

# The Economist

---

Why Hispanics like Trump

---

China, Iran and Russia v the West

---

Don't curb abortion pills

---

AI music: algorithm and blues

---

MARCH 23RD-29TH 2024

---

# ISRAEL ALONE



## Science & technology



### AI and music

## Algorithm and blues

### Artificial intelligence promises to transform music-making—for better and worse

**I**N THE DYSTOPIA of George Orwell's novel "1984", Big Brother numbs the masses with the help of a "versificator", a machine designed to automatically generate the lyrics to popular tunes, thereby ridding society of human creativity. Today, numerous artificial-intelligence (AI) models churn out, some free of charge, the music itself. Unsurprisingly, many fear a world flooded with generic and emotionally barren tunes, with human musicians edged out in the process. Yet there are brighter signs, too, that AI may well drive a boom in musical creativity.

AI music-making is nothing new. The first, so-called "rules-based", models date to the 1950s. These were built by painstakingly translating principles of music theory into algorithmic instructions and probability tables to determine note and chord progressions. The outputs were musically sound but creatively limited. Ed Newton-

Rex, an industry veteran who designed one such model for Jukedeck, a London firm he founded in 2012, describes that approach as good for the day but irrelevant now.

The clearest demonstration that times have changed came in August 2023. That is when Meta, a social-media giant, released the source code for AudioCraft, a suite of large "generative" music models built using machine learning. AI outfits worldwide promptly set about using Meta's software to train new music generators, many with additional code folded in. One AudioCraft model, MusicGen, analysed patterns in

some 400,000 recordings with a collective duration of almost 28 months to come up with 3.3bn "parameters", or variables, that enables the algorithm to generate patterns of sounds in response to prompts. The space this creates for genuinely new AI compositions is unprecedented.

Such models are also getting easier to use. In September Stability AI, a firm based in London at which Mr Newton-Rex worked until recently, released a model, Stable Audio, trained on some 800,000 tracks. Users guide it by entering text and audio clips. This makes it easy to upload, say, a guitar solo and have it recomposed in jazzy piano, perhaps with a vinyl playback feel. Audio prompts are a big deal for two reasons, says Oliver Bown of Australia's University of New South Wales. First, even skilled musicians struggle to put music into words. Second, because most musical training data are only cursorily tagged, even a large model may not understand a request for, say, a four-bar bridge in ragtime progression (the style familiar from Scott Joplin's "The Entertainer").

The potential, clearly, is vast. But many in the industry remain sceptical. One widespread sentiment is that AI will never produce true music. That's because, as a musician friend recently told Yossef Adi, an engineer at Meta's AI lab in Tel Aviv, "no

#### → ALSO IN THIS SECTION

**73 Can AI improve football tactics?**

**74 Starship's latest launch**

**75 Why XL bullies are so aggressive**



▶ one broke its heart". That may be true, but some AI firms reckon that they have found a way to retain and reproduce the "unique musical fingerprint" of their musician users, as LifeScore, a company founded near London, puts it. LifeScore's AI limits itself to recomposing the elements of a user's original recordings in ways that maintain the music's feel, rather than turning them into something radically new.

It takes about a day to plug into LifeScore's model the dozens of individually recorded vocal and instrumental microphone tracks, or stems, that go into producing an original song. Once that's done, however, the software, developed at a cost of some \$10m, can rework each stem into a new tempo, key or genre within a couple of seconds. The song's artists, present during the process, choose which remixes to keep. Manually remixing a hit track has traditionally taken one or more highly paid specialists weeks.

LifeScore, says Tom Gruber, a co-founder, is "literally swamped with requests" from clients including Sony Music, Universal Music Group and Warner Music Group. An original release is typically turned into anywhere from a handful to a dozen remixes. But one client aims to release a dizzying 6,000 or so AI versions of an original track, each targeting a different market. Artists including Pink Floyd's David Gilmour and Tom Gaebel, a German pop singer, use LifeScore's AI to power websites that allow fans to generate, with a few clicks, new remixes adapted to personal tastes.

### The beat of a different drum

If this seems like dizzying progress, it's worth noting that AI's impact on music is still in its early days. Legal uncertainties over the use of copyrighted recordings to train models have slowed development. Outfits that have coughed up for licensing fees note that this can get expensive. To save on that cost, MusicGen's training set mostly sidestepped hits, says Dr Adi. Though output is pretty good, he adds, the model is not yet "artistic enough" to generate narratively complete songs. Harmonic misalignments are common. OpenAI, a San Francisco firm, for its part, says its MuseNet model struggles to pull off "odd pairings", such as a Chopin style that incorporates bass and drums.

In time, bigger training sets of better music will largely overcome such shortcomings, developers reckon. A Stability AI spokesperson says that while Stable Audio's top duration for coherently structured music—"intro, development and outro"—is now about 90 seconds, upgrades will produce longer pieces with "full musicality". But judging music AI by its ability to crank out polished tracks mostly misses the point. The technology's greatest promise, for now at least, lies elsewhere.

Part of it is the empowerment of amateurs. AI handles technical tasks beyond many people's capabilities and means. As a result, AI is drawing legions of newbies into music-making. This is a boon for experimentation by what Simon Cross, head of products at Native Instruments, a firm based in Berlin, calls "bedroom producers".

Consider RX, a Native Instruments AI "assistant" that corrects errors in things

like pitch and timing. For the latter, software time-shifts notes by cutting out or inserting slivers of sound with matching timbre, a process called "dynamic time-warping". The company's AI also determines what mixing and mastering processes were performed on a song of a user's choosing. It then replicates, or at least approximates, the same expensive processing on the user's own creations. Boomy, an online "mu- ▶▶

### AI and football

## Back of the neural net

### AI can improve corner-kick tactics. Coaches should pay attention

**T**O ZHE WANG at Google DeepMind, an artificial-intelligence (AI) company, corner kicks are like games of chess. Partly because both feature opposing sides poised to react to a single imminent move. But also, no doubt, because they too may be revolutionised by AI.

AI models thrive where there is abundant data. Football more than satisfies this requirement. Elite players wear vests that measure heart rate, position, speed and force exerted; team analysts watch hours of footage to tally possession percentages and numbers of passes, shots and goals. In a paper published on March 19th in *Nature Communications*, Mr Wang and his colleagues worked with staff at Liverpool Football Club to feed this data into a statistical model known as a graph neural network (GNN). They then were able to use this model to predict which on-field player would wind up making first contact with the

ball with levels of accuracy similar to human experts.

Coaches around the world should be pacing in front of their dugouts in excitement at the news. During a football match's standard 90 minutes, it is rare to encounter the same situation twice. But corner kicks, which are eminently repeatable, are the focus of hours of specialist training. And doing well in them pays off. Arsenal, sitting at the top of the Premier League when this research paper was published, have scored 13 goals from corners out of a season total of 70. Small advantages in this part of the game can make a big difference.

DeepMind's GNN works by establishing statistical relationships between data collected from all 22 players during 7,176 corner kicks. Each corner routine was eventually represented as a vector in a 352-dimensional space: impossible for a human to visualise, but easy for a computer to process. Corners that unfolded similarly wound up close together in that space, allowing the model to make predictions about new set pieces.

The model was also capable of using this analysis to suggest new tactics. And, based on ratings given by five experts from Liverpool, AI-generated corner-kick tactics were just as good as those suggested by human coaches. In fact, shown 50 pairs of corners, one real and one with an AI-suggested improvement, 90% of the AI suggestions found favour with the majority of judges.

Mr Wang, who confesses he is "no football fan", says that the sport offers a safe and controllable test-bed to develop helpful AI technologies that might one day be used in health care or defence. After all, football is not a matter of life and death. It is, as Bill Shankly, a former Liverpool manager, once observed, much more important than that.



Whip it into the black box

► sic automation” platform for what Alex Mitchell, its CEO, describes as “low-friction” song production with text prompts, has more than 2m users. The company, based in Berkeley, California, uploads users’ (vetted) creations to streaming services and collects a cut of revenues.

AI serves professionals, too. The soundtracks to “Barbie” and “Oppenheimer” were cleaned up in post-production with RX, for example. Another application area is “style transfer”, in which models transform music recorded with one instrument into sounds that seem to come from a different one, often with a twist or two requested by the user. Style transfers are also used for voice. A model developed by a startup in London called Voice-Swap slices up sounds sung by (remunerated) pro-

fessional singers and rearranges the slivers into lyrics written by the service’s users, who pay licensing fees for the rights to sell the resulting tracks. And AI tools already exist to recreate singers’ voices in other languages. Vocaloid, a voice-synthesising tool from Yamaha, a Japanese instrument manufacturer, is one of many that can use a translation sung by a native speaker as a template for an AI to imitate as it rearranges, modifies and stitches together tiny snippets of the original singer’s voice.

Accomplished musicians now widely tap MusicGen and its competitors as sources of “infinite inspirations”, the better to alight upon promising composition ideas, says Meta’s Dr Adi. Whether such inspiration pays off will, ultimately, be up to the listener to decide. ■

cally. During the second test flight, in November, the Starship upper stage successfully separated from the booster around three minutes into the flight, and continued to fly for another eight minutes, reaching an altitude of 149km. But both vehicles then suffered failures that led to their rapid unscheduled disassembly.

The third flight, once again from SpaceX’s Starbase near Boca Chica in Texas, also led to the loss of both vehicles. But this time the upper stage reached orbit, at an altitude of 230km. SpaceX’s aim is that both vehicles should eventually be fully reusable, flying back to Earth to land, as its Falcon 9 boosters already do. For these initial test flights, however, its plan has been for the Super Heavy booster to practise a “soft” splashdown at sea, slowing its descent using rocket engines before slipping beneath the waves, while the Starship swoops halfway around the Earth, re-enters the atmosphere and makes a “hard” splashdown (ie, crashes into the sea from a great height).

Yet even these more limited goals have so far proved overambitious. During the test flight in November, the Super Heavy booster exploded when it attempted to reignite its engines in preparation for splashdown. This time round, it successfully completed its “boostback burn”, which carried it to its intended splashdown point in the Gulf of Mexico. But when it reignited its engines for its landing burn, to slow its descent velocity to almost zero, something went wrong, causing the vehicle to explode a few hundred metres above the water.

The Starship, however, reached a low-Earth orbit that took it across southern Africa. While in this so-called coast phase, it carried out a series of test operations: opening and closing its payload door (Starship will ultimately be able to carry as much as 150 tonnes of payload into orbit); transferring propellant between two on-board tanks (a warm-up for future Starship-to-Starship refuelling in orbit, a requirement for missions to the Moon and Mars); and, finally, attempting a controlled re-entry. On-board cameras showed the glowing pink plasma expected of a spacecraft re-entering the atmosphere.

But contact was lost soon afterwards, and SpaceX announced that the vehicle had disintegrated. (Whether this was because of a failure of its heat-shielding tiles, loss of control of the Starship or some other reason remains unclear.)

SpaceX has been mocked for the tendency of its test flights to end in expensive firework shows, but such criticism misunderstands the company’s approach to rocket development. It prefers to iterate quickly and learn from failure, rather than painstakingly plan for a perfect launch—which, in the rocket industry, almost never happens (as the loss of the Kairos rocket ►►

## Space flight


# Fireworks display

## After two failed attempts, Elon Musk’s Starship finally made it into orbit

THE WORLD’S largest rocket has flown again, and its uncrewed test flight on March 14th, like the two previous ones, ended in “rapid unscheduled disassembly” (ie, catastrophic explosions). But the upper stage of Starship, built by SpaceX, Elon Musk’s rocket company, reached orbit for the first time and completed several test operations before being destroyed while re-entering Earth’s atmosphere. In a defiant post on X, the social network formerly known as Twitter, Mr Musk insisted that “Starship will make life multiplanetary.”

Perhaps. This latest flight was at least another small step in that direction.

Starship is in fact two vehicles: the Super Heavy booster stage, a behemoth 71 metres tall with 33 engines, and the 50-metre Starship upper stage, which separates from the booster at an altitude of around 70km (44 miles) to proceed to orbit. The combined vehicle’s first test flight, in April 2023, ended roughly four minutes after lift-off, with the self-destruction of the entire spacecraft, after stage-separation failed and the rocket began to corkscrew errati-



What goes up must undergo rapid unscheduled disassembly