do things at a scale that simply wasn't possible before.

520 8th Avenue, Suite 2005, New York, NY 10018

[americanorchestras.org](http://americanorchestras.org)

The better news is that AI can make tiny organizations super powerful. Sam Altman, founder of OpenAI says that it's only a matter of time before the first billion-dollar startup with only a single employee is created. That's the scale that we're talking about. But using AI requires some adjustment to our mindset. We need some ways to process and think about what's about to happen. And that's what I hope to do.

There's lots of hyperbole and lots of embarrassing AI fails, as Simon said, that add glue to your pizza. I'm not an AI utopian, I'm not an AI doomer. But I do think that there are real threats to — downsides for perhaps millions of us who will need to figure out how to adapt. But I also see in all this some potentially transformative upsides that will benefit us all. So today, I want to focus on culture, on the arts, and most specifically orchestras. And I hope you'll come out of the session with four things.

[0:03:49.6]

First, my best attempt is to suggest a conceptual framework thinking about what AI offers us, and why it's so revolutionary. Second, to suggest some ways AI will change us creatively or artistically, our ways of doing business. And then finally, some policy issues that we have to work out. And lastly, I'd like to offer some things that you can do right now to start thinking about the opportunities and, just important, prepare for the threats. I want to stress; I think AI is going to radically and fundamentally change the ways that our culture and our economy works.

We don't have the choice to ignore it. Think about how incredibly different our world is now from pre-internet days. Technology changed, to be sure, but the more transformational thing was culture. How we communicate, what we pay attention to, how we get information, who we trust. The iPhone has only been around since 2007. Think about the transformation just since then.

I think we can learn some important lessons from the digital revolution. I think we need to learn those lessons. Because I think many of the current ills in our culture and in our institutions, and particularly for the health of arts and artists, can be traced back to mistakes we made, and are only really now recently understanding as we document the harms to mental health, particularly among young people, and to the institutions that support a healthy society.

I believe that, in all revolutions, there are winners and losers, and standing on the sidelines is not an option. I prefer to find the opportunities and mitigate the threats. If, as many are saying, this revolution in machine learning is on the scale of a new industrial revolution, and it looks to me like it is, we need some context, some ways to process it, give it shape. I'm going to suggest three of them here. First, we need to think about a different scale of change. We think things move fast now.

520 8th Avenue, Suite 2005, New York, NY 10018

[americanorchestras.org](http://americanorchestras.org)

Moore's Law, Intel's founder, Gordon Moore's prediction in 1965, the computer processor speed would double every 18 months, is tough enough to keep up with. Think about Apollo 11 going to the moon. The computer processor power in that is less than a fraction of what we have in our typical smartphones that we carry around with us.

[0:06:29.0]

Processing speed is about a million times more powerful today. And just as significantly, the cost of that processing is about a million times cheaper than it was. The thing is, AI compresses Moore's Law timeframe, and instead of doubling it every year or so, change difficult enough to process but comprehensible, a new exponential change cycle kicks in now every month or every week, and it's difficult for us to even comprehend. So think about this for a second. Thirty steps on this stage. I can kind of predict where I'm going to be in 30 steps. But if each of those steps is exponential, I don't land at the next room or even outside the conference center, I end up on the moon. How do we even comprehend something like that?

ChatGPT-3 couldn't pass the SATs. ChatGPT-3.5 passed the Bar exam, beating 90 percent of the humans that take it. We're now at ChatGPT-4. So I want to suggest that a real frame for thinking about what AI does, what its real significance, here is as a universal translator. This is Javier Milei. Last December, he got elected President of Argentina, and in January he went to Davos to deliver a speech.

JAVIER MILEI: [RECORDING] When we embrace freedom in 1860, in 35 years, we became the world's first dominant power. Thirty-five years we became the first world power. While, when we embrace collectivism, over the past 100 years, we saw how our citizens began to systematically impoverish themselves until they fell to the 140th position in the world.

[0:08:25.3]

DOUGLAS: Thing is, he's speaking Spanish. This is doing simultaneous translation into about a dozen languages. In his own voice, with his own inflections, right? This is transformational. AI — at its developer conference just a couple of weeks ago, Google unveiled a reportedly not ready for primetime version of this for all of us to use. Now think of the cultural implications. Anyone will now be able to talk to anyone in any language and have a conversation. Currently, only three percent of books in languages not written in English are translated into English. Now everything is accessible. And just like that, the Tower of Babel disappears.

520 8th Avenue, Suite 2005, New York, NY 10018

[americanorchestras.org](http://americanorchestras.org)

But let's carry this further. Movies can be produced in every language without subtitles or dubbing, using the actor's own voices, their faces synced to whatever language they're speaking. The same with music, books, theater. You thought globalism was dead. AI opens up culture from every corner of the earth to make it accessible in a way that it never was. This will create as many problems as it solves, of course, but think about the shifting culture and the change in mindset that it will require in the ways we deal with one another. But I think it goes even further.

[0:09:53.7]

The real promise of the revolution in AI from my perspective is that, for the first time, we have a universal translator that not just translates human languages but translates everything into language that we can understand. So I think the real frame is this. To an AI model, a picture is data, sound and music are data, as is traditional spoken or written language. That data is translatable, interchangeable, and most important, linkable and actionable. That means for the first time that video, music, sound, movement, image can interact in common language.

Now, why is that important? Until now, in order to find music, for instance, we had to tag it with text in what we call metadata. And our metadata was always an imperfect description of a piece of music, composer, title, artist, etcetera, using text. To find music we wanted, we could only search for the tags. But if every medium is data, and it translates in common language, instead of depending on tags or descriptions, we'll be able for the first time to see into music, or video, or pictures, for exactly what we want.

For example, finding a specific phrase or tambour or chord progression and a piece of music changes our ability to find and interact with the creativity around us. In a way, until now, maybe we've been seeing or hearing in 2D. And all of a sudden, we're in a multidimensional space in which music, sound, picture and image interact in new ways. Beyond media, this revolution allows data across medicine, climate, biology, science, to interact in new ways. And going even further, descriptor data can now link with behavior data, as autonomous vehicles now do.

[0:11:54.9]

But think about it in the context of art or music interacting with the audience or potential audience for that art. Nvidia just announced this week that it created a data model of the Earth's weather that will enable forecasts that are a thousand times better than the mathematical models that are now used for producing weather forecasts. The implications for everything from agriculture to predicting natural disasters are revolutionary. Talk about a change in mindset.

520 8th Avenue, Suite 2005, New York, NY 10018

[americanorchestras.org](http://americanorchestras.org)

The third concept I want to talk about is the notion of synthetic artists, synthetic music, and synthetic realities. The kind of AI we're talking about is generative AI. You're already well familiar with what are called traditional AI. Spell checkers, streaming, social media algorithms, pavement centers — sensors that change the traffic lights, etcetera. But generative AI actually makes things. Large language models and neural networks train on trillions of bits of data, studies the rules of how that data interacts, then uses those rules to create new things. Say, read every Hemingway novel, learn the style, and turn out new Hemingway stories. But not just one, maybe 100 or 200 before lunch. And it accepts input, so you can suggest a plot or a character or locations. And Hemingway bot will happily create.

So a kind of a collaborative partner for you, the audience member. It's already happening in music, where platforms like Spotify report they take down hundreds of thousands of synthetic songs every day. You can see the problem. AI floods the market, optimizing the curation algorithms of a platform like Spotify. And it vastly dilutes the ability of musicians to be found, and importantly, to make a living. And it potentially transforms the extraordinary into the mundane.

[0:14:05.3]

The commonplace. In a culture built around the extraordinary, identifying what the extraordinary thing is. What happens when everything is at a peak? Our definition of extraordinary will by definition have to evolve. Hold on. Sorry, I think I went back here. Okay. One of the things AI is brilliant at is having comprehensive body of specialized technical knowledge at its command. Where it takes humans a career of experience and education to develop specialized skills for a task, machines have already trained on and acquired them.

For example, if you're a composer, you want to write a piece of music for orchestra, you need years of study and work and a ton of specialized knowledge. How to read music, understand style, design rules, chord structures, melody, how to notate, orchestration, technical capability of every instrument, and know how they sound together. And then you have to hire musicians, book a studio or concert hall, and hope that you can draw an audience. If it all works really, really, really, really well, maybe you can have a career.

How long does it take from having the notion you might want to be a composer to having your first music played by an orchestra? Five years? Ten? Fifteen? And even after you've had all that training and experience and you're well established, how long does it take from when you have an idea for a piece in your head until you actually hear it on a stage? A month? A year? Five years? Between the

520 8th Avenue, Suite 2005, New York, NY 10018

[americanorchestras.org](http://americanorchestras.org)

notion of composing to actually being a composer who makes music, how many technical hurdles along the way frustrate or discouraged or dissuade somebody from actually realizing it?

[0:16:06.3]

If AI is really a universal translator, then instead of having to learn how to notate, I can describe exactly what I want in language that I speak. Instead of needing to know orchestration, I can build a piece around sound by sound by telling the AI what I want. And instead of translating what I hear in my head to imprecise traditional notation, kind of code, I can endlessly describe and tweak and create the perfect realization of what I hear in my head. And if then I want a real orchestra to perform it, I can print it out, and they can know exactly what I intended.

Now, instead of, what, 5,000, 10,000, maybe even 100,000 composers who are capable of composing for an orchestra, now there are millions. And sure, most of what they make will be dreck, just wretched. But then again, if history is any guide, this will have a transformative effect on the ways that the art form will evolve. Just look at the impact of young musicians being able to learn from clips on YouTube. The average technical facility of today's musicians easily beats those of my generation. Photoshop turned anyone into a sophisticated photo editor.

I've been playing with a new music generator called Udio. It launched at the beginning of April this year, maybe six weeks ago. And there are now several of these AI music generators. But I can describe music I want to create, and within seconds Udio creates music, offering me ideas, which I can modify or re-modify with voice prompts. It's really rudimentary at the moment, but again, thinking exponentially, it's easy to see how sophisticated this will quickly become.

[0:17:55.0]

OpenAI launched Sora in February, which does the same thing for creating video and blew everybody away, as did Dali, and Stable Diffusion for images last year, and ChatGPT for text in November of 2022. If you're a traditional orchestra, how are you going to position yourself? Is this a threat? Or is it an intriguing possibility that revitalizes the art form? I think it's actually the ladder.

Here's another scenario. My AI has just ingested the entire catalog of every orchestral recording ever, including broadcasts. Scratch that, every performance of classical music ever. It now understands not only how the music is put together and the styles of composers, but it also recognizes stylistic differences between one performer and the next, one orchestra and the next, and the various conductors who lead them. And now it can create new performances based on all of

this. This is, by the way, already happening in pop music, where there are synthetic convincing Taylor Swifts and Drakes all over the internet.

So now I'd like to hear that famous Leon Fleisher, George Szell 1958 performance of Brahms D Minor Concerto, not with that Cleveland Orchestra, but with the Cleveland Orchestra of today. How might it be different? What might I learn? Or maybe I'd like to create an entirely new performance using Chicago's brass, Philadelphia strings, and Makela conducting. Think this won't happen? It isn't yet, and I've actually tried. [LAUGHING] But in the next year or two, absolutely.

One can imagine synthetic soloists, virtual Lang Lang's and Yuja Wangs out on tour, playing real instruments and interacting in real time with orchestras. Or perhaps an entirely synthetic soloist, for whom no technical feat is an obstacle. The BBC Symphony Orchestra has already released a full library of samples of all of its instruments free for anyone to use, with a DAW, it's a digital audio workstation. Anyone with Ableton or Cubase or Pro Tools can write sophisticated orchestra pieces. Throw in a performance touch sensitive fader, and you can create and perform complex orchestral scores in real time.

[0:20:17.9]

DAWs are astonishing, allowing you to create, in anything, sounds that you want, and process them in any form. Now put a generative AI on top of that, and you don't even have to know how to work the DAW software. Just describe what you want and what instrumentation. I'd love to show you how this might work in music, but we're still in early days of coupling chatbots to music. Instead, I want to show you a video that Google released two weeks ago that demonstrates the current state of the art in speech of the kind of thing that I'm talking about. The AI only needs a two second sample of anyone's voice to be able to credibly replicate it. And again, remember the exponential curve and process. This is, after all, a Google demonstration and typically really, really cheesy.

[VIDEO PLAYS]

SPEAKER: [IN PROGRESS] doing?

SPEAKER: I'm doing fantastic. Thanks for asking. How about you?

SPEAKER: Pretty good.

520 8th Avenue, Suite 2005, New York, NY 10018

[americanorchestras.org](http://americanorchestras.org)

SPEAKER: What's up?

SPEAKER: So my friend, Barry here, he's been having trouble sleeping lately. And I want you to tell him a bedtime story about robots and love.

SPEAKER: Oh, a bedtime story about robots and love? I got you covered. Gather round, Barrett. Once upon a time, in a world not too different from ours, there was a robot named Bite. Bite was a curious robot, always exploring —

SPEAKER: I know you started this story, but I want a little bit more emotion in your voice. A little bit more drama.

SPEAKER: Got it. Let's add some drama. Once upon a time, in a world not too different from ours, there was a robot named Bite.

SPEAKER: No, no, no, ChatGPT, I really want maximal emotion, like maximal expressiveness, much more than you were doing before.

[0:21:59.9]

SPEAKER: Understood. Let's amplify the drama. Once upon a time, in a world not too different from ours, there was a robot named Bite. Bite was —

SPEAKER: Can you do this in a robotic voice now?

SPEAKER: Initiating dramatic robotic voice. Once upon a time.

[VIDEO ENDS]

DOUGLAS: Okay. [LAUGHING] So you can see that already that the kind of voice sampling is really sophisticated. And for music it's getting to be that way too. Because remember, everything is just data. For artists, this opens new questions about what they're creating and on what terms. If before,

520 8th Avenue, Suite 2005, New York, NY 10018

[americanorchestras.org](http://americanorchestras.org)



artists created performances, now they have to think about their unique essence. If an AI can be your surrogate, something you create that extends you, what could you co-create with that maybe even your audience or other artists collaborate with?

It opens all sorts of possibilities for what you can do creatively and what your career might look like. One artist who has been thinking about this a lot is Holly Herndon, an American composer who lives in Berlin. She created an AI that she calls Holly Plus, which she trained on her own voice. She then made Holly Plus available for others to use. And Holly Plus sings in languages she doesn't speak, can perform technical melismas that she can't do. Here she explains.

HOLLY HERNDON: [RECORDING; IN PROGRESS] Holly Plus to create a wide range of vocals that I didn't sing from a set of recorded phrases that I did sing. I like to think of spawning as a kind of 21st century corollary to the musical tradition of sampling, which had a really big impact on both music and intellectual property. But I think spawning is far more exciting and potentially really weird.

[0:23:50.6]

So for example, with sampling, usually you copy and remix a recording by someone else to create something new. But with spawning, you can perform as someone else based on trained information about them. And as an artist, this is making me rethink my own past work, as not only my archive, but potentially also I myself could become reanimated with AI. This also opens up the question of how we deal with a collective human archive. If we can reanimate old media. It opens up really big ethical and intellectual property questions that require entirely new conceptual and legal frameworks.

One way that I like to think about intellectual property I call identity play. So rather than limiting the use of my voice, I'm creating instruments to allow as many people as possible to create music with me and even as me. That's why I've made versions of Holly Plus freely available for anyone to use online.

DOUGLAS: Okay, so, Holly calls the technique tambour transfer, in which the essence of one artist is kind of used as a graft on another. And for an example, here's a short clip of FUR singing as himself through one mic, and Holly Plus in the other.

[VIDEO PLAYS; SINGING]

520 8th Avenue, Suite 2005, New York, NY 10018

[americanorchestras.org](http://americanorchestras.org)

DOUGLAS: Holly rejects the term AI and prefers the notion of collective intelligence, a term I'll get back to in a few minutes. So what kinds of hybrid orchestral collaborations might become the norm? And how will this extend and transform the art form? So now we come to the business part of this, because it's going to transform the way that you do business out in the world. The digital age has been really, really bad for creative people. Brutal for arts organizations like orchestras. Traditional arts organizations have been trapped in cost disease in which the cost for performances like orchestra concerts continue to go up, while in the digital space, distribution costs fall to almost nothing.

[0:26:28.7]

The online audience is scaled to massive numbers, millions, hundreds of millions, even billions, while traditional performing arts see their ability to reach through markets flooded with cheap, distracting content. Increasingly difficult. The problems are so self-evidently intractable that most of the major foundations and corporations that used to routinely support traditional art forms have pulled away. Yes, they're focusing more on social issues. But at least some of them have concluded that traditional arts institutions have a viability problem.

I attribute much of the bad state we're into the insertion of big tech as middlemen who exact heavy tolls on creative work of all kinds and have captured the basic societal utilities upon which the modern economy runs. Forty-five years of no enforcement of antitrust has shriveled choice and competition and has concentrated online markets to a few players, with practically no afterthought on big tech part. The cultural infrastructure that supported traditional arts, culture, and journalism was undermined and swept away.

I can't stress this enough. Big tech inserted itself between creators and their audiences, charging significant tolls and making it impossible to not, if not impossible, for traditional culture to be financially viable. But I think there's actually a glimmer of hope. With recent suits brought — brought against Google, Apple, Amazon, Facebook, and now Ticketmaster, and Live Nation, a rejuvenated Biden FTC under Lena Khan has begun the arduous work of breaking up some of these monopolies. The decaying of social networks, which have — have, as the writer Cory Doctorow observed, en-shit-ifies the internet with ads, bots and trolls, offers opportunity to rebuild the creative marketplace.

[0:28:35.3]

And maybe, just maybe, AI has a real opportunity to reinvent the business infrastructure. So an example. In the 1950s, if you looked at big orchestra charts, philanthropy departments barely

520 8th Avenue, Suite 2005, New York, NY 10018

[americanorchestras.org](http://americanorchestras.org)

existed, marketing departments were thin, and educational programs were modest. Today, as ticket sales cover less and less of the typical budget, marketing and development is a main — development is a major part of the operations of any orchestra. And yet, the reach of orchestras into the general population, the ability to get in front of highly fragmented, highly segmented market, has gone down.

The typical arts organization is plagued by a lack of resources and the ability to reach potential audience. The typical big brand wouldn't dream of launching a product without marketing budgets that typically outstripped what it costs to make the movie or cereal or car, to do otherwise would all but guarantee failure. In the arts, though, marketing and engagement capacity can barely get in the game. We confuse the power of quality with the realities of structural access to the market that has just gradually shut us down. That's a fatal mistake.

[0:29:58.3]

Customers are now talking about something called exponential organizations, in which employees aren't hired for particular specialties, but for their ability to be creative journalists — generalists, pardon me. I've got journalism on the brain. Let me explain. So today's workplace is organized around specialties. Employees who are designers or fundraisers or marketers or data analysts or bean counters. Each of these jobs depends on a body of knowledge and experience. But seeing inside that knowledge is real work.

The average white-collar employee in America spends 32 days a year on data retrieval of one kind or another. A significant cost. And the data is almost always inexact or incomplete, and almost always not of the moment. Companies are starting to build what are called digital twins. These are data models of an organization that measure every aspect, not only of the company itself, but relevant contextual data about the market in which you operate and about your customers.

Because everything is data for an AI, nothing has to be in structured databases, which means that all the data can talk to one another and learn and evolve in real time. And AI agents can spin up bots on the fly to query and get answers and solve problems. Think of the dashboards you use to monitor ticket sales or marketing or schedules or audience behavior. And then imagine a model that not only is able to interact, but to model and align behaviors with strategies.

Instead of thinking of software as tools, think of AI as skills which interact dynamically, with agents calculating tasks, then calling upon the appropriate expertise and tools to solve or implement. Organizations go from being black boxes that require real expertise to see inside, to digital models

520 8th Avenue, Suite 2005, New York, NY 10018

[americanorchestras.org](http://americanorchestras.org)

that you can interact with, try out scenarios on, and react to in real time. Such models will ultimately be networked to be able to anticipate market shifts, economic forces, and importantly, understand in detail taste, shifts in taste, and consumer behavior.

[0:32:26.0]

Use cases abound. One big shipping company used AI plugged into real time weather and current information to steer the ships autonomously as they go around the world, saving about \$2 million a year per ship in fuel costs. That's big savings. Every one of us now carries personal sensor devices that record where we are, what we're doing, and what we pay attention to. The cost of getting in front of a customer or a potential customer at any point in the day, knowing enough about that person to know what and where to offer them, is not 10 years off, it's now.

I talked about, earlier, the idea of AI being a universal translator, and of a shift from needing to tag in metadata, to being able to directly able to compare the data itself. But that doesn't mean that descriptors or metadata goes away. But for the first time, we have the ability for the creative metadata to be able to interact and integrate directly with the user metadata even in real time. To personalize audience experience inside a concert hall or lobby based on what the audience member needs. To provide contextual assistance, to interact with customers as they're buying their tickets, choose the best experience.

In the future, to interact with your audience's digital twin assistant to alert them when something is important to them. Or even use facial recognition to figure out where in a performance you're actually registering with your audience. The possibilities to extend your reach is not in a scattershot, hit or missed broadcast way, or something dependent on someone else's opaque algorithmic platform, but to be able to simultaneously have one on one interactions with thousands of potential audience members.

[0:34:15.2]

The revolution in marketing is — for example, isn't the opportunity to craft a thousand different mass — messages for a thousand different people, but to be able to interact, engage, and persuade a thousand customers on their own terms, not just feed them unwanted messages. There's a cliché about us shaping our tools and then our tools shaping us. The notion that we address problems based on tools we have created. So a hammer is everything — to everything looks like a nail.

So think of institutions as tools. We shape them in the form of the kinds of things we want them to do, but the more complex the problems, the more layers an institution grows, the further away it can get from the ultimate thing that it's trying to do. The typical institution has silos stuffed with expertise, but the silos often don't talk to each other. AI turns that traditional organization chart sideways. Suddenly, we don't need specialties. We don't need number crunchers or marketing analysts. AI can support all of that kind of activity.

But that doesn't mean we don't need employees. It allows us to redeploy talent and right-size for our own resources. Instead of needing experts with technical skills to separate and bury in silos, we'll need creative generalists. This shifts our work from skills to projects. That's a profound shift.

With that in mind, how might you change the tool of the institution in getting things done? How might the new tool change what the institution is able to do and support its core functions? Let's go even further. How might the institution go from being a tool setup to dynamic — of dynamic skills that interact with its — to one that goes to a dynamic tool that interacts with its musicians and its community at radically lower cost? What if the core mission of the new institution shifted from not only being — building a better performance, a better concert, but building a better audience?

[0:36:27.5]

I think it might be transformational. For much of the last year, I've been running around the country trying to talk to policymakers to pay attention to this in the arts. So far, I've had no — no takers, but I fully understand the obstacles. So I'm going to make the case to you. Maybe it resonates. We got killed by the digital revolution in the past 25 years. In the absence of smart thinking about what was at stake, a new industrial revolution, no less, the new players who understand what's happening and think strategically become our overlords. In the case of the shift to digital, big tech heavily subsidized by cheap money and corporate investors, decimated existing infrastructure and inserted themselves between creative industries and their audiences, becoming toll collectors.

As they consolidated into massive monopolies, they have systematically starved culture that doesn't fit their definitions of algorithmic success. Changing opaque rules on a whim and repeatedly defining the only culture that matters as maximal distraction. The result is an impoverished civil discourse, a coarsening of our public life, and politics, and the marginalization of art and artists. The lesson. In the absence of engagement by creative industries, big tech writes its own rules, and the rest of us have to live under them, no matter how much they are to our disadvantage.

[0:37:59.0]

520 8th Avenue, Suite 2005, New York, NY 10018

[americanorchestras.org](http://americanorchestras.org)

Some examples. When music metadata was first being worked out in the '90s, the structure was built around pop music. Classical music metadata didn't fit pop music. Making the logic of how actual listeners go out and find it very, very difficult. It wasn't until 2015, when Adagio came along and wrote a new schema, and now people can find music in the way that they want to. But in the meantime, a whole generation of listeners disappeared because of all these digital speed bumps that were erected.

Ninety percent of all web traffic goes through Google Search. So being found on Google is not just important, it's existential, and the difference between being viable or being invisible is not ranking high enough on the Google algorithms. Being on page 20 guarantees that you're invisible. And so there's a whole industry of search engine optimization built around helping you do that.

But Google changes those algorithms in a blink of an eye. A couple of months ago, they changed the algorithm, and overnight, hundreds of thousands of websites saw their traffic declined by 90 percent. Some of these sites had been around for 20 years. And chilling stories are no less true, but the ability to find culture, it's just that the poverty of that infrastructure largely happened while no one was looking. And we keep thinking that it's a content quantity — quality issue, or an engagement issue, when it isn't.

[0:39:42.0]

So what are the policy issues we ought to be paying attention to? And I realize you're all struggling just to keep the doors open and deal with today's challenge. That's why this is difficult. So some issues right now. Copyright. You might have heard that all sorts of artists are suing. But copyright rules were never intended for AI. And so these artists are going to lose, right? And right now, there are dozens of groups in Washington D.C. lobbying for what the new copyright rules are going to be.

Are orchestras involved with that? Are artists involved with that? Not to any meaningful extent. So how do you copyright the essence of an artist? These are issues that are going to have to get worked out. And unless we're part of that conversation, it's going to happen in ways we might not like. A few weeks ago, AI started offering AI answers to questions when you search, right?

As a user, this is actually pretty cool. It saves you a whole lot of time. Instead of sending you to a website, it sends you the answer, right? But the problem is, is that if you don't ever have to visit the website, then the people making the content, which the AI is getting the answers from, have no way of being able to monetize it. And so what happens to the web? It collapses. Same thing is going to

520 8th Avenue, Suite 2005, New York, NY 10018

[americanorchestras.org](http://americanorchestras.org)

happen, the ability for orchestras and artists to be found online. This is profoundly important, and people aren't talking about it enough.

Synthetics. To go back to my earlier example of a listener calling up synthetic performances. What are the issues around using synthetics? Who owns them? Who owns the content that they make? How does that interact? How do we exchange that? What's the ownership of a piece of culture? These are all important issues that we haven't yet dealt with. Each of these issues deserves a talk all on its own.

When we started, I promised some ideas on what you might do right now. If you had real resources lying around, I'd say start engaging AI people. But I mean, start making relationships with experts working and thinking about AI, so you can begin to figure out how you might be able to use it. Almost all Fortune 500 companies now have not only chief AI officers in their c-suites, but AI departments beginning to incorporate AI on all their product workflows. Why? Because the cost savings are enormous. The gains and productivity and potential markets are significant. Your goals may be more modest, but in time, not far off, AI is going to be part of every transaction and business and artistic relationship you have.

[0:42:42.4]

You can start to understand and prepare for it now, or struggle to keep up later. But since you don't have endless R&D resources, and your capacity for risk is low, I have a few suggestions. One, don't get distracted by shiny AI toys. You'll be seduced and it will lead you down unproductive paths. Two, invest in your mindset. You need a philosophy, a point of view, informed by your values and what you are trying to accomplish. But you also need to have an understanding of the possibilities. Now, more than ever, it's important to know why you exist, and what it is that you're trying to accomplish.

Three, problems first, solutions later. Don't reach for solutions that you don't have problems for. That happens a lot in any kind of technological revolution. Four, start playing. The way to get comfortable with AI is to play with it. You will find things that are useful. More important, it will expand your worldview and how you might interact with others. You'll learn to speak AI. And five, still unsure of how you might use AI? Ask it, be as specific as possible, engage it in conversation. It will suggest ways in which it might help you.

[0:44:09.0]

520 8th Avenue, Suite 2005, New York, NY 10018

[americanorchestras.org](http://americanorchestras.org)

Just ask it a simple question. I'm an orchestra of this size, I have this, how could you help me? It will give you suggestions, believe me, things that you didn't think about. So that's actually a really cool thing to do. And lastly, as you may have noticed, I love talking about this stuff, and I'm eager to learn more. So if you want to chat, or ask a question, or maybe even argue with me about some or many of the things I've said, drop me an email at [McLennan@ArtsJournal.com](mailto:McLennan@ArtsJournal.com), and I'll have my AI talk to your AI.

And in the meantime, go to the website and sign up for my newsletter. I write about all of these issues and will continue to do it more or less weekly going forward. So thank you. I realized that was a lot of information and a lot of ideas in a very short time. I appreciate your patience. Thank you.

[APPLAUSE]

SPEAKER: Okay, that was amazing, Doug. Thank you very much. So do we have some questions? Can we take some questions? I've got a microphone here, and I'm going to move it around. I'm going to be the roving mic. Mark Pemberton has a question.

MARK PEMBERTON: Hello, yes. I'm Mark Pemberton, come all the way from London to listen to this excellent presentation.

DOUGLAS: Thank you.

MARK: And it's not a question, but it's more to develop your discussion point around legislative frameworks, regulations. And I accept that they're not — there are — your political system will make it hard perhaps to create a regulatory framework because of the influence of these big tech companies. But the European Union may ride to the rescue. It is very clearly going to implement a legislative framework where copyrighted material must be credited, artists must be remunerated.

[0:46:10.6]

And because at the moment it is wild west, and it is simply harvesting copyrighted material with no controls whatsoever. Now, whether the EU in itself can, therefore, hold back the tide, while the rest of the world does not introduce regulations, is open to question. But I think that it is recognizing that outside of the USA there are other political entities that could create those controls.

520 8th Avenue, Suite 2005, New York, NY 10018

[americanorchestras.org](http://americanorchestras.org)



DOUGLAS: Yes, I'm very conflicted about the EU. The EU has passed a number of measures to protect copyright and, you know, protect privacy. I view some of them as very heavy handed, and frankly, not that effective. And so I worry about a system in which we try and legislate from top down to — to essentially, like, create a regulatory framework for this. Because we don't actually — like as I've been trying to explain, the mindset changes, right? Like, the way we have to think about this sort of changes.

And often those kinds of regulations are built on the last set of things that — that didn't go so well for us. And so the top-down kind of things, I worry about those a lot. That said, it has certainly got big tech's attention. And that's a — that's a really good thing. But I do think that we have to be very careful incrementally. I'm not arguing from a policy standpoint that we have to figure it all out, and that we have to have a map, and then here are all the rules of the road.

[0:47:59.1]

But for instance, how we didn't enforce our antitrust, and let these companies be so large and buy up all their competitors, and all the adjacent businesses, that has led to all sorts of just really horrible things. We have a lot of laws that could apply to this already. So I agree with you. I think — I think the EU is helping enormously. But even those things, like we're going to charge Google for putting new stories on their — that show up in their search things. It's a very, very bad idea. Because it's an aftermarket tax on something that, frankly, they don't really care about, right? But if instead, you reformed the ad market that they also own, which is an absolute monopoly, both for the buyer and the seller and the marketplace, then suddenly that problem goes away in a significant way. So it's not — it's not easy.

SPEAKER: I'm coming over there. It's a big room.

RUTHIE: Hi, my name is Ruthie, I'm an artist. I am very curious to hear about how you see these AI tools as a way to generate new art and champion new artists. Like it's so cool to imagine booking synthetic Yuja Wang, but how does the new Yuja Wang find a path into the industry when everybody can book synthetic Yuja Wang, and everybody can use free Holly Plus? So I would love to hear ideas for how this can create a vibrant young community of artists as opposed to a world where we all see synthetic ABBA, you know, for the rest of our lives. So I'd love to hear your thoughts on that.

[0:49:49.8]

DOUGLAS: Right. So I think anytime you expand the tools, anytime you make things possible that

520 8th Avenue, Suite 2005, New York, NY 10018

[americanorchestras.org](http://americanorchestras.org)

weren't possible before, like, you know, I have a terrible visual sense, but put me in Photoshop or something and suddenly I look like a world class designer, right? And — and so, you know, I have a terrible singing voice. I have lots of ideas for songs, right? But if I could create a voice that would be sort of synthetic, but that is of my own creation that I could shape, then all of a sudden, you freed me to be able to think creatively about how to make songs, how to perform them, how to interact with other people, right?

So putting tools in the hands of artists is a way of expanding their ability to do things that only so far could be in the imagination. When you — when you give these tools to everybody, you're going to get a whole lot of really unimaginative kind of, you know, ho-hum stuff. But people who have inspiration artists are going to use those tools and create things that we can't even imagine right now, in ways that we don't know. So I'm very excited about, you know, where this takes us, how it evolves.

Like the issues that we have right now are not the issues we're going to have the day after tomorrow, right? Some of the creative technical issues that we have to cope with right now, that will solve them, but then there will be others that will show up. So I think that this gives us — if you have any notion that the creativity has sort of stalled in any kind of way, nothing puts fire on it by giving you a new capability, and setting you loose to try and figure out how to use it.

[0:52:07.0]

SPEAKER: Thank you. Thank you, Doug. I'm one of the Toulmin composers here, commissioning composers. I love what you're saying about — on my other side of the creative work, I'm also a creativity scholar. I've spent 11 years investigating what the nature of extraordinary artistic creativity and the AI evolution and revolution also make me excited. Because of the learning curve for composers, we're really looking at that as a reality. The learning curve for composers to be able to create extraordinary symphonic works for orchestras across the nation is much more likely with the AI technology.

But I'll just share maybe one thing that really struck me in 11 years of investigating extraordinary artistic creativity, is that we — the fact we're human beings is the fact that we're biological and not architectural. And I think putting that into a framework that we are biological, we're already seeing the fact that AI technology, it will weed out a lot of the average really quickly. And that is a great opportunity for artists and composers to think what makes something that is so human based that is irreplaceable.

And one more thing. Let me take out my notes. Composing as a creative practice, and the best of us composers, need something more than simply the right notes, or the knowledge of the orchestration. There is something that is ineffable, that makes me want to write a symphony, and this is about feelings between us that is ineffable, that cannot be put into words. I just want to be a reminder today for everyone here that we composers work every day, we labor our inspirations and feelings into our symphonic works for you and to connect with your audiences. All of that, it sprang from feelings, and how much we want to connect with everyone who's putting on good shoes to come to your orchestra performances, that they want to feel elated, they want to feel good as human beings to be part of this great art form that we all believe in.

[0:54:51.6]

DOUGLAS: So thank you. Yeah, you're exactly right. There's actually a whole — I just did a talk in Los Angeles about the idea of all of this technology and the speed-up and all of this. I actually think that it's an opportunity for us to take back and reinforce the humanity of what it means to be an artist, what it means to be part of a community, what it means to interact with one another. I kind of look at where we are now as some sort of, you know, I'm sorry, hellscape of content that is artificially fed to us, that — that is not good for us.

And the decay of all of those social network platforms and the possibilities of AI being able to grab back those things that I think are the essence of the human experience, the individual experience. I think we all have that need still to want to gather together, to be together, and to interact together. And so I don't see this as a replacement for any of that. I see it as the ability to take back that and make it more central, more primary to the things that we're trying to accomplish.

SPEAKER: But Doug, this question is really interesting to me, because it just begs another one, which is terrifying, which is, you know, she talked about the one thing that we have, which is unique, which is that we're biological and we have emotion. But can we imagine a world in the future when AI can replicate that? Where AI can replicate emotion? I mean, that's — that's the question that, for me, it begs. So what's your answer to that?

DOUGLAS: Yeah, so — so I use the metaphor of the slow food movement, okay? So in the '60s when fast food came out, and suddenly you could predict what you were going to get when you went into a McDonald's, or a Kentucky Fried Chicken, or whatever. And people just went nuts for it. It grew like crazy. But after a while, people started to realize, oh wait, it's not very good for you. Doesn't taste all that good. It's familiar, that's good. But you know, there's something sort of soulless about it that doesn't nourish us, right?

520 8th Avenue, Suite 2005, New York, NY 10018

[americanorchestras.org](http://americanorchestras.org)

[0:57:14.7]

And so that began the slow food movement. And the whole idea of slow food is, it's important about the culture around eating. You want to know the provenance of how the food was grown, where it was grown, who did it. The art of getting together and sharing as part of the meal and everything. And now, slow food has become a movement that sort of emulates in all sorts of other spheres, right? That's that need to create culture around it. So when you say an AI — you know, I use the example of Hemingway stories because they've actually already done that, and they're actually pretty damn good.

[0:57:58.8]

I think there's going to be endless supply of that stuff. But there already is endless supply of really crappy art out there, or art that's sort of on purpose, and that is just sort of wallpaper, you know? When you can replicate the Mona Lisa in a print in a jillion different ways, it ceases to have any kind of meaning to you in a way. But we still value that thing that is the — the one-to-one interaction. And if I used AI to create that, that's still input. I don't know that that completely answers it, because I don't know that there is an answer right at the moment. But that's my belief in what the essential part of being an artist and that art experience in a shared space is.

SPEAKER: Next question. Bill Neri.

BILL: Hi, everybody. Bill Neri from the Sphinx Organization. I just had a really quick question for the room actually, if I may. I just want to do a temperature check on the industry. Who here runs or is part of an organization that is currently leveraging AI tools as a centerpiece of their workflow or product? Thank you.

DOUGLAS: That's more than I expected. Interesting.

SPEAKER: Okay, next question. Over here.

EMMA: Thanks. Hi, I'm Emma from the Erie Philharmonic. My question is, how do you like combat or how do you like to be proactive with DEI initiatives and AI? Because I know there are some biases currently with the various AI platforms. How are you proactive in ensuring that your content and your output has a positive impact with your DEI initiatives, and making sure that there's no prejudice?

520 8th Avenue, Suite 2005, New York, NY 10018

[americanorchestras.org](http://americanorchestras.org)

DOUGLAS: Ah, okay, so that's — that's actually a really good question.

SPEAKER: Also, if I could just add to that, not creating harm is another — another words I would use there. Thank you for that question.

[0:59:55.0]

DOUGLAS: Yeah, part of the problem when you train AIs on the entire internet is that you get all — you scoop up all the bad stuff as well. And so that — that pizza glue thing, that was from a post on Reddit, and somebody was making a joke, and it didn't understand, the AI didn't understand that it was a joke, and so it took it seriously, right? There's also sort of embedded cultural prejudices in the internet, obviously. But also the way some of these things are trained.

And so for instance, I think it was — might have been Google who took its AI down for a little bit, because they asked for a picture of the founding fathers, and it was all people of color, right? There's nothing wrong with that, but it's not accurate. And so there is need to be concerned about that. The way AI is starting to evolve is that you have these massive, large language models that live. There — there are five like really super big ones right now. And they take about 200 million, 300 million dollars a year just to train and operate, and the electrical costs and all of that. So only the biggest companies can afford to do that.

But what's also being developed now are local language models in which you train it on your own data, where you — all of your data stays within your local computer. But it trains its abilities based on one or more of these large models. So you have — you have the opportunity to be able to train your AI in a way that is consistent or congruent with your own values, with your own data, and all of that. So, and again, think exponentially. This is going to be — you know, five months from now is going to be an entirely different problem around that. But yeah, you're right. I mean, there have been any number of problems of thinking about what the — what the equity in all of this is.

SPEAKER: Yeah.

KOVEN SMITH: Hi, Koven Smith from the Knight Foundation. I wonder if you could talk a little bit about the generational problem. I'm thinking of Ted Chang's article, that ChatGPT is a blurry — blurry JPEG of the web, saying that as these models start to incorporate their own output as input for the next generation of training data, what — you know, what effect that is going to end up having. And I think a lot of the discourse there has been more on kind of the text-based models, less so on, you

520 8th Avenue, Suite 2005, New York, NY 10018

[americanorchestras.org](http://americanorchestras.org)

know, visual and auditory art. So I'd just like to hear your — your thoughts on where that might be going.

DOUGLAS: Thank you. Unbelievably, we've kind of run out of data, right? The internet has just so much data on it, but they've now crawled everything that they have access to. And so the next question is, where do we get the next set of data, right? So one of the first things that they started doing was feeding it synthetic data, by which the output of the AIs, it now adds to the internet stuff that they've scraped, and then the new AIs train on the synthetic AI, right? And what they've discovered is that the hallucinations just get to be extraordinary.

[1:03:49.1]

It's like — they actually call it pollution, right? The data sets become polluted. And — and so there's been a back off of that. So where is the next data going to come from? It's going to come from our phones, from our sensors, from motion sensors, from — in the real world, so that we'll be able to interact with the real world in real time. And we'll be able to create things that will, if this happens here, make this physical action happen over there, right? So that's — that — if you think about the amount of data out there, the amount of sensors out there, the amount of data that we're — we're able to now even collect, it exponentially increases the amount of data.

But that still doesn't answer the question of, how do you make the data that we've already scraped and get rid of that synthetic data so that it becomes more in focus and isn't just the blurry jpeg of that. And I think they're working on that. But right now, I — you know, I don't hear a lot of people even understanding all the issues. So until they start to get a little more into focus, we'll start addressing those things. It's just we haven't got there yet.

SPEAKER: One last question. I'm going to go over here. And then we'll wrap up.

MAARTEN: Hi, Maarten from Symphony.Live. I also write a newsletter called MusicX about music and tech. Thanks, Doug. I just wanted to sort of highlight a few things. It was a good introduction. But there had seemed to me to be a little disconnect between the way that you talk about sort of the big tech companies, Meta, and the fact that, you know, you then quote Sam Altman without also sort of making that into a problem. Because OpenAI is probably worse than any of the big tech companies, right?

And you talked about Udio. But Udio is trained on copyrighted data. They claim it's not the case. And they — if you read the stuff that they come out with, it's all based on fair use, and it's bullshit, yeah?  
520 8th Avenue, Suite 2005, New York, NY 10018

[americanorchestras.org](http://americanorchestras.org)

So if you use Suno, if you use Udio, you're — you're working on copyrighted data. And we should make a big, big, big problem out of that. And there's a lot of initiatives out there. So you mentioned Holly. Holly's great. If you don't know Holly Herndon, check her out, her music is fantastic, her thinking is — is good.

[1:06:29.4]

Spawning is coming out with a model where you can transparently see what data it has been trained on, right? So go check that out. Another example is Fairly Trained, which is US-based initiatives from — at Newton Rex. Check that out. They're making a certificate to show that something is ethically transparently trained data, yeah? There's another one from Ireland, which is AI:OK. Also follow along with that. Those are the kinds of things that I would encourage everyone in the room to kind of follow along with, right? And on top of that, do play around with all of this stuff, because that's the way you understand how it kind of works and what it can do for you. So just wanted to add that.

DOUGLAS: Yeah, no, I think that's a real point. I mean, for anything that I said here today, I can make a counter argument to it that it's, oh, this is — this is not good, right? And we haven't worked out what is fair play. I guess I should have turned off my phone. [LAUGHING] Or maybe it's the AI calling me.

SPEAKER: Well, you can go and answer it because we're going to wrap up now.

DOUGLAS: Okay. Yeah, so — so, you know, I don't mean to be definitive here. I've skipped over an enormous amount of information. I — you know, you have to balance sort of like how much can you give over every issue you want to go deeper into. So yeah, all of these issues, you're absolutely right. The ethically trained AIs are a real issue. But we don't yet even really understand what ethically trained means, right? Licensing by the New York Times or Wall Street Journal. Wall Street Journal just made a \$250 million licensing deal with OpenAI to license its — its stories. But what does that actually mean? If website traffic completely goes away, does that \$250 million over five years replace that revenue? Not even slightly. So it's all very complicated.

[1:08:31.9]

SPEAKER: Well, the conclusion is we — we know less than we need to, but maybe we know a little bit more than we did an hour and fifteen minutes ago. So for that, we thank you Doug McLennan. Thank you very much.

520 8th Avenue, Suite 2005, New York, NY 10018

[americanorchestras.org](http://americanorchestras.org)

DOUGLAS: Thank you.

[APPLAUSE]

### END OF TRANSCRIPT ###